UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

Facing an Altered Future: Essays on the Economic Impacts of Climate Change and Adaptation

Permalink

https://escholarship.org/uc/item/0ss9b490

Author

Hultgren, Andrew

Publication Date

2020

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Facing an Altered Future:
Essays on the Economic Impacts of Climate Change and Adaptation

By

Andrew Bruce Hultgren

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Maximilian Auffhammer, Co-chair
Professor Solomon Hsiang, Co-chair
Professor James Sallee
Professor David Anthoff

Summer 2020

Copyright 2020

Andrew B. Hultgren

All rights reserved

Abstract

Facing an Altered Future:
Essays on the Economic Impacts of Climate Change and Adaptation

by

Andrew B. Hultgren

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Maximilian Auffhammer, Co-chair
Professor Solomon Hsiang, Co-chair

In this dissertation, the reader will find four essays examining the empirical effects of future climate change, while accounting for both adaptation as well as its costs. In three of these essays, we uncover climate impacts that result in a substantially altered future from our present reality: crop yields in sharp decline (while accounting for costly adaptation); cropped areas shifted away from present-day locations; and increased human mortality, with costly adaptations undertaken to avoid even further excess deaths. In one essay, we estimate that the impact of climate change on global energy consumption will be small in aggregate due to offsetting effects from more hot days but fewer cold days, though regional disparities may be important.

Several conclusions follow. First, accounting for both adaptation and its costs is important for grounding an accurate estimate of the social cost of carbon. Second, performing these analyses at a global scale and with high spatial resolution reveals important sources of heterogeneity – both in impacts as well as in the magnitude of adaptation. Third, results such as these may be useful inputs to the Integrated Assessment Modeling community, improving the empirical grounding of the damage functions employed.

In these essays, we seek to include information from as globally representative a sample as possible, and where possible estimate heterogeneous effects for the global rich versus the global poor. We also seek to graphically illustrate projected effects so that non-expert readers can locate themselves within the empirical estimates we produce.

Acknowledgments

Kira Dominguez-Hultgren: thank you for being my advisor, sounding board, support system, critic, comedy relief, and lifelong friend and partner. I love you. Also, one day, I will beat you at Agricola.

Pace and Satya Dominguez-Hultgren: thank you for keeping me grounded in the wider world beyond academia. Thank you for being board game and movie partners, amazing and entertaining performers, willing learners, and for putting up with – and even loving – your sometimes distracted parents. I love you both, and am so proud of you.

My co-chairs, Max Auffhammer and Sol Hsiang: You opened my eyes to an entire world of research and have been ready partners in its exploration, rather than letting me hesitate at the threshold. You have been an incredible advising team; I am deeply indebted to you both.

Jim Sallee: Thank you for believing in me, pushing me, and walking with me into the unknown. You made my advising trio, and I am deeply indebted to you, as well.

My cohort: You are the BEST cohort anyone could ask for. Period. Thank you for being my supporters and constructive critics, and for tons of laughter! I will always look forward to our next beer, and our next paper.

Thank you to the members of the Climate Impact Lab and Global Policy Lab. What a dynamic, engaging, and amazing group of people to work and play with! Tamma Carleton in particular: you have been an incredible advisor, example, and teammate. Thank you.

Thank you to the ARE, Econ, EI, and GSPP faculty for your instruction, critiques, and kindnesses. You have formed me in ways I simply could not have imagined possible. In particular, thanks to Michael Anderson, Severin Borenstein, Lucas Davis, Alain De Janvry, Meredith Fowlie, Larry Karp, Jeff Perloff, David Romer, Betty Sadoulet, Leo Simon, Reed Walker, and Brian Wright. Particular thanks to Steve Raphael for early encouragement, training, and advice while I was an MPP student; to Michael Greenstone for expecting excellence and teaching me to (hopefully) expect it of myself; and to David Anthoff, for thoughtful late-stage feedback.

Thank you to Carmen Karahalios and Diana Lazo for always being willing to help in a pinch, and for ready words of kindness and encouragement.

Finally, I thank my family (on both sides) and God for patience, grace, and love, no matter my failings. Special thanks to my parents Bruce and Trice Hultgren, for cultivating a sense of curiosity from a young age.

Dissertation Committee

Maximilian Auffhammer (co-chair)
University of California, Berkeley
Department of Agricultural and Resource Economics
Energy Institute

Solomon Hsiang (co-chair)
University of California, Berkeley
Goldman School of Public Policy
Global Policy Lab

James Sallee
University of California, Berkeley
Department of Agricultural and Resource Economics
Energy Institute

David Anthoff
University of California, Berkeley
Energy and Resources Group
Society, Environment, and Economics Lab

Chapter 1: Introduction

Future climate change is a problem with the potential for civilization-altering consequences: end-of-century sea level rise under a high emissions scenario could leave over 600 million people at risk of coastal flooding, based only on present-day population (Kulp and Strauss, 2019); weather-induced spikes in the price of wheat may already have played a key role in the "Arab Spring" regional destabilization (Johnstone and Mazo, 2011); and carbon-cycle feedbacks could irreversibly push the planet into a "hothouse" pathway of 4-5°C warming and 10-60 meters of sea level rise – a climate not seen since the Mid-Miocene period 15 million years ago (Steffen et al., 2018). Clearly, understanding the economic impacts of climate change and appropriately pricing the greenhouse gas externality is critically important to the welfare of future generations (Interagency Working Group on Socal Cost of Carbon, 2010; National Academies of Sciences, Engineering, and Medicine, 2017).

The last decade has seen growth in the literature estimating the effects of weather shocks on a range of economic and social outcomes (Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Feng, Krueger, and Oppenheimer, 2010; Auffhammer and Aroonruengsawat, 2011; Deschênes and Greenstone, 2011; Hsiang, Meng, and Cane, 2011; Dell, Jones, and Olken, 2012; Graff Zivin and Neidell, 2014). However, understanding how agents might adapt to slow-moving changes in climate (Mendelsohn, Nordhaus, and Shaw, 1994) has proven to be difficult (Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016). Importantly, agents may be able to take a range of actions to mitigate the harm of a slow-moving change in temperature that would be unavailable to them in the context of a short-run weather shock (Mendelsohn, Nordhaus, and Shaw, 1994; Hsiang, 2016). If these adaptive actions effectively mitigate the effects of extreme heat on social outcomes of interest, then empirical estimates derived from short-run weather shocks may overstate the economic losses associated with future climate change. However, adaptive actions are not costless (Schlenker, Roberts, and Lobell, 2013); if they were, then rational agents would already have undertaken them. Understanding the capacity for adaptation to climatic changes as well as the costs of adaptive actions is, therefore, of first-order importance in determining optimal climate policy today.

This dissertation comprises a body of work focused on empirically estimating the scope for human adaptation to slow-moving climatic changes, including both the magnitude of mitigated impacts afforded by adaptation as well as the implied costs of adaptive actions, at a global scale. Three chapters focus on estimating climate impacts, net of adaptation and its costs, on three important aspects of human society: global grain yields, global energy consumption, and non-market-valued human mortality. An additional chapter on agriculture and climate focuses on an important but poorly understood potential mechanism for adaptation: the poleward shifting of cropped areas to follow optimal growing climates.

Across these studies, two key patterns emerge. First, adaptation is, in general, important. Confirming the intuition of Mendelsohn, Nordhaus, and Shaw (1994), we find that adaptation mitigates the effects of weather shocks on grain yields and on human mortality, and that cropped areas do shift with a shifting climate. In addition, we find that energy consumption helps enable adaptation, with moderately increased hot-day demand for electricity in rela-

tively hot places, and substantially increased cold-day demand for non-electricity (heating) fuels in relatively cold places. Second, adaptation is costly: substantial losses due to climate change – net of adaptation and its costs – persist.

The key findings in each study are briefly summarized below.

1.1 Chapter Summaries

The Impacts of Climate Change on Global Grain Production Accounting for Adaptation

This study presents the first granular and globally comprehensive analysis of crop yields under climate change, accounting for farmer adaptation and its costs along the critical dimensions of slow-moving changes to temperature and precipitation, income growth, and irrigation. To do this, we assemble once of the largest datasets of subnational crop production available to researchers, covering over 10,000 administrative units from 55 countries and representing two thirds of global crop caloric production.

Across crops, we find consistent patterns of adaptation to climatic changes and through growth in incomes and access to irrigation. However, we also find that globally widespread yield declines persist, in spite of these adaptive efforts. We show this by adapting simple machine learning techniques to a causal inference context, enabling us to for the first time systematically assess a rich set of channels through which temperature and precipitation shocks might causally affect yields in each of our crops, while also systematically assessing the extent to which farmers might use income growth, irrigation, or expectations of long-run temperature and precipitation to mitigate the extent of these yield losses.

Projecting our empirical estimates into the future, we find that even after accounting for costly adaptation, RCP 8.5 end-of-century yield losses are projected to be severe in the key cropped areas of the U.S. grain belt and Eastern China (40% or greater losses), important in the major cropped areas of South America and Australia (10% losses), and mixed in Europe, Sub-Saharan Africa, India, and Southeast Asia.

Overall, while costly adaptation is protective for the major grain producing regions of the world relative to no adaptation at all, yield impacts in spite of adaptation are still generally negative and for many regions of the world they are substantially so.

Climate Change and Crop Choice: Evidence from the United States

This study examines how farmers adjust the land they allocate to crops when climatic conditions change. We analyze subnational longitudinal data for two globally important cereals from one of the world's major agricultural regions, the United States. Consistent with theory, we find that more land is allocated to producing a specific crop when the climatological conditions for growing that crop improve, although the timescales for these adjustments are long, on the order of 20 or more years. When we estimate the effectiveness of land use change as an adaptive strategy, we find that these adjustments reclaim 5-10% of climate-induced losses to producer surplus. The availability of adaptation via land use fails to generate more value because enduring climate-induced yield losses is less costly than altering land use practices for most farmers, indicating that land use adjustment costs may

be high.

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

This paper develops the first globally comprehensive and empirically grounded estimates of mortality risk due to future temperature increases caused by climate change. Using 40 countries' subnational data, we estimate age-specific mortality-temperature relationships that enable both extrapolation to countries without data and projection into future years while accounting for adaptation. We uncover a U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for the elderly, that is flattened by both higher incomes and adaptation to local climate (e.g., robust heating systems in cold climates and cooling systems in hot climates). Further, we develop a revealed preference approach to recover unobserved adaptation costs. We combine these components with 33 high-resolution climate simulations that together capture scientific uncertainty about the degree of future temperature change. Under a high emissions scenario, we estimate the mean increase in mortality risk is valued at roughly 3.2% of global GDP in 2100, with today's cold locations benefiting and damages being especially large in today's poor and/or hot locations. Finally, we estimate that the release of an additional ton of CO2 today will cause mean [interquartile range] damages of $36.6 [-$7.8, $73.0] under a high emissions scenario and $17.1 [-$24.7, $53.6] under a moderate scenario, using a 2% discount rate that is justified by US Treasury rates over the last two decades. Globally, these empirically grounded estimates substantially exceed the previous literature's estimates that lacked similar empirical grounding, suggesting that revision of the estimated economic damage from climate change is warranted.

Estimating a Social Cost of Carbon for Global Energy Consumption

The global marginal damage caused by emitting a single ton of carbon dioxide (CO_2), or its equivalent, is key to climate policy,Nordhaus (1992); Interagency Working Group on Socal Cost of Carbon (2010); National Academies of Sciences, Engineering, and Medicine (2017) but our current understanding of its value is based on spatially-coarse theoretical-numerical modelsStern (2006); Tol (2009); Pindyck (2013) that are not tightly linked to dataNational Academies of Sciences, Engineering, and Medicine (2017); Burke et al. (2016). We develop the first architecture that integrates best-available data, econometrics, and climate science to estimate climate damages worldwide at the local level, as well as aggregated global marginal damages. Here we apply this architecture to construct the first global empirical estimates of the impact of climate change on total non-transport end-use energy consumption, one of the most uncertain impacts in current models.Anthoff and Tol (2013) At end-of-century, we project annual global electricity consumption to rise roughly 4 EJ (1100 TWh, 6% of current global consumption) for each 1°C increase in global mean temperature, reflecting increased cooling demand, while direct consumption of other fuels declines 10.1 EJ (6% of current global consumption) per 1°C, reflecting reduced heating. Together, these estimates indicate that emission of 1 ton of CO_2 today produces global net savings in future aggregate energy consumption of about $1 in net present value (3% discount rate). This finding is largely driven by a sharply nonlinear relationship between income and temperature-induced energy consumption, which indicates that for most of the 21^{st} century, much of the world is expected

to remain too poor to increase energy consumption in response to warmer temperatures. By end-of-century, emerging economies in the tropics (e.g. India) are projected to increase electricity consumption dramatically, but these rising costs are offset by heating reductions in the wealthy economies of North America and Europe.

1.2 Conclusions and Directions for Future Research

Relying on the Ricardian approach, the seminal work of Mendelsohn, Nordhaus, and Shaw (1994) found that climate impacts on U.S. agricultural profits, net of adaptation, would likely only be modestly negative and possibly beneficial. Using fixed effects estimators more robust to omitted variables bias (Deschênes and Greenstone, 2007), the subsequent literature has found negative climate impacts across a range of outcomes, though estimates are generally unable to account for adaptation (Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016).

Similar to Auffhammer (2018), the studies presented here combine elements of both the Ricardian and within-estimator approaches: we use cross-sectional variation in long-run climate to uncover heterogeneity in the response to short-run (within-estimator) weather shocks. In so doing, we find that adaptation does indeed mitigate the effects of weather relative to the effect of an unanticipated shock.

However, for the first time we also uncover the implied costs of these adaptive actions. In contrast to Mendelsohn, Nordhaus, and Shaw 1994, we find that accounting for both adaptation and its costs, negative effects of climate change persist and are widespread. Severe ($\sim$40%) yield losses persist for major grain growing regions of the world; the costly migration of cropped areas is substantial, but only offsets $\sim$10% of producer losses; human mortality effects persist and are valued at 3.2% of end-of-century global GDP. Only in the case of global energy consumption do we find gains from climate change – in this case increases in the consumption of electricity on hot days as locations become hotter on average is offset by the decreased consumption of heating fuels as cold places both warm and experience fewer cold days.

Several conclusions follow. First, accounting for both adaptation and its costs is clearly important for grounding an accurate estimate of the social cost of carbon (SCC). Second, performing these analyses at a global scale and with high spatial resolution reveals important sources of heterogeneity – both in impacts as well as in the magnitude of adaptation – that regionally focused or geographically coarse studies might miss. Third, results such as these may be useful inputs to the Integrated Assessment Modeling community (Nordhaus, 1992; Tol, 1997; Hope, 2002), who may be able to incorporate estimates like these into their damage functions, improving their empirical grounding.

A shortcoming of the approaches applied in this dissertation is that the variation used to estimate adaptive responses is typically cross-sectional in nature. That is, we use cross-sectional variation in long-run climate to parameterize the heterogeneity in the short-run response function. If omitted variables are correlated with heterogeneity in the short-run response function and the outcome of interest, these omitted variables could be a source of bias in our estimates of the magnitude of adaptation and its implied costs.

Future research can take a number of different directions. Valuation of the grain yield losses presented here will be important for their incorporation into an estimate of a partial-SCC. Additional economic outcomes can be investigated, such as labor supply, human migration, or vector-borne diseases. Risk preferences can be incorporated to reflect curvature in the social welfare function of a planner, or in the utility function of a representative consumer. Finally, following Burke and Emerick (2016); Shrader (2018); Hagerty (2020), natural experiments over dimensions of adaptation (or the information set that forms expectations) can be used to improve identification in our estimates of adaptation.

This concludes the summary of the dissertation research itself. Given the extraordinary nature of the moment in which I am completing my doctorate, I turn below to a brief reflection on a potential role for economics research in the present social context. This reflection is of a more editorial nature; the disinterested reader may, of course, disregard it.

Fiat Lux.[1]
– Motto, U.C. Berkeley

"If you do away with the yoke of oppression,
    with the pointing finger and malicious talk,
and if you spend yourselves in behalf of the hungry
    and satisfy the needs of the oppressed,
then your light will rise in the darkness,
    and your night will become like the noonday."
– Isaiah Chapter 58 (Israel, 8^{th} century BCE)

Reflection: Research during a period of social upheaval[2]

I am completing my doctorate during a time of turmoil without precedent in recent decades: the emergence of COVID-19 has led to widespread unemployment, economic upheaval, and a new leading cause of death; the murder of George Floyd and subsequent Black Lives Matter protests have brought into sharp relief the failing of U.S. society to confront ongoing systemic racism; and the Me Too movement has brought the same sharp focus on systemic sexism.

What, at a time like this, could possibly be the point of engaging in academic research?

Like many, I believe that research can be used effectively to improve policy and enhance social welfare. *Fiat Lux*, and may the knowledge it affords lead to better-informed decisions. However, the policy-making process takes many inputs; information is but one of them. Optimal policy may be compromised to obtain votes. Politicians may be concerned with how they will finance the coming political campaign. Regulators may be swayed by the persistence of a well-organized lobby. The currying of political influence with the objective of shaping policy to an individually favorable outcome is a long-standing human practice: in the chaos of the policy-making process, it is those with the means to protect and expand their private positions who frequently come out as winners. When hard choices need to be made, those without a voice at the table are distinctly disadvantaged.

This is where we as researchers can be of service: through our research, we can help give voice to the voiceless. Let me be clear, I am not advocating anything less than the clear-eyed, unswerving pursuit of truth. We should be our own harshest critics, and then look to our colleagues and peer review to test the soundness of our findings even further. However, as experts, we can do more to make our findings accessible to those who might benefit from them but who may lack the means to place those findings in their own contexts. In so doing, we can focus the beam of light created by newfound knowledge and illuminate particularly those parts of society that have been conveniently ignored in the policy-making process.

[1] "Let there be light."

[2] I thank Kira Dominguez-Hultgren and my dissertation committee for helpful comments on this reflection.

What might research that gives voice to the voiceless look like? I propose a non-exhaustive and certainly flawed set of characteristics, to which I hope others will add.[3]

Estimate heterogeneity over historically marginalized groups, where possible.

First, researchers might ask whether the question at hand exhibits heterogeneity over historically marginalized groups, and a sufficiently powered empirical analysis could estimate such heterogeneous treatment effects. For example, in Chapter 4 of this dissertation, Figure 10, we use our estimates of heterogeneous treatment effects to show how human deaths (Panel A) and mortality adaptation costs (Panel B) are differentially distributed over the global poor vs. the global rich. In particular, we find that today's wealthy countries are largely able to pay to adapt away the mortality effects of future climate change; however we project today's poor countries to largely bear the mortality effects of climate change through human deaths, rather than costly adaptation.

Create figures for non-expert readers.

Second, we might present our results in such a way that a non-expert reader can locate themselves within the data. Maps of intuitively accessible outcomes can be particularly effective tools in this regard. For example, in Chapter 2 of this dissertation, Figure 4, we display global maps of projected crop yield losses or gains (net of adaptation and its costs) under climate change. Such a map – especially if communicated through the media – makes it possible for policymakers and the public in Argentina or Malaysia or Tanzania to see how climate change is projected to affect the yields of key global grains in their home towns.

Train in science communication.

Third, we could treat science communication as an integral part of the research process, and explicitly train our doctoral students in it. Thoughtfully engaging the press with an accurate and well-honed message – including accessible graphics and illustrative analogies – can increase the chances that research findings will be heard and used accurately by those most affected by a given policy. The papers in this dissertation have not yet reached the publication stage, so I cannot speak to the science communication and media preparation process for them. However, a recent paper on the effectiveness of COVID-19 anti-contagion policies has been published, and our team spent many hours honing our ability to communicate the key findings of the paper,[4] including discussing word pictures, creating new graphics, and repeatedly giving mock-interviews to practice the accurate yet accessible communication of our key results.

[3]This discussion is inspired by many interactions, however two critical sources stand out. One is the science communication training incorporated into the NSF DS421 traineeship, including a particularly insightful talk given by Trevor Houser, along with other formal communication workshops. The other is preparation with the Global Policy Lab for the release of our paper on coronavirus anti-contagion policies, which included reading the book *Escape from the Ivory Tower: A Guide to Making Your Science Matter* by Nancy Baron.

[4]http://www.globalpolicy.science/covid19

Four essays: Giving voice to future generations

In this dissertation, the reader will find four essays examining the empirical effects and costs of future climate change; that is, seeking to give voice to the future generations who will bear the consequences of our collective decisions today. In three of these essays, we uncover climate impacts that result in a substantially altered future from our present reality: crop yields in sharp decline (while accounting for costly adaptation); cropped areas shifted away from present-day locations; and increased human mortality, with costly adaptations undertaken to avoid even further excess deaths. In one essay, we estimate that the impact of climate change on global energy consumption will be small in aggregate, though regional disparities may be important. In these essays, we seek to include information from as globally representative a sample as possible, and where possible estimate heterogeneous effects for the global rich versus the global poor. We also seek to graphically illustrate projected effects so that individuals can locate themselves, and perhaps their own children, within the empirical estimates we produce. We cannot force global leaders to act on this information. But perhaps constituents and local leaders, inspired by accurate and compelling media reporting, can.

I leave for reflection the words recorded by the Hebrew prophet Isaiah, whose reinterpretation of *Fiat Lux* will, I hope, never cease to structure my own motivation to engage in research.

Andrew B. Hultgren
August 14, 2020

References

Anthoff, David and Richard SJ Tol. 2013. "The uncertainty about the social cost of carbon: A decomposition analysis using FUND." *Climatic Change* 117 (3):515–530.

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

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

Auffhammer, Maximilian and Wolfram Schlenker. 2014. "Empirical studies on agricultural impacts and adaptation." *Energy Economics* 46:555–561.

Burke, Marshall, Melanie Craxton, Charles D Kolstad, Chikara Onda, Hunt Allcott, Erin Baker, Lint Barrage, Richard Carson, Kenneth Gillingham, Josh Graff-Zivin et al. 2016. "Opportunities for advances in climate change economics." *Science* 352 (6283):292–293.

Burke, Marshall and Kyle Emerick. 2016. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8 (3):106–40. URL https://www.aeaweb.org/articles?id=10.1257/pol.20130025.

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

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

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

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

Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer. 2010. "Linkages among climate change, crop yields and Mexico–US cross-border migration." *Proceedings of the National Academy of Sciences* 107 (32):14257–14262. URL https://www.pnas.org/content/107/32/14257.

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.

Hagerty, Nick. 2020. "The Scope for Climate Adaptation: Evidence from Water Scarcity in Irrigated Agriculture." *Working Paper* .

Hope, Chris. 2002. "The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern." *Integrated Assessment* 6 (1). URL https://journals.sfu.ca/int_assess/index.php/iaj/article/view/227.

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

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

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

Johnstone, Sarah and Jeffrey Mazo. 2011. "Global Warming and the Arab Spring." *Survival* 53 (2):11–17. URL https://doi.org/10.1080/00396338.2011.571006.

Kulp, Scott and Benjamin Strauss. 2019. "New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding." *Nature Communications* 10 (4844):1–12. URL https://www.nature.com/articles/s41467-019-12808-z.

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.

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.

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

Pindyck, Robert S. 2013. "Climate change policy: What do the models tell us?" *Journal of Economic Literature* 51 (3):860–872.

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

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

Shrader, Jeffrey. 2018. "Expectations and Adaptation to Environmental Risks." *SSRN Working Paper* .

Steffen, Will, Johan Rockström, Katherine Richardson, Timothy M. Lenton, Carl Folke, Diana Liverman, Colin P. Summerhayes, Anthony D. Barnosky, Sarah E. Cornell, Michel Crucifix, Jonathan F. Donges, Ingo Fetzer, Steven J. Lade, Marten Scheffer, Ricarda Winkelmann, and Hans Joachim Schellnhuber. 2018. "Trajectories of the Earth System in the Anthropocene." *Proceedings of the National Academy of Sciences* 115 (33):8252–8259. URL https://www.pnas.org/content/115/33/8252.

Stern, Nicholas. 2006. "Stern review report on the economics of climate change." .

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.

Tol, Richard SJ. 2009. "The economic effects of climate change." *Journal of economic perspectives* 23 (2):29–51.

Chapter 2: The Impacts of Climate Change on Global Grain Production Accounting for Adaptation

Chapter Summary

In this opening chapter, we present the first granular and globally comprehensive analysis of crop yields under climate change, accounting for farmer adaptation and its costs along the critical dimensions of slow-moving changes to temperature and precipitation, income growth, and irrigation.[1] To do this, we assemble once of the largest datasets of subnational crop production available to researchers, covering over 10,000 administrative units from 55 countries and representing two thirds of global crop caloric production.

Across crops, we find consistent patterns of adaptation to climatic changes and through growth in incomes and access to irrigation. However, we also find that globally widespread yield declines persist, in spite of these adaptive efforts. We show this by adapting simple machine learning techniques to a causal inference context, enabling us to for the first time systematically assess a rich set of channels through which temperature and precipitation shocks might causally affect yields in each of our crops, while also systematically assessing the extent to which farmers might use income growth, irrigation, or expectations of long-run temperature and precipitation to mitigate the extent of these yield losses.

Projecting our empirical estimates into the future, we find that even after accounting for costly adaptation, RCP 8.5 end-of-century yield losses are projected to be severe in the key cropped areas of the U.S. grain belt and Eastern China (40% or greater losses), important in the major cropped areas of South America and Australia (10% losses), and mixed in Europe, Sub-Saharan Africa, India, and Southeast Asia.

Overall, while costly adaptation is protective for the major grain producing regions of the world relative to no adaptation at all, yield impacts in spite of adaptation are still generally negative and for many regions of the world they are substantially so.

[1] This material first appeared as a working paper of the same title, with authors Andrew Hultgren, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Amir Jina, Robert Kopp, Kelly McClusker, Ishan Nath, James Rising, Ashwin Rode, and Jiacan Yuan. This project is an output of the Climate Impact Lab consortium that gratefully acknowledges funding from the Carnegie Corporation, Energy Policy Institute of Chicago (EPIC), International Growth Centre, National Science Foundation (#SES1463644), Sloan Foundation, and Tata Center for Development. We thank Laura Alcocer, Thomas Bearpark, Trinetta Chong, Greg Dobbels, Radhika Goyal, Simon Greenhill, Dylan Hogan, Azhar Hussain, Theodor Kulczycki, Maya Norman, Sebastien Phan, Yuqi Song, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and we thank Jack Chang, Megan Landín, and Terin Mayer for excellent project management. We acknowledge the World Climate Research Programme's Working Group on Coupled Modeling, which is responsible for CMIP, and we thank the climate modeling groups for producing and making available their model output. For CMIP, the U.S. Department of Energy's Program for Climate Model Diagnosis and Intercomparison provides coordinating support and led development of software infrastructure in partnership with the Global Organization for Earth System Science Portals.

2.1 Introduction

One of the most threatening impacts of climate change centers on human food systems (Lobell et al., 2008; Lobell, Schlenker, and Costa-Roberts, 2011). Absent any adaptation, global grain production is anticipated to decline sharply under climate change (Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010; Welch et al., 2010; Lobell, Schlenker, and Costa-Roberts, 2011; Auffhammer, Ramanathan, and Vincent, 2012; Burke and Emerick, 2016; Gammans, Mérel, and Ortiz-Bobea, 2016; Fishman, 2016; Burgess et al., 2017; Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016). However, it is not known to what extent producers can mitigate against the effects of slow-moving and foreseeable climatic changes, especially as global incomes grow and adaptive resources increase; indeed the limited empirical evidence (Butler and Huybers, 2013; Burke and Emerick, 2016) on farmer adaptation has shown mixed results (Auffhammer and Schlenker, 2014).

We present the first granular and globally comprehensive analysis of crop yields under climate change, accounting for farmer adaptation and its costs along the critical dimensions of slow-moving changes to temperature and precipitation, income growth, and irrigation. To do this, we assemble once of the largest datasets of subnational crop production available to researchers, covering over 10,000 administrative units from 55 countries and representing two thirds of global crop caloric production.

Across crops, we find consistent patterns of adaptation to climatic changes and through growth in incomes and access to irrigation. However, we also find that globally widespread yield declines persist, in spite of these adaptive efforts. We show this by adapting simple machine learning techniques to a causal inference context, enabling us to for the first time systematically assess a rich set of channels through which temperature and precipitation shocks might causally affect yields in each of our crops, while also systematically assessing the extent to which farmers might use income growth, irrigation, or expectations of long-run climate to mitigate the extent of these yield losses.

Our approach to adaptation and its costs (Carleton et al., 2019) does not require us to observe which adaptive measures farmers might employ, nor does it require us to observe the costs of these measures. We employ reduced-form empirical techniques to uncover the magnitude of farmer adaptation without requiring knowledge of the particular measures farmers might take across the broad span of geographies and socioeconomic contexts represented in our data.

This analysis overcomes three key hurdles that have faced the literature on climate change and agriculture. First, rather than focusing on a single country or region, our broad data represents yield responses to weather across a wide range of climactic and socioeconomic contexts, and its granularity enables us to match to high-resolution historical temperature and precipitation data, allowing us to conduct robust causal inference at a global scale. Second, this rich data uniquely enables us to systematically distinguish which weather channels are most impactful on crop yields across a broad set of key global crops – weather channels that have only been evaluated piecemeal in the existing literature. Finally, we are for the first time able to recover reduced-form effects of farmer adaptive efforts, as well as the costs of these efforts, as they have been practiced in the full range of global climactic and socioeconomic contexts.

2.2 Data

We assemble one of the largest datasets on subnational agricultural production available to researchers, acquiring data from the statistical offices of 55 countries and assembling a panel of observations from over 10,000 administrative units spanning all continents and a wide range of socioeconomic contexts (Figure 2.1a). Our data covers the four major global grains (maize, rice, soy, and wheat[2]) and two regionally important crops (cassava and sorghum), representing two-thirds of global crop caloric production.

We match this granular crop production data to global, high resolution (0.1 degree) weather data from the Global Mean Forcing Dataset for Land Surface Modeling version 1 (Sheffield, Goteti, and Wood, 2006) reanalysis product, providing us with measures of daily maximum and minimum temperatures and daily precipitation at high resolution going back to the year 1948. Matching these datasets gives a panel of highly resolved yields and weather with global coverage, enabling the systematic analyses described below.

Additional data is described in the Supplementary Methods and in Appendix Table 2.A.1.

2.3 Summary of Our Approach

Beyond the coarse measures of growing season temperature and precipitation, a range of finely resolved weather measures have been examined in various studies for their potential effect on grain yields, including measures such as extreme heat degree days, vapor pressure deficit, drought extremes, or the count of rainy days (Schlenker and Roberts, 2009; Auffhammer, Ramanathan, and Vincent, 2012; Roberts, Schlenker, and Eyer, 2013; Fishman, 2016). However, because these weather measures are all correlated with one another and have not been systematically evaluated, it is unclear which are most determinative of crop yields. Further, we seek to allow for a rich set of interaction terms with our weather measures, to capture the various dimensions (long run temperature, long run precipitation, income, and irrigation) over which farmers might make adaptive decisions that modulate the effects of weather on yields. To resolve these issues while also reducing the combinatorial space of the problem, we break our analysis up into two cross validation steps.

Selection of weather measures

In the first cross validation step, we use k-fold cross validation to systematically evaluate a broad range of weather measures. We select those weather measures that: 1) are important determinants of yields, and 2) are able to be reliably projected into the future by climate models. In this step, weather terms are each interacted with a fixed set of covariates: growing season average daily maximum temperature ("long-run temperature"), $\log(GDP_{pc})$ ("income"), and the share of cropped area equipped for irrigation ("irrigation"). This fixed set of covariate interactions allows us to capture some of the flexibility that will be required of the model in the second cross validation step, without having to simultaneously explore the full set of potential combinations of weather measures and covariate interactions.

[2]The wheat analysis is not included at present.

Selection of adaptation dimensions
In our second cross validation step, select a set of covariate interactions to model farmer adaptation, where covariates reflect the information set over which farmers make their adaptive decisions. For example, if a given location is on average quite hot, farmers are likely to chose to plant crop varietals that are more resilient to heat stress, anticipating that the coming growing season will be similar to past growing seasons. In this cross validation step, we take as given the set of weather parameters carried forward from the first step. We allow for the previously described interactions with income, irrigation, and long-run temperature, as well as growing season average monthly precipitation ("long-run precipitation") and an interaction between long-run temperature and long-run precipitation. Further, we allow for the nonlinear precipitation response to vary over the growing season (under month-grouping determined through a sequential set of F-tests) or to be constant; and we allow for interaction terms between degree-day responses and precipitation. In this step, we use a country-fold cross validation procedure, in which each fold of the data reflects all data from a single country. Here, the test of model performance is the ability to model yields for a country which is held out of the data. This matches one of the intended uses of our interaction surface, which is to model the response of yields to weather for locations with no yield data available.

Cross validation in a causal inference context
Cross validation is most typically applied in prediction problems where causal inference is not the primary interest. We adapt the cross validation procedure in the following manner in order to apply it to our causal inference problem. We first select our controls for unobservables (fixed effects and time trends) so as to isolate plausibly random variation in our set of weather measures. We do not select these controls in our cross validation procedure, as the objective in this step is to isolate random variation in weather shocks for the purpose of causal inference; the objective is not to best predict yields through our controls. Next, relying on the Frisch-Waugh (1933) and Lovell (1963) result for partitioned regression, we partial out the variation in our outcomes y_{it} and our vector of regressors $\mathbf{X}_{it}$ that is associated with our controls for unobservables $\mathbf{\Phi}_{it}$.

$$y_{it} = \mathbf{\Phi}_{it} + \tilde{y}_{it} \tag{2.1}$$

$$\mathbf{X}_{it} = \mathbf{\Phi}_{it} + \tilde{\mathbf{X}}_{it} \tag{2.2}$$

We then use the residual variation $\tilde{y}_{it}$ and $\tilde{\mathbf{X}}_{it}$ in our cross validation procedures, regressing $\tilde{y}_{it}$ on elements of $\tilde{\mathbf{X}}_{it}$. In our specification, y_{it} is log yields in administrative unit i and year t, and $\mathbf{X}_{it}$ is the vector of weather measures shown in Figure 2.1b, with potential interactions as described above. $\mathbf{\Phi}_{it}$ is a vector of controls composed of administrative unit i (typically county) fixed effects, absorbing time-invariant confounds such as soil quality; country-year fixed effects, absorbing price and trade policy confounds; and state/province quadratic time trends, absorbing slow-moving, local confounds such as the diffusion of new technologies.

2.4 Results

Selection of weather measures

Figure 2.1b presents the results of our cross validation procedure for the selection of weather measures, which searched over more than 2,800 models across all five crops. Of all the weather measures examined, degree days (top row, constructed following Schlenker and Roberts 2009) most strongly improve model out-of-sample fit across all crops except rice, where fit improvements are present but more modest. Inclusion of linear (third row) and quadratic (fourth row) precipitation also substantially improve out-of-sample fit across all crops, with more modest gains from higher order polynomial terms. (Note: inclusion of a higher order polynomial term implies inclusion of the lower order terms.) For all crops, we carry the degree day terms and quadratic precipitation polynomials forward to the selection of adaptation dimensions.

Daily minimum temperatures and vapor-pressure deficit contribute quite modest gains in fit, and we generally do not carry them forward in the analysis. The exception is rice; some improvement in fit is observed with t_{min} (see Figure 2.2c), and given the balance of the literature on minimum temperature effects on rice yields (Peng et al., 2004; Welch et al., 2010; Auffhammer, Ramanathan, and Vincent, 2012; Dong et al., 2014) we do carry this term forward to our second step cross-validation selection of adaptation dimensions.

We include results for terms that are not well-projected by existing climate models as well, in part to point out where advances in climate modeling might yield large information gains for society. Of particular note is the importance of the growing season drought measure for cassava (bottom row, far right). Cassava yields are not very responsive to weather overall, but drought appears to be a relatively important determining factor in yields for this regionally important crop. Drought is also of some importance in understanding rice yields, as is the number of rainy days during the growing season. Finally, the number of growing season rainy days appears to be a moderately import determinant of maize yields, with lesser contributions from the number of extreme rain days and the amount of extreme rain the crop is exposed to.

Overall, Figure 2.1b is a surprisingly strong affirmation of the seminal Schlenker and Roberts 2009 model. The inclusion of additional weather terms for different crops does further improve model fit, but the strongest gains clearly result from the degree day terms and precipitation quadratics. And, even more surprisingly, this generally holds across the more broad set of crops examined in this study.

The general importance of the degree day and quadratic precipitation measures can also be seen in Figure 2.2, which displays the rank-ordering of model out-of-sample fits for all models examined in the cross validation procedure for each crop. In each plot, the model that we carry forward is vertically offset only slightly below the best-fit "projectable" model (the highest red dot). Because we will take the weather measures selected here and carry them forward into a rich interaction surface, parsimony at this stage is important. Figure 2.2 reinforces our finding that the degree day and quadratic precipitation measures are the weather measures that are most determinative of yields for the crops in our study, and so will be those measures over which adaptation will be most important to model.

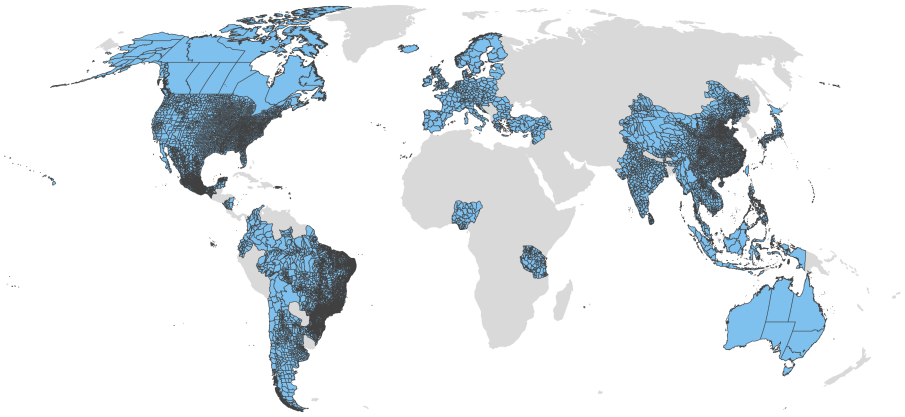

(a)

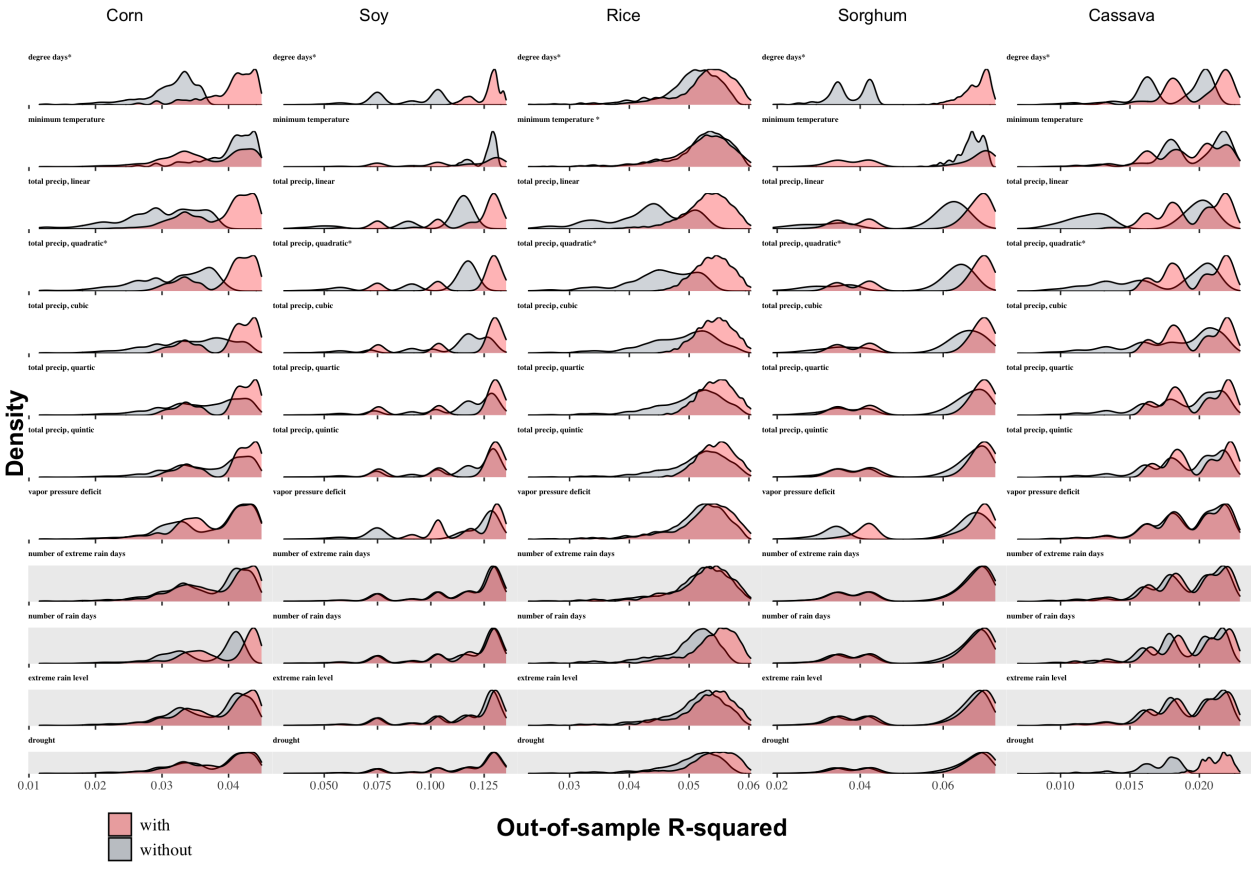

(b)

Figure 2.1: Panel (a) Coverage of our subnational yield data. Most country data is reported at the equivalent county level for that country. Many administrative units are too small to be clearly visible. Panel (b) The results of our first cross validation procedure for selecting weather variables. Pink represents the kernal density of out-of-sample R^2 for all model permutations including a given weather term; grey represents the analogous density for all model permutation exluding a given term. Columns are crops, rows are weather terms. Starred terms were carried forward to the next cross validation procedure. The bottom four rows represent weather terms that cannot be reliably projected by existing climate models.

Dimensions of adaptation

Figure 2.3 displays the temperature and precipitation response surfaces that are selected by our second step cross validation procedure. The top panel displays the temperature response surfaces estimated for each of our crops. The first row shows how the temperature response is modulated across the terciles of long-run temperature in our data (blue, lower tercile; yellow, mean; red, upper tercile). Likewise, the second row shows the the temperature response is modulated across the terciles of long-run precipitation in our data. Across crops, we see that the temperature response function flattens for locations that are hotter on average (first row, blue to red). For maize, soy, rice, and cassava the temperature response also flattens for places that are drier on average (second row, red to blue).

The third and fourth rows shows how the temperature response is modulated across the terciles of average income and access to irrigation. Income can have competing effects on a crop's temperature response function: on the one hand, it is possible that income captures average borrowing and liquidity constraints, and technology availability, to farmers in our sample. In this case, temperature response functions might become steeper as incomes rise, as wealthier farmers are able to access higher yielding (but more temperature sensitive) seed varietals of a given crop. We see this behavior in maize and sorghum. On the other hand, income may allow farmers to better access protective technologies, such as better tilling and irrigation technologies, or increased ability to pay for additional labor, which could flatten the temperature response function. This appears to be the case for rice and cassava. Regarding irrigation, gains from access to irrigation almost always manifest as increased gains from moderate temperatures, which is consistent with the enhanced ability of a crop to keep stomata open and transpire CO_2 for photosynthetic activity when it is not water stressed (see Osakabe et al 2014 for an overview).

The bottom panel of Figure 2.3 show how the precipitation response function is modulated by long-run average temperature (top row) and long-run average precipitation (second row). For maize, soy, and rice, precipitation responses differ by growing season phases; precipitation in the middle of the growing season is particularly important for these three grains and is shown here. For these middle portions of the growing season, the yield losses from a dry month are larger for places that are on average hotter (top row, blue to red) and drier (second row, red to blue). Cassava and sorghum are less responsive to precipitation overall, though the pattern of increased losses from a dry month in places that are on average drier does appear to hold for these crops as well. For all crops, greater access to income (third row) and irrigation (fourth row) flatten the precipitation response, indicating that both income and irrigation are protective against yield losses during a particularly dry month.

Projected climate impacts

Figure 2.4 displays the results of projected log yields at the year 2090 under RCP 8.5 using the CCSM4 climate model (a "middle of the road" warming model, Gent et al. 2011), for each of the crops we study. The first column displays yield impacts if adaptation does not occur; that is, if the temperature and precipitation response functions are fixed at their present day (2010) values. In this scenario, we find widespread and substantial yield losses for all crops in nearly all major growing regions of the world, with losses in major growing regions ranging from 25 log points (22% losses) to upwards of 150 log points (78% losses).

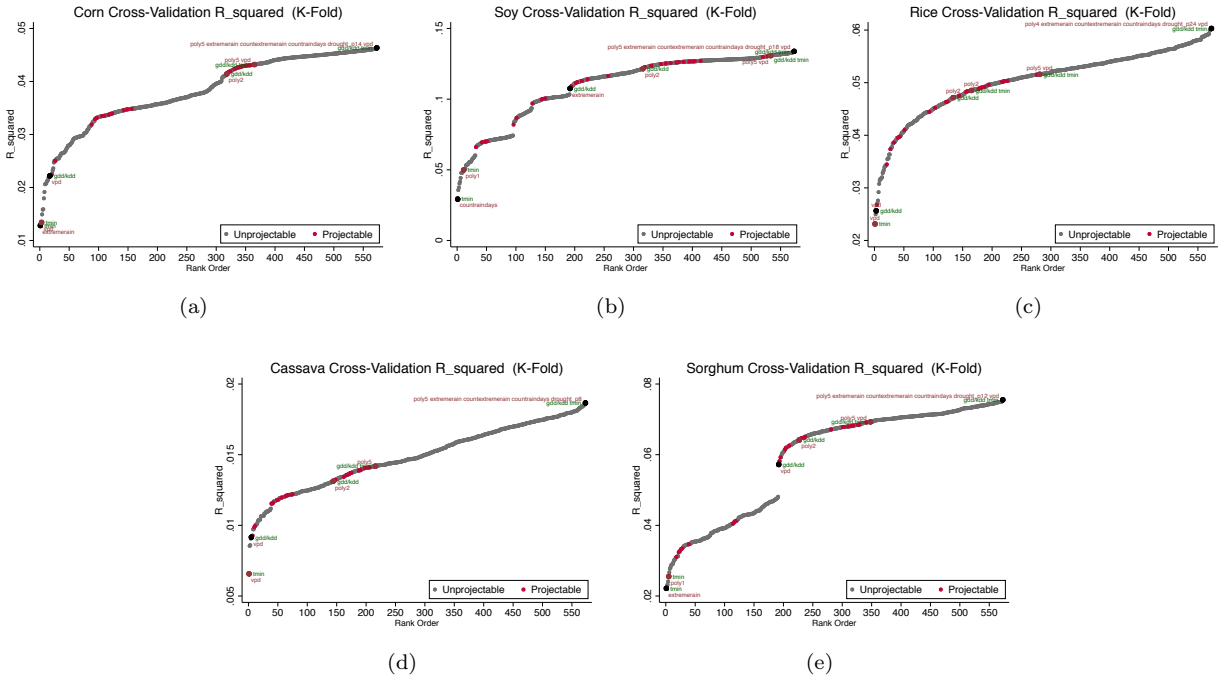

Figure 2.2: Rank ordering of out-of-sample R^2 performance of different models in the first step cross validation (selection of weather measures). Models where all weather measures are projectable by climate models are plotted as red dots, those that contain at least one detailed precipitation measure that is not well-projected by climate models are plotted as grey dots. Generally, the model carried forward for selection of adaptation dimensions is the model labeled "gdd/kdd, poly2," with exception of rice, where we also carry forward t_{min}. Plots are for: (a) Maize, (b) Soy, (c) Rice, (d) Cassava, (e) Sorghum.

The one exception is rice in Northern Europe, which may benefit from climate change.

The second column displays yield impacts allowing for adaptation, but assuming it is costless (Butler and Huybers, 2013; Schlenker, Roberts, and Lobell, 2013). Here we can see that adaptation in some places fully mitigates the impacts of climate change, even flipping the impact from negative to positive in parts of South America, Central Africa, Southeast Asia, and the U.S. Elsewhere, adaptation increases exposure to climate impacts, in places like Northern Europe, Central Asia, and the Middle East.

The third column of Figure 2.4 displays yield impacts allowing for adaptation and incorporating our estimates of the costs of adaptation (in terms of yields). Here we see that places that gained from adaptation by reducing their exposure to climate change lose from costs, and those places that increase their climate exposure under adaptation gain from costs. Intuitively, a "gain" from costs can be thought of as an upward shift in average yields, as might result from switching to a seed varietal that is more exposed to extreme heat but also higher yielding on average.

In this third column of Figure 2.4, we see for all grains for which costs have been estimated (maize, soy, and rice) that most of the major crop growing regions of the world see projected yield losses under climate change, even after accounting for costly adaptation. Losses in the grain belt of the U.S. and Eastern China are severe, around 50 log points (40%) for maize and even higher for soy. Losses in the major cropped regions of South America are more moderate for maize and soy, around 10%. Soy impacts in Europe are moderate to severe (10-40%) while for maize they are mixed. Maize in Sub-Saharan Africa, and India are mixed and near zero; soy may see some yield gains in these regions. Rice, primarily grown in India and Southeast Asia, sees mixed impacts in these regions. Yields in North American show moderate declines, while those in Europe show moderate gains.

Relying on the Ricardian approach, the seminal work of Mendelsohn, Nordhaus, and Shaw (1994) found that climate impacts on U.S. agricultural profits, net of adaptation, would be modestly negative and possibly beneficial. Using fixed effects estimators more robust to omitted variables bias (Deschênes and Greenstone, 2007), the subsequent literature has found negative climate impacts across a range of contexts, though estimates are generally unable to account for adaptation (Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016).

Similar to Auffhammer (2018), the study presented here combines elements of both the Ricardian and within-estimator approaches: we use cross-sectional variation in long-run climate to uncover heterogeneity in the response to short-run (within-estimator) weather shocks. In so doing, we find that adaptation does indeed mitigate the effects of weather relative to the effect of an unanticipated shock.

However, for the first time we also uncover the implied costs of these adaptive actions in agricultural yields. In contrast to Mendelsohn, Nordhaus, and Shaw 1994, we find that accounting for both adaptation and its costs, negative effects of climate change persist and are widespread.

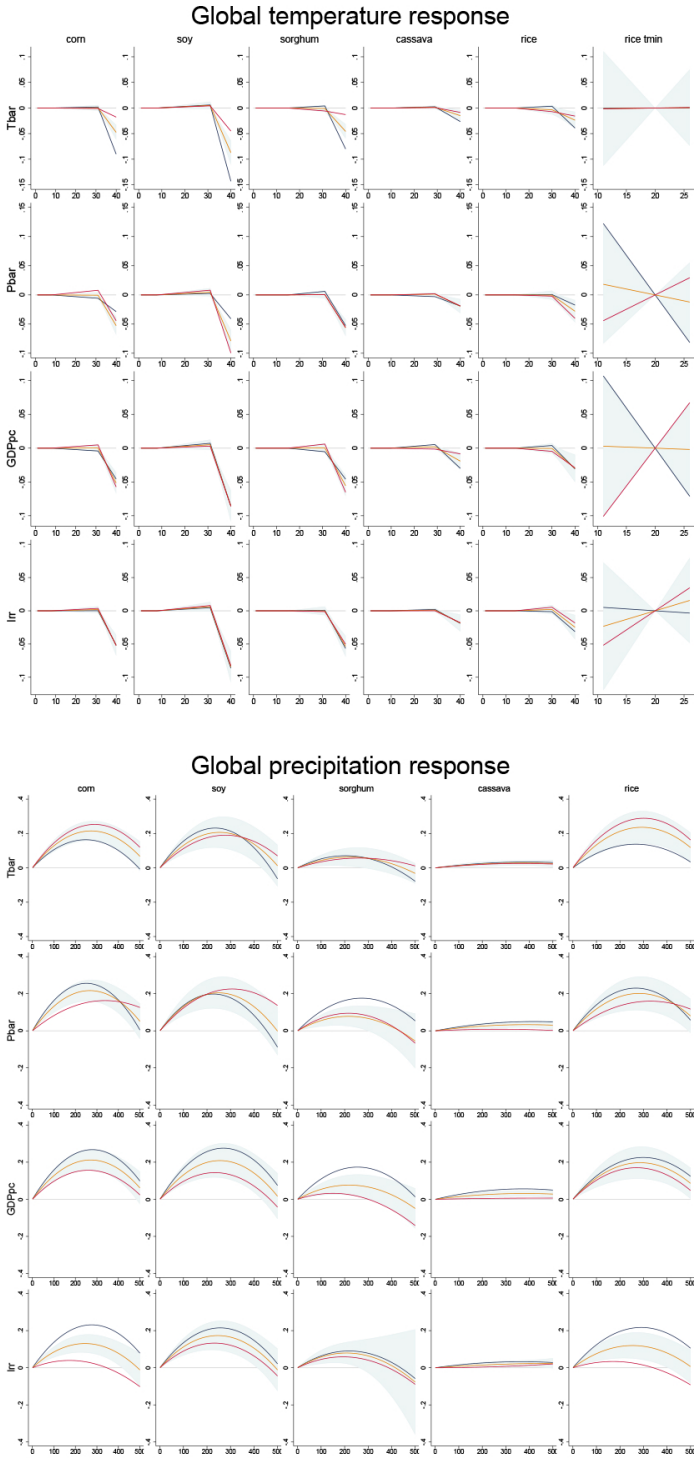

Figure 2.3: Results from the second cross validation step: empirically estimated temperature and precipitation response surfaces, by crop and by adaptation parameter. Columns are crops, rows are adaptation parameters (covariates). Yellow: all covariates evaluated at the sample mean; blue: covariate of interest evaluated at its lower tercile; red: covariate of interest evaluated at its upper tercile. Note the top panel includes an extra column for rice (t_{min}). Standard errors robust to correlated shocks within state and within country-year.

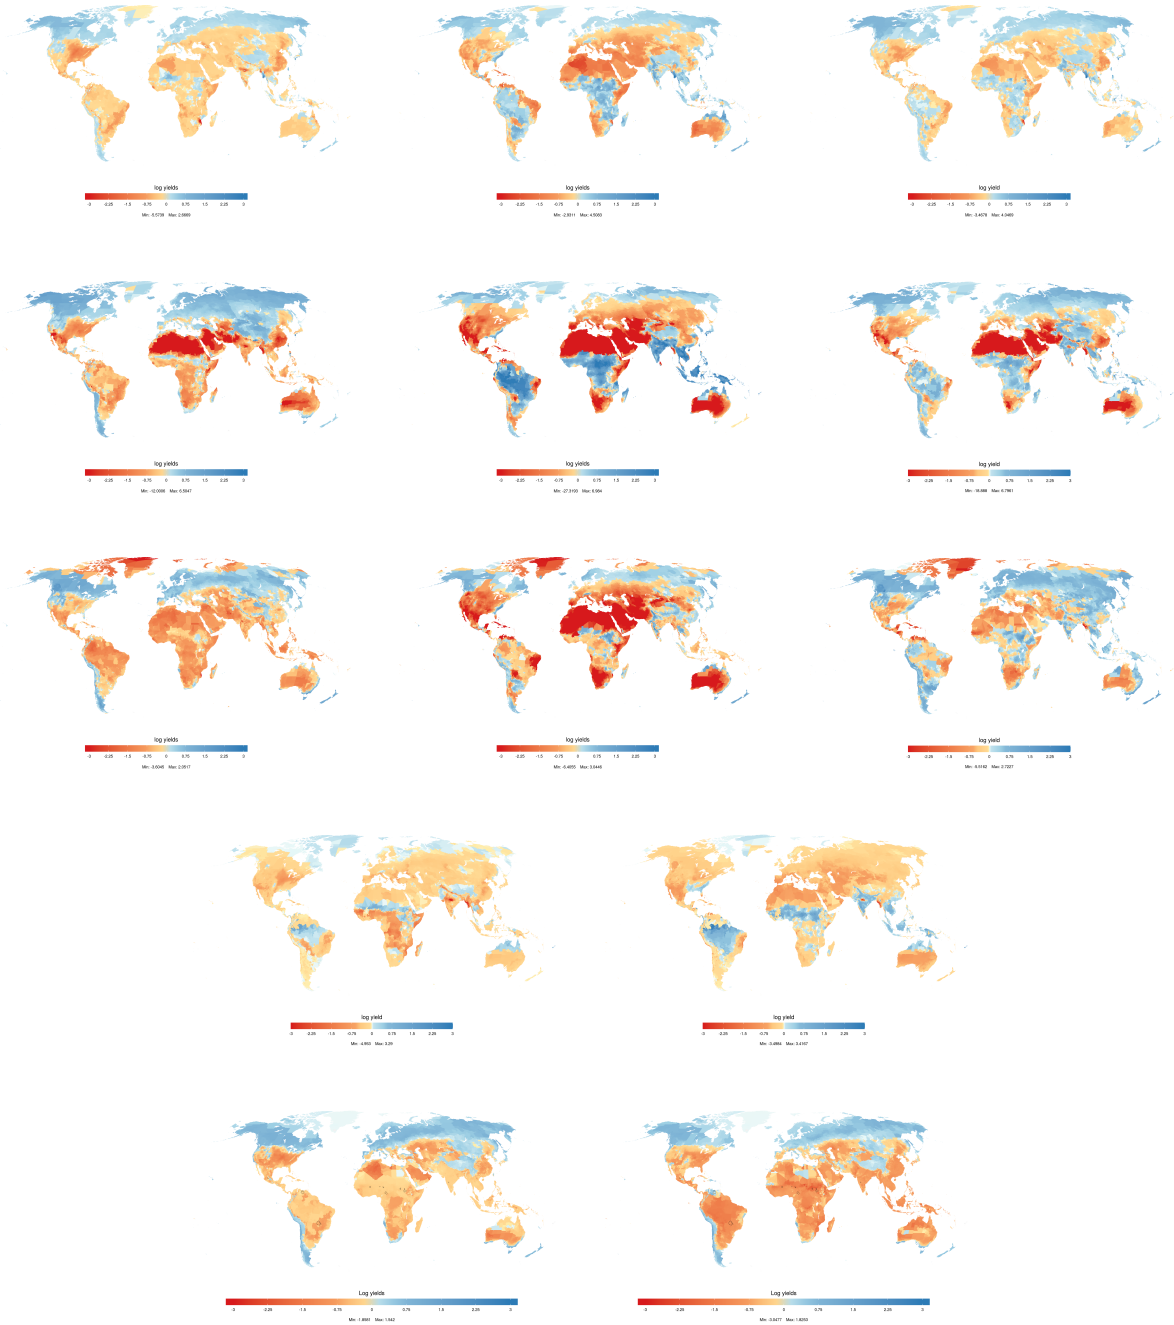

(a) (b)

Figure 2.4: Projected climate impacts under: (a) No adaptation, (b) Income and climate adaptation, (c) Income and climate adaptation including adaptation costs. Rows are maize, soy, rice, cassava, and sorghum. Incomes taken from Shared Socioeconomic Pathway 3 (SSP3), emissions from the RCP 8.5 warming scenario. Note: costs for sorghum and cassava have not yet been projected. All projections use the CCSM4 climate model except the rice projections, for which CCSM4 was not available at the time of writing and CESM1-BGC was used instead.

2.5 Supplementary Methods

Here we provide additional details for each of the steps in our analysis.

Weather terms
Here we briefly describe the construction of the weather terms that entered our first-stage cross validation procedure. For all weather terms, weighting and aggregation to administrative units were done as described in Carleton et al. 2019 (but using cropped area weights from SAGE rather than population weights).

Degree day terms were constructed following Schlenker and Roberts 2009, with kink points selected for each crop by a 10-fold cross validation search over potential kinks. Construction of vapor pressure deficit followed Roberts, Schlenker, and Eyer 2013. Extreme rain was identified by constructing the 1980-2010 rainfall distribution for each grid cell in the GMFD daily precipitation data, and aggregating the amount of rainfall or number of days of rainfall in excess of the 95th percentile of this distribution. Drought was defined using a 10-fold cross validation search over percentile cutoffs in the administrative-unit growing season precipitation distribution.

Cross validation for the selection of weather measures
When using cross validation to select weather measures, each model was required to have at least one temperature term (degree days and/or tmin) and one "precipitation" term (precipitation polynomials, amount of extreme rain, count of extreme rain days, count of rain days, and/or vapor pressure deficit).

Cross validation for the selection of adaptation dimensions
Before the cross validation selection of adaptation dimensions (covariate interactions), potential groupings of months were determined, to allow the effect of precipitation to optionally differ by phase of the growing season. Month groupings were determined by constructing the following sequence of F-tests:

1. Test for a differential precipitation response between the first month and all remaining months of the growing season.

2. If the result from (1) indicates a significant difference, insert a "phase" break in the growing season and restart this procedure with the remaining months.

3. If not, group this month with any earlier months and restart with the next month in the season.

Whether to use growing season total precipitation or to allow the precipitation response to vary by these phases of the growing season was then selected as part of the second cross validation procedure.

In the cross validation selection of adaptation dimensions (covariate interactions), models were required to include at least long-run temperature, long-run precipitation, income, and irrigation interactions. The search was to additionally include long-run temperature x long-run precipitation interactions, degree day x precipitation terms, and whether these models

should allow precipitation to enter as a growing season total or differentially by phase of the growing season. Early work suggested that the main covariate interactions were always included in the final model selection, so to save computational time they were simply taken as required. We may in the future relax that requirement to show the full search.

Projections
Projections were conducted as described in Carleton et al. 2019.

Adaptation costs
Our adaptation costs estimation followed Carleton et al. 2019. Because the context differs (farmer profit maximization versus individuals maximizing lifetime utility), the cost theory is briefly outlined in the Appendix 2.B.

Chapter Review and Looking Ahead

In this first chapter, we presented the first granular and globally comprehensive analysis of crop yields under climate change, accounting for farmer adaptation and its costs. Overall, while costly adaptation is protective for the major grain producing regions of the world relative to no adaptation at all, yield impacts in spite of adaptation are still generally negative under the RCP 8.5 high emissions scenario. Further, for many important grain-producing regions of the world, yield impacts from climate change are projected to be severe.

In the next chapter, we will examine a particular channel of climate adaptation in agricultural production; a channel which is widely discussed but for which there is little empirical evidence: the migration of cropped areas as climatic suitability shifts.

References

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

Auffhammer, Maximilian, V. Ramanathan, and Jeffrey R. Vincent. 2012. "Climate change, the monsoon, and rice yield in India." *Climatic Change* 111 (33):411–424. URL https://doi.org/10.1007/s10584-011-0208-4.

Auffhammer, Maximilian and Wolfram Schlenker. 2014. "Empirical studies on agricultural impacts and adaptation." *Energy Economics* 46:555–561.

Burgess, Robin, Olivier Deschênes, David Donaldson, and Michael Greenstone. 2017. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." *Working Paper* URL https://epic.uchicago.edu/wp-content/uploads/2019/07/Publication-9.pdf.

Burke, Marshall and Kyle Emerick. 2016. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8 (3):106–40. URL https://www.aeaweb.org/articles?id=10.1257/pol.20130025.

Butler, Ethan and Peter Huybers. 2013. "Adaptation of US maize to temperature variations." *Nature Climate Change* 3 (1):68–72. URL https://doi.org/10.1038/nclimate1585.

Carleton, Tamma, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E Kopp, Kelly McCusker, Ishan Nath et al. 2019. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2018-51* URL https://ssrn.com/abstract=3224365.

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

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

Dong, Wenjun, Jin Chen, Lili Wang, Yunlu Tian, Bin Zhang, Yongcai Lai, Ying Meng, Chunrong Qian, and Jia Guo. 2014. "Impacts of nighttime post-anthesis warming on rice productivity and grain quality in East China." *The Crop Journal* 2 (1):63 – 69. URL http://www.sciencedirect.com/science/article/pii/S2214514113000378.

Fishman, Ram. 2016. "More uneven distributions overturn benefits of higher precipitation for crop yields." *Environmental Research Letters* 11 (2):024004. URL https://iopscience.iop.org/article/10.1088/1748-9326/11/2/024004/meta.

Gammans, Matthew, Pierre Mérel, and Ariel Ortiz-Bobea. 2016. "The impact of climate change on cereal yields: Statistical evidence from France." 2016 Annual Meeting, July 31-August 2, Boston, Massachusetts 236322, Agricultural and Applied Economics Association. URL https://ideas.repec.org/p/ags/aaea16/236322.html.

Gent, Peter R., Gokhan Danabasoglu, Leo J. Donner, Marika M. Holland, Elizabeth C. Hunke, Steve R. Jayne, David M. Lawrence, Richard B. Neale, Philip J. Rasch, Mariana Vertenstein, Patrick H. Worley, Zong-Liang Yang, and Minghua Zhang. 2011. "The Community Climate System Model Version 4." *Journal of Climate* 24 (19):4973–4991. URL https://doi.org/10.1175/2011JCLI4083.1.

Lobell, David B., Marshall B. Burke, Claudia Tebaldi, Michael D. Mastrandrea, Walter P. Falcon, and Rosamond L. Naylor. 2008. "Prioritizing Climate Change Adaptation Needs for Food Security in 2030." *Science* 319 (5863):607–610. URL https://science.sciencemag.org/content/319/5863/607.

Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate Trends and Global Crop Production Since 1980." *Science* 333 (6042):616–620. URL https://science.sciencemag.org/content/333/6042/616.

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.

Peng, Shaobing, Jianliang Huang, John E. Sheehy, Rebecca C. Laza, Romeo M. Visperas, Xuhua Zhong, Grace S. Centeno, Gurdev S. Khush, and Kenneth G. Cassman. 2004. "Rice yields decline with higher night temperature from global warming." *Proceedings of the National Academy of Sciences* 101 (27):9971–9975. URL https://www.pnas.org/content/101/27/9971.

Roberts, Michael J., Wolfram Schlenker, and Jonathan Eyer. 2013. "Agronomic Weather Measures in Econometric Models of Crop Yield with Implications for Climate Change." *American Journal of Agricultural Economics* 95 (2):236–243. URL http://www.jstor.org/stable/23358387.

Schlenker, Wolfram and David B Lobell. 2010. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5 (1):014010. URL https://iopscience.iop.org/article/10.1088/1748-9326/5/1/014010/meta.

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

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

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

Welch, Jarrod R., Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. 2010. "Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures." *Proceedings of the National Academy of Sciences* 107 (33):14562–14567. URL https://www.pnas.org/content/107/33/14562.

Appendix

2.A Data

The table below describes the various datasets used in this analysis. For further details, see Carleton et al. 2019.

Dataset	Variables	Resolution
Historical		
GMFD v1	Temp, precip	0.25°×0.25°
National Stat. Agencies	Crop production	ADM2, ADM1
Gennaioli et al. (2014)	GDP	ADM1
NASA DMSP-OLS	Nighttime lights	30 arcsecond
SAGE	Cropped area & growing seasons	0.1°×0.1°
FAO Aquastat	Irrigation (area equipped)	0.1°×0.1°
Future		
NASA NEX-GDDP	Temp, precip	0.25°×0.25°
OECD Env-Growth	GDP	National

Table 2.A.1: Sources of data used in this analysis.

2.B Adaptation costs

Profit maximization

The following assumes adaptation actions $\mathbf{b}$ are continuous, do not have fixed costs, and are not durable (do not persist over time). Area planted and prices are held fixed, for simplicity.

Let profits π be determined by area planted A, yield per acre y, prices p, and costs per acre c. Denote climate $\mathcal{C}$ as the vector $(\bar{T}, \bar{P})$. Then in each period, profit is given by

$$\pi = Ay(T, P, \mathbf{b}, \alpha)p - Ac(\mathbf{b}, \alpha)$$

Where α represents a vector of weather-invariant cost and/or yield shifters (such as soil type). Profit maximizing farmers choose b at the start of the growing season, giving the following equalization of marginal benefits and marginal costs of adaptation:

$$\begin{aligned} \frac{\partial \pi}{\partial b} &= A\left[\frac{\partial E[y]}{\partial b}p - \frac{\partial c}{\partial b}\right] = 0 \\ \frac{\partial E[y]}{\partial b}p &= \frac{\partial c}{\partial b} \\ \implies b^* &= b^*(\mathcal{C}, E[T], E[P]) \end{aligned} \tag{2.B.1}$$

Climate Change

To think about damages from climate change, we will need an expression for $\frac{d\pi}{d\bar{T}}$.

$$\frac{d\pi}{d\bar{T}} = A\left[\left(\frac{\partial y}{\partial T}\frac{\partial T}{\partial \bar{T}} + \frac{\partial y}{\partial b}\frac{\partial b^*}{\partial \bar{T}}\right)p - \frac{\partial c}{\partial b}\frac{\partial b^*}{\partial \bar{T}}\right]$$

Plugging in[3] for $\frac{\partial c}{\partial b}$ from Eq. 2.B.1:

[3] Taking the familiar total derivative:

$$\frac{dy}{d\bar{T}} = \underbrace{\frac{\partial y}{\partial b}\frac{\partial b^*}{\partial \bar{T}}}_{\text{Indirect effect}} + \underbrace{\frac{\partial y}{\partial T}\frac{\partial T}{\partial \bar{T}}}_{\text{Direct effect}}$$

So,

$$\frac{\partial y}{\partial b}\frac{\partial b^*}{\partial \bar{T}} = \frac{dy}{d\bar{T}} - \frac{\partial y}{\partial T}\frac{\partial T}{\partial \bar{T}}$$

That is, the marginal gains from adaptation to $\bar{T}$ are equal to the observable fully adapted effects of climate change minus the losses when the long run adaptation channel is shut down. This, of course, equals the marginal cost of adaptation under the Eq. 2.B.1. Note that empirically we cannot disentangle the direct effects of temperature changes from short-run adjustments, because we only observe end-of-season yields. This means our adaptation costs are estimated net of any short-run adjustments farmers might make.

$$\frac{d\pi}{d\bar{T}} = A\Bigg[\underbrace{\frac{dy}{d\bar{T}}p}_{\substack{\text{Fully adapted effects} \\ \text{of climate change}}} - \underbrace{\left(\frac{dy}{d\bar{T}} - \frac{\partial y}{\partial T}\frac{\partial T}{\partial \bar{T}}\right)p}_{\text{Estimated adaptation costs}} \Bigg] \tag{2.B.2}$$

Note that this equation embeds the estimated adaptation costs for a marginal change in $\bar{T}$. A similar expression exists for $\bar{P}$.

We now turn to damages from climate change. Damages from climate change over the period from time 1 to time 2 are given by:

$$\pi(T(\mathcal{C}_2), b(\mathcal{C}_2, I_2)) - \pi(T(\mathcal{C}_1), b(\mathcal{C}_1, I_2))$$

$$= \int_{\mathcal{C}_1}^{\mathcal{C}_2} \frac{d\pi}{d\mathcal{C}} d\mathcal{C}$$

$$= \int_{\bar{T}_1}^{\bar{T}_2} \int_{\bar{P}_1}^{\bar{P}_2} \frac{d\pi}{d\bar{T}} + \frac{d\pi}{d\bar{P}}\ d\bar{P} d\bar{T}$$

Plugging in from Eq. 2.B.2 gives the final form of climate damages.

Chapter 3: Climate Change and Crop Choice: Evidence from the United States

Chapter Summary

In this chapter, we examine how farmers adjust the land they allocate to crops when climatic conditions change.[1] We analyze subnational longitudinal data for two globally important cereals from one of the world's major agricultural regions, the United States. Consistent with theory, we find that more land is allocated to producing a specific crop when the climatological conditions for growing that crop improve, although the timescales for these adjustments are long, on the order of 20 or more years. When we estimate the effectiveness of land use change as an adaptive strategy, we find that these adjustments reclaim 5-10% of climate-induced losses to producer surplus. The availability of adaptation via land use fails to generate more value because enduring climate-induced yield losses is less costly than altering land use practices for most farmers, indicating that land use adjustment costs may be high.

3.1 Introduction

Understanding how the global food supply responds to climatic changes has important implications for all societies. An extensive literature has examined how weather fluctuations affect agricultural production in the short-run (Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010; Welch et al., 2010; Lobell, Schlenker, and Costa-Roberts, 2011; Auffhammer, Ramanathan, and Vincent, 2012; Burke and Emerick, 2016; Gammans, Mérel, and Ortiz-Bobea, 2016; Fishman, 2016; Burgess et al., 2017; Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016) and has widely found substantial losses under future warming. However, these estimates hold farmer adaptive decisions fixed: understanding the full impact of a persistent change in the climate requires that we understand farmers' adjustments in the long-run (Mendelsohn, Nordhaus, and Shaw, 1994). Estimating the scope for farmer adaptation to slow-moving climatic changes has proven difficult (Auffhammer and Schlenker, 2014). Some evidence suggests that cost-effective adjustments are available in the long run (Mendelsohn, Nordhaus, and Shaw, 1994; Olmstead and Rhode, 2011); other evidence suggests that adaptive adjustments are too costly to employ (Schlenker and Roberts, 2009; Burke and Emerick, 2016), suggesting that short-run effects may be a good approximation for long-run effects. In this paper, we examine the historical scale and scope of a specific type of adaptation that is frequently raised by both academics and policy-makers: land-use changes that lead to the migration of crops.

[1] This material first appeared as a working paper of the same title, with authors Andrew Hultgren, Solomon Hsiang, David Lobell, Michael Roberts, and Wolfram Schlenker. We are grateful to Maximillian Auffhammer, Larry Karp, Jon Proctor, and Jerrod Welch for helpful comments.

It is often suggested that if climatic zones shift in space, the crops that perform well in those climatic zones will follow (Patt et al., 2010), yet this intuition has not been widely tested (Olmstead and Rhode, 2011). We wish to empirically test the extent to which this intuition holds in data; and if it does hold, to characterize the speed and effectiveness of this response. Specifically, we examine two major cereal crops at subnational scales in on of the world's major agricultural regions: the United States. We begin by estimating crop yield response functions for two key global grains, maize and soy, so that we can summarize historical climate changes in terms of their effects on yields. We then model changes in the area planted with each crop as a function of yield changes induced by climate. We implement this second step using two different approaches. First, we use the cross-section to examine whether trends in area planted reflect trends in climate-induced yield anomalies. Second, we use panel models to estimate the area planted response to an abrupt and permanent change in a climate-induced yield anomaly. Both of our approaches lead to the conclusion that farmers adjust their land use practices in response to climatic changes: if climate changes cause yields in a specific crop to decrease, then farmers will plant less of that crop. Further, these adjustments may large in the long run, with crop areas changing by as much as $\sim$20% in response to a 1% change in predicted yields. However, they tend to be slow, with adjustment periods ranging from 10 years to more than 25 years. This suggests that farmers face substantial frictions to land use adjustments, that suboptimal land use allocations may persist for decades, and that when adjustments do occur farmers may still face a substantial loss in surplus.

Our findings provide support for the notion that crop land distributions will shift in response to climate changes, and may shift at socially substantial scales. However our analysis points to a second important conclusion: on its own, the movement of crop lands does little to mitigate the welfare losses from climatic changes. Because the pecuniary and opportunity costs of agricultural production are variable across space, the first farmers to reduce their cultivated area in response to falling yields are marginal producers who have little to lose (or gain) by changing their land use practice. In contrast, inframarginal farmers facing adjustment costs may not find it optimal to change their land use practices, and may instead lose a portion of their surplus. Building on this insight, we apply our empirical estimates to calculate the welfare value of crop movements, and we find that long run land use changes only eliminate 5-10% of the losses from marginal climatic changes; the remaining 90-95% of losses persist unmitigated.

We begin by presenting *prima facie* evidence that farmers face adjustment costs that affect their land use decisions, preventing the distribution of crops from instantly and perfectly adjusting to every climatic shift.

3.1.1 Prima facie evidence from the global cross-section

Exogenous environmental factors, such as local climate, play an important role in the allocation of land to agricultural production. However, there are other constraints on agricultural land use, such as access to transportation or irrigation infrastructure, land protections based on cultural or biodiversity value, or geopolitical boundaries embedding competing policy-maker priorities. Furthermore, high opportunity costs of using land for agriculture may

prevent agricultural production on land that has good environmental properties. Because these costs and constraints may or may not be strongly linked to climate or other environmental factors, it is unclear whether land will be reallocated with respect to agriculture if climatological conditions change. In a frictionless world with no costs to land reallocation, we would expect every change in climate to be reflected in changes to agricultural land use, since farmers would quickly re-optimize their planting practices to maximize profits given a slightly adjusted climate. But if the costs and constraints to land use change are substantial, we expect to observe imperfect correlations between climatological suitability and agricultural land use, since incremental changes in climate will not always alter the optimality of a given planting practice.

The importance of non-climatological factors in determining land is easily observed in the global cross-section. Using high-resolution global data (Figure 3.A.1) on the suitability of land for agriculture (Ramankutty et al., 2008), and the fraction of land allocated to agriculture (Ramankutty and Foley, 1998), we treat every pixel as a separate observation and estimate a nonlinear regression of land fraction planted on crop suitability (Figure 3.1, top panel). We observe a highly statistically significant linear correlation, indicating that the influence of environmental factors on land use is clearly detectable. However, this high statistical significance does not imply that environmental factors explain all of the variation in land use. In the lower panel of Figure 3.1, we plot the raw data used to estimate this regression and it is immediately obvious that residual variations around the central tendency are large. It seems implausible that all of the residual variance is due to measurement error, suggesting that this unexplained variance in land use must be driven by factors other than climate and soil. This residual variance reflects the costs and constraints of agricultural land use.

3.2 Data

3.2.1 Climate data

We obtain temperature and precipitation data from the National Center for Atmospheric Research reconstruction (GFMD) which contains daily temperatures and precipitation for every $0.1^\circ \times 0.1^\circ$ pixel during 1950-2012 (Sheffield, Goteti, and Wood, 2006). For each pixel-crop pair, we compute the average temperature and precipitation rate during each month of the growing season[2] of each crop. We obtain pixel-specific growing season dates from the University of Wisconsin Center for Sustainability and the Global Environment (SAGE) (Sacks et al., 2010). Within each subnational unit, we then average growing season temperature and precipitation across pixels using as weights the distribution of croplands in 2000, also obtained from SAGE (Ramankutty et al., 2008). For small administrative units, this crop-weighting is relatively unimportant since these units are similar in scale to the size of GMFD pixels. However, for large administrative units, this weighting is important because crops may be present in only a small fraction of a large administrative unit.

[2]We define growing season to be from the first day of planting to the last day of the harvest.

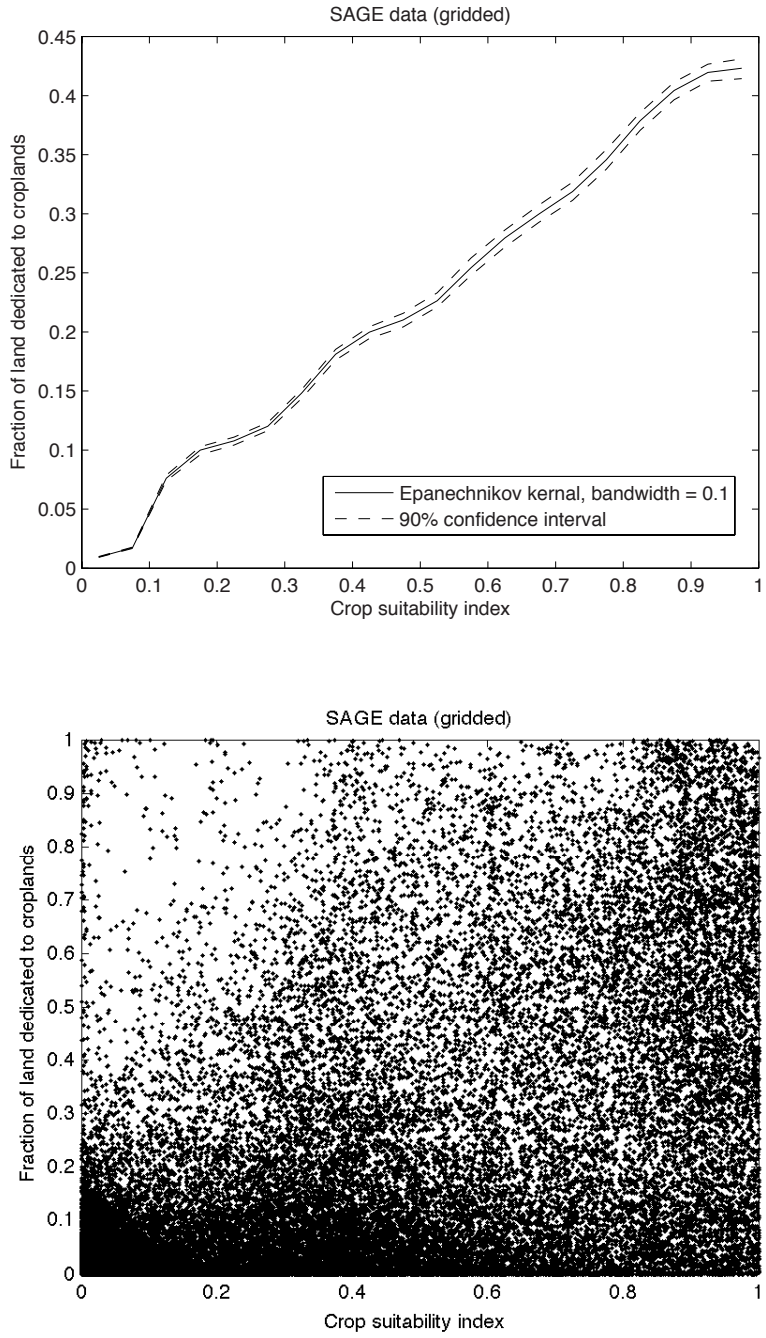

Figure 3.1: **Prima facie evidence of frictions to land use changes:** the strong but imperfect correlation between area planted and cropland suitability. Top: Non-linear cross-sectional regression (Nadaraya, 1964; Watson, 1964) of area cropped on a crop suitability index based on climate and soil (1 is best). Data same as in Figure 1. Standard errors are bootstrapped. Bottom: Despite the strong correlation between cropped area and cropland suitability, a large amount of residual variance in cropped area is unexplained by cropland suitability.

3.2.2 Agricultural data

We assemble subnational data files for U.S. counties from the USDA National Agricultural Statistics Service, and restrict our attention to the major crops maize and soy. We plot maps of cropped area fraction in Figure 3.2, where the coloring denotes the fraction of area cropped in 2000.

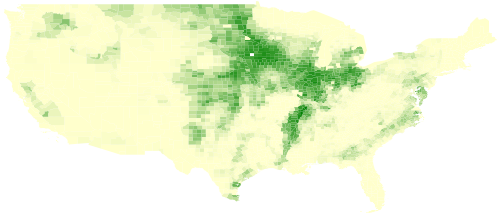

Figure 3.2: **Cropped area fraction.** Colors represent the fraction of area planted with any crop (Ramankutty et al., 2008).

3.3 First stage: Summarizing climate with yields

We begin our analysis by collapsing growing season climate variables, temperature and precipitation, into predicted crop yield estimates $\hat{Y}$. We do this for three reasons. First, temperature and precipitation both influence crop yields but may vary independently. However, because we are primarily interested in the land-use response to yields, it is unnecessarily difficult to keep track of both parameters independently when $\hat{Y}$ can be used to summarize their joint effects. Second, both temperature and precipitation influence yields in non-linear ways, making it difficult to compare the marginal effects of a 1C change in temperature or a 1 mm/month change in rainfall when different locations have different initial climates. It is more natural to compute the marginal effect of a 1% change in yields. Finally, different crops respond differently to temperature and precipitation changes, making across-crop comparisons of effect-sizes difficult if temperature and precipitation are retained as the independent variables. In contrast, it is intuitive to compare the effects of a 1% change in soy yields and a 1% change in maize yields.

We compute yield perturbations due to climatic fluctuations $\hat{Y}'_C$ with a flexible statistical model that we estimate from the data. We estimate percentage changes in yields around a location-specific baseline by using a model that contains administrative unit fixed effects (Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010). We also flexibly account for common trends, such as technological changes, and common shocks, such as the El Niño-Southern Oscillation, by including year fixed effects. For crop k and sample s, we estimate the model

$$Y_{it}^{ks} = f^{ks}(T_{it}^{k}) + g^{ks}(P_{it}^{k}) + \mu_i^{ks} + \gamma_t^{ks} + \epsilon_{it}^{ks} \tag{3.1}$$

where observations are indexed by their administrative unit i and year t. Temperature T^k and precipitation P^k are specific to the growing season of crop k. The effect of temperature and precipitation are unique to each crop and are described by $f^{ks}(.)$ and $g^{ks}(.)$, respectively. We model the effects of temperature and precipitation as additively separable for parsimony,

however we model both response functions flexibly using restricted cubic splines with six knots placed over the support of the data. Administrative unit fixed effects are μ and year fixed effects are γ. ϵ is a disturbance.

Because we describe Y using the natural logarithm of yields, our estimates $\hat{f}$ and $\hat{g}$ represent location-specific fractional deviations in yields due to temperature and precipitation. Figures 3.3 and 3.4 plot these estimated response functions for each crop. Consistent with earlier findings, yields tend to exhibit an optimum at moderate temperatures and moderate precipitation, with sharp declines occurring at high temperatures (Schlenker and Roberts, 2009; Schlenker and Lobell, 2010). More moderate declines tend to occur at low temperatures and very low or very high levels of precipitation.

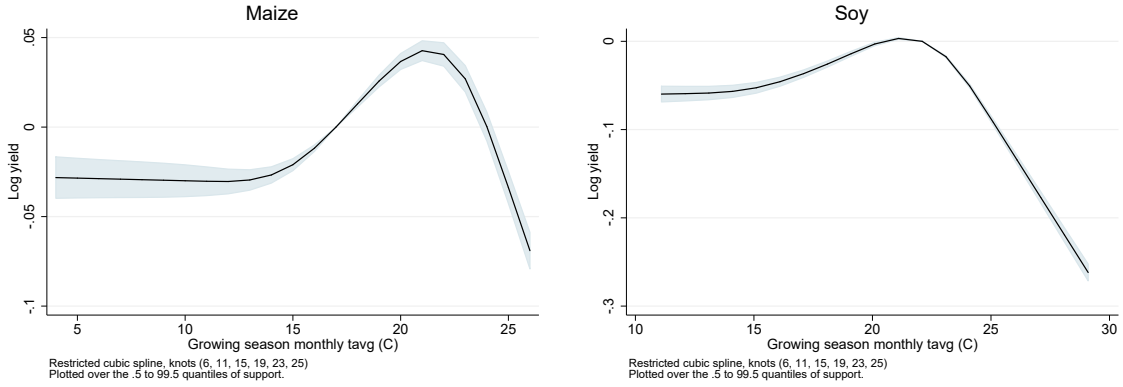

(a) (b)

Figure 3.3: **Crop yield responses to temperature.** Horizontal axis is average monthly temperature over the growing season of each crop in degrees Celsius. Vertical axis is log yields. Estimates are restricted cubic splines with a common set of six knots per weather variable. Grey region is $\pm 2\sigma$, standard errors robust to state-year correlated shocks.

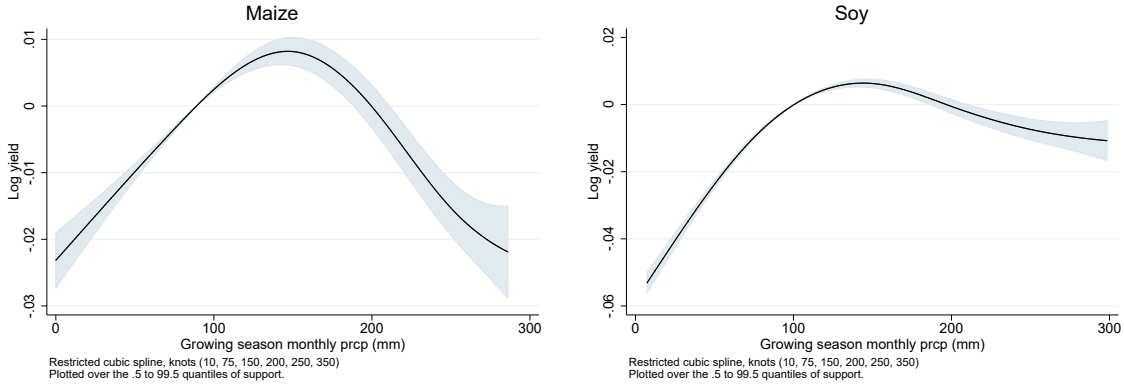

(a) (b)

Figure 3.4: **Crop yield responses to precipitation.** Horizontal axis is average monthly precipitation over the growing season of each crop in mm. Vertical axis is log yields. Estimates are restricted cubic splines with a common set of six knots per weather variable. Grey region is $\pm 2\sigma$, standard errors robust to state-year correlated shocks.

Because we would like to isolate the percentage yield deviations due only to climatological

fluctuations (Y'_C in Equation 3.B.2), we use only $\hat{f}$ and $\hat{g}$ to predict historical climate-driven deviations in yields by constructing

$$\hat{Y}^{\prime *}_{Cit} = \hat{f}^{ks}(T^k_{it}) + \hat{g}^{ks}(P^k_{it}). \tag{3.2}$$

We then use these estimates for $\hat{Y}^{\prime *}_{Cit}$ to identify R^{-1} in Equation 3.B.3.

We estimate climate-driven changes in the land allocated to farming crop using two approaches. First, we correlate climate driven trends in yields with trends in area planted. In this cross-sectional approach, we examine whether gradual changes in average yields are associated with gradual changes in area planted. In a second approach, we use the within administrative unit variation in yields to estimate a distributed lag model. Using this longitudinal approach, we explicitly estimate the transitional dynamics of land use under an abrupt and permanent change in climate.

3.4 Cross-sectional results: Land use trends vs. climate trends

In recent decades, many locations have exhibited substantial trends in yields that can be attributed to gradual changes in local climatological conditions (Lobell, Schlenker, and Costa-Roberts, 2011). We examine whether these gradual changes in yields are reflected in the gradual growth or contraction of cropped area. To do this, we first must estimate a location-specific trend in predicted yield perturbations $\hat{Y}^{\prime k}_{Cit}$, which we take to be exogenous. For each administrative unit and each crop, we estimate a linear yield trend $\hat{YT}^k_i$

$$\hat{Y}^{\prime k}_{Cit} = Y^k_{0i} + YT^k_i + \epsilon^k_{it}. \tag{3.3}$$

Because $\hat{YT}^k_i$ is in log points relative to the administrative unit fixed effect in Equation 3.1, YT is in units of "log points per year."

Next, we estimate the trend in area planted $\hat{AT}^k_i$ for the same administrative unit and crop

$$\hat{A}^k_{it} = A^k_{0i} + AT^k_i + \epsilon^k_{it}. \tag{3.4}$$

where A^k is the log of area planted with crop k. Similar to YT, the units of AT are "log points per year."

After estimating Equations 3.3 and 3.4, we have a single observation for the pair of trends for each location-crop pair (i, k). This new dataset is a cross-sectional sample of trend estimates $(\hat{YT}^k_i, \hat{AT}^k_i)$. Here it is important to note that for each sample s, we estimate the trend in predicted yields for the entire sample of years for which we have area planted data, but we also include the ten years preceding the agricultural sample. We do this because trends in climate will be difficult for farmers to detect above sampling noise, so if farmers are responding trends in their climate, they will likely base their cropping choices on trends over several preceding years. Thus each observation contains two trends, however $\hat{YT}$ is estimated for a window of time ten years longer than $\hat{AT}$.

Using the cross section of $(\hat{YT}_i^k,\ \hat{AT}_i^k)^s$ for crop k in sample s, we regress the trend in area planted on the trend in predicted yields:

$$AT_i^k = AT_0^{sk} + \beta^{sk} \times YT_i^k + \epsilon_i^k, \tag{3.5}$$

where $\hat{AT}_0^{sk}$ is the average trend in area planted for reasons unrelated to climate and β^{sk} is the marginal effect of a trending climate on trends in area planted, our parameter of interest.

If farmers decrease their area planted in response to a falling trend in their yields (Equation 3.B.1), we would expect β^{sk} to be positive. We acknowledge that Equation 3.5 describes a cross-sectional estimate, so it is possible that some unobserved omitted variable will be correlated with both trends in predicted yields and trends in area planted. However, we argue that systematic biases will be unlikely to persist across crops since the response function for each crop is a unique nonlinear function of both temperature and precipitation (Figures 3.3 and 3.4). Furthermore, secular trends in area planted, which might be driven by development or population growth, will be described by the constant term $\hat{AT}_0^{sk}$ and should not influence our estimate of β^{sk}.

We present ordinary least-squares (OLS) estimates of Equation 3.5 for both crops in Figure 3.5. We also present nonlinear estimates of both the mean and median tendencies. Both cross-sectional estimates exhibit a positive correlation as predicted by Equation 3.B.1. The estimated elasticity for soy (19.7, heteroskedasticity-robust s.e. 3.77) suggests a large response of area planted to trends in climate: if predicted yields rise by 1% per year then area planted is predicted to rise by 20% per year. The estimate for maize (6.2, heteroskedasticity-robust s.e. 2.27) suggests a still substantial, but more moderated, area response.

To put the magnitude of this response in perspective, we note that our estimated trends in predicted yields are on the order of $\pm 0.1\%$ per year (see histograms in Figure 3.5). Taking 0.1% per year as a representative trend in predicted yields, these estimates suggest that a representative administrative unit (with $\beta = 13$) would experience changes in area planted on the scale of $\pm 1.3\%$ per year due to climate changes.

Broadly speaking, while we acknowledge that these cross-sectional estimate are imperfectly identified, we find the consistency of positive correlations across crops as suggestive evidence that farmers respond to climate changes by adjusting their use of land (Equation 3.B.1). We now turn to a longitudinal approach that is better identified because it utilizes within-location variation in weather and area planted (Deschênes and Greenstone, 2007). It is less clear whether or not farmers respond to high-frequency variations in weather the same way that they respond to low-frequency trends in climate; however, if the responses are similar, the longitudinal approach has the added benefit that it allows us to explicitly estimate the rate at which farmers adjust their land use choices.

3.5 Longitudinal results: Land use vs. weather

We estimate the land use response to exogenous and idiosyncratic variations in weather by using a within-estimator similar to the approach used in Equation 3.1, but capturing the

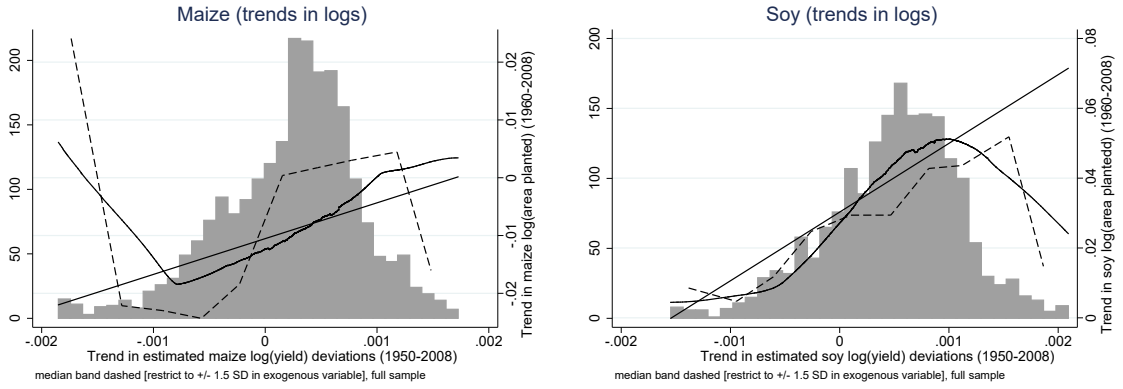

(a) (b)

Figure 3.5: **Trends in area planted vs. trends in yields.** Horizontal axis is the trend in log predicted yields due to climate (trend in $\hat{Y}'_C$), units are in log-points per year. Vertical axis is the trend in log area planted, units are log-points per year. All plots are cross-sectional, with each subnational unit represented only once. Samples are restricted to ± 1.5 standard deviations in the independant variable (grey histograms show distribution of estimates). Solid lines are mean regressions: OLS and local linear regressions (lowess). Dashed line is a median band plot.

dynamic response of farmers to a permanent shift in weather. We implement this approach by estimating the distributed lag model

$$A_{it}^{k} = \sum_{L=-3}^{20} \left[\psi_{t-L}^{sk} \times \hat{Y}'^{k}_{Ci,t-L} \right] \mu_i^{ks} + \rho_{pt}^{ks} + \epsilon_{it}^{ks},$$

where A is again log of area planted and $\hat{Y}'^{k}_{Ci,t-L}$ is the predicted perturbation in log yields (Equation 3.2) at a lag L relative to the current observation. μ is an administrative unit fixed effect that absorbs unobserved differences between locations and ρ is a state-year fixed effect that flexibly accounts for local secular trends, such as differential rates of technological diffusion, and common shocks, such as changes in agricultural trade policy. ϵ is a disturbance, which we allow to have an arbitrary correlation structure within a given state-year. We estimate Equation 3.5 separately for each crop and compute the sum of coefficients

$$\hat{\Omega}^{sk}(x) = \sum_{L=0}^{x} \hat{\psi}_{t-L}^{sk} \tag{3.6}$$

because we are interested in the effect of permanent climate changes. $\hat{\Omega}(x)$ represents the cumulative effect on area planted x years after an abrupt and permanent change in yields caused by an abrupt and permanent change in weather (that is, climate).

In a world with no frictions, we would expect that farmers instantly adjust to a permanent change in climate. This would imply that $\Omega(x)$ is discontinuous at zero and flat thereafter. However, if there are frictions that prevent farmers from instantly adjusting the area of land they farm, for example they may be uncertain whether an abrupt change in climate will persist, or there may be costly search and contracting requirements to adjust land use, then $\Omega(x)$ will be a more gradual function. If our estimates in the previous section were accurate

estimates of long-run adjustments, we would expect that $\hat{\beta}^{sk} = \hat{\Omega}^{sk}(\infty)$.

In Figure 3.6 we plot $\hat{\Omega}(x)$, the cumulative effect on area planted of an abrupt and permanent 1% yield increase, as estimated by applying an instrumental variables estimator of Equation 3.5, in which the lagged temperature and precipitation splines of Equation 3.1 are used as instruments for the lagged values of yields $\hat{Y}^{ık}_{Ci,t-L}$. Results for maize and soy both suggest that area planted exhibits a persistent increase in response to a permanent change in weather-induced yields. Responses, however, are more muted than the cross sectional estimates, with elasticities of 1-3% rather than 6-20%. This could be due to omitted variables in the cross-sectional estimates, or it could be that farmers receive a noisier signal from the variation used in the distributed lag estimator (a single year deviation in weather) versus that used in the cross-sectional estimator (a long run trend in climate). In addition, it should be noted that pre-trends present in the maize estimate suggest that temporally trending effects have not yet been fully controlled for.

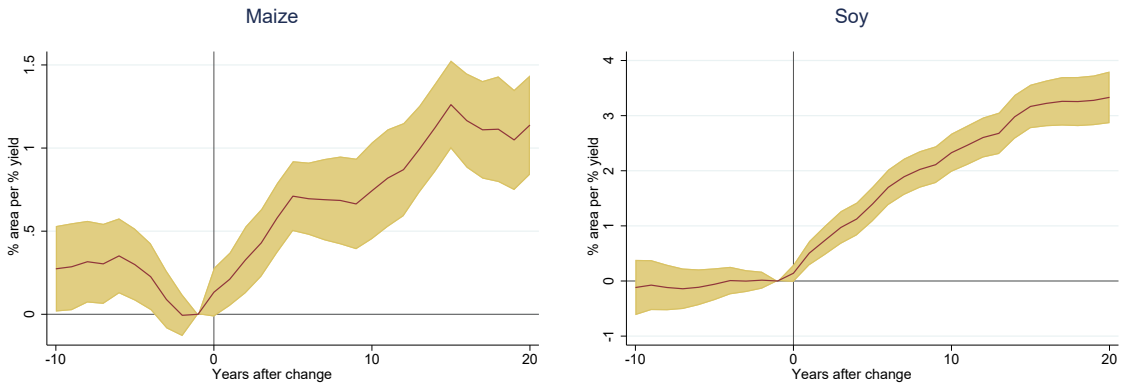

(a) (b)

Figure 3.6: **Transition dynamics of the land use response to a permanent yield change induced by climate.** Horizontal axis is years relative to a permanent 1% yield increase (induced by climate) that occurs in year zero. Vertical axis is the cumulative effect on area planted (% change) relative to the first lead year. Estimates are from Equation 3.5. Grey region is $\pm 2\sigma$, standard errors robust to state-year correlated shocks.

3.6 Projections

To illustrate the magnitude of the cropped area changes we find, we use the estimates of Equation 3.5 to project cropped area changes at the year 2050 under the RCP 8.5 warming scenario. To do this, we take temperature and precipitation projections from the CCSM4 climate model (Gent et al., 2011) and predict the associated yield trend for each administrative unit using 2010 to 2050 projected weather and the estimates in Figures 3.3 and 3.4. We then use the estimates in Figure 3.5 to predict the cropped area trend for each crop.

Figure 3.7 shows the projected cropped area changes by 2050 for each crop, as a percentage of the total administrative unit land area. Cropped area losses are widespread across the heavily cropped areas of the grain belt, southern Mississippi River, and Southeastern coastal states (compare to Figure 3.2), with the grain belt experiencing the largest cropped

area losses. Note that many parts of the U.S. do not grow maize or soy, and so the change in cropped area in these counties is zero.

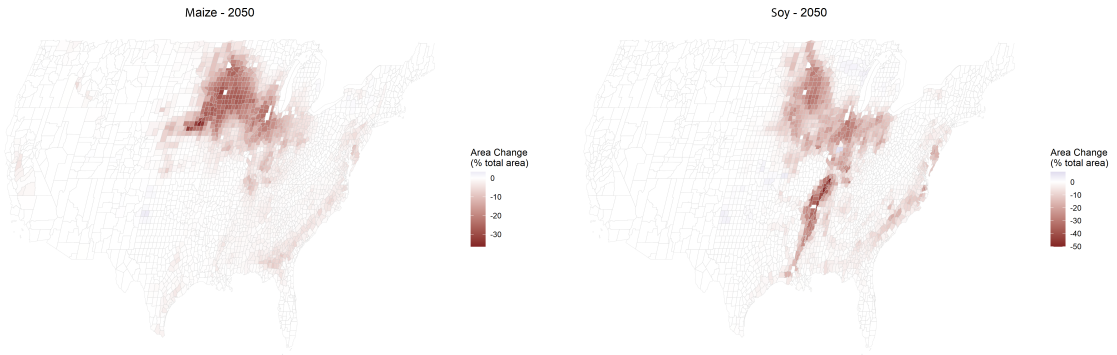

Figure 3.7: **Projections of crop migration under the RCP 8.5 warming scenario.** Projected changes in area cropped by 2050 under the RCP 8.5 warming scenario, using the CCSM4 climate model. Projections based on empirical estimates of Equation 3.5.

3.7 Discussion

Using data from one of the world's major agricultural regions, we measure changes in agricultural land use in response to climatic changes. We find that farmers respond to declining yields by reducing their planted area, and to increasing yields by increasing their planted area. That is, we find that cropped areas do qualitatively "migrate" with the climate, with substantial losses in cropped area projected for almost all maize- and soy-growing regions of the U.S.

However, quantitatively, we find that these adjustments have a limited impact on the welfare losses that producers suffer from climate change. In Equation 3.B.7 we present a simple statistic α to summarize the effectiveness of land use change as an adaptive strategy in response to climate changes: the fraction of gross climate-induced losses that are avoided by land use change. Recall that the mean elasticity across all estimates of Equation 3.5 is 13%; this implies a mean value of α across our estimates of 6.5%. This representative value suggests that if climatic changes caused a marginal reduction in producer surplus via yield losses, in the long run farmers would change their planting practices to avoid 6.5% of these damages. This is small number, however its small size is consistent with the basic intuition of how land use changes can mitigate climate change losses. Note that in Figure 3.B.1, the value of avoided damages was represented by the area of the triangle $c-d-e$. Because the change area planted $A_C'^*$ is small, this triangle ends up being small compared to the gross loss represented by rectangle $a-c-d-f$. The only situation in which α approaches a large value is if the slope of R becomes very shallow, however our empirical estimates suggest that the slope of R is substantially steeper. These results are consistent with analyses from other contexts that suggest the high costs of otherwise effective adaptive strategies limits their value for mitigating climate-induced losses (Hsiang and Narita, 2012; Anttila-Hughes and Hsiang, 2013).

This study is a partial equilibrium analysis, a fact that limits its scope in important ways. First, we have not explicitly modeled how farmers adjust when the costs of land use changes themselves respond to climate changes. This may be important because the opportunity cost of planting one crop may be the value from planting another crop which is also affected by the climate. Also, we have not accounted for responses in productive factors other than yields, such as changes in the labor market (Hornbeck, 2012; Feng, Krueger, and Oppenheimer, 2010; Graff Zivin and Neidell, 2014), which may respond to either climatic changes or land use changes. In addition, we have not accounted for possible changes in prices which may alter both producer and consumer surplus as well as affect farmers' land use decisions; this may be especially important in the dynamic models we estimate. Finally, we note that because the land use adjustments we estimate occur slowly, their value as an adaptive strategy declines if their benefits, which mainly accrue in the future, are discounted. Future work should address these issues.

Chapter Review and Looking Ahead

In this chapter, we examined a particular channel of climate adaptation in agricultural production; a channel which is widely discussed but for which there is little empirical evidence: the migration of cropped areas as climatic suitability shifts. Analyzing subnational longitudinal data for two globally important cereals from one of the world's major agricultural regions, we found that more land is allocated to producing a specific crop when the climatological conditions for growing that crop improve, although the timescales for these adjustments are long, on the order of 20 or more years. Further, we found that these adjustments reclaim only 5-10% of climate-induced losses to producer surplus. The availability of adaptation via land use fails to generate more value because enduring climate-induced yield losses is less costly than altering land use practices for most farmers, indicating that land use adjustment costs may be high.

In the next chapter, we will shift away from agriculture to examine how climate change and adaptation affect a different outcome of substantial social importance: human mortality.

References

Anttila-Hughes, Jesse and Solomon Hsiang. 2013. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." *SSRN Working Paper* URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2220501.

Auffhammer, Maximilian, V. Ramanathan, and Jeffrey R. Vincent. 2012. "Climate change, the monsoon, and rice yield in India." *Climatic Change* 111 (33):411–424. URL https://doi.org/10.1007/s10584-011-0208-4.

Auffhammer, Maximilian and Wolfram Schlenker. 2014. "Empirical studies on agricultural impacts and adaptation." *Energy Economics* 46:555–561.

Burgess, Robin, Olivier Deschênes, David Donaldson, and Michael Greenstone. 2017. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." *Working Paper* URL https://epic.uchicago.edu/wp-content/uploads/2019/07/Publication-9.pdf.

Burke, Marshall and Kyle Emerick. 2016. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8 (3):106–40. URL https://www.aeaweb.org/articles?id=10.1257/pol.20130025.

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

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

Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer. 2010. "Linkages among climate change, crop yields and Mexico–US cross-border migration." *Proceedings of the National Academy of Sciences* 107 (32):14257–14262. URL https://www.pnas.org/content/107/32/14257.

Fishman, Ram. 2016. "More uneven distributions overturn benefits of higher precipitation for crop yields." *Environmental Research Letters* 11 (2):024004. URL https://iopscience.iop.org/article/10.1088/1748-9326/11/2/024004/meta.

Gammans, Matthew, Pierre Mérel, and Ariel Ortiz-Bobea. 2016. "The impact of climate change on cereal yields: Statistical evidence from France." 2016 Annual Meeting, July 31-August 2, Boston, Massachusetts 236322, Agricultural and Applied Economics Association. URL https://ideas.repec.org/p/ags/aaea16/236322.html.

Gent, Peter R., Gokhan Danabasoglu, Leo J. Donner, Marika M. Holland, Elizabeth C. Hunke, Steve R. Jayne, David M. Lawrence, Richard B. Neale, Philip J. Rasch, Mariana Vertenstein, Patrick H. Worley, Zong-Liang Yang, and Minghua Zhang. 2011. "The Community Climate System Model Version 4." *Journal of Climate* 24 (19):4973–4991. URL https://doi.org/10.1175/2011JCLI4083.1.

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.

Hornbeck, Richard. 2012. "The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe." *American Economic Review* 102 (4):1477–1507. URL https://www.aeaweb.org/articles?id=10.1257/aer.102.4.1477.

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

Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate Trends and Global Crop Production Since 1980." *Science* 333 (6042):616–620. URL https://science.sciencemag.org/content/333/6042/616.

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.

Nadaraya, E. A. 1964. "On Estimating Regression." *Theory of Probability & Its Applications* 9 (1):141–142. URL https://doi.org/10.1137/1109020.

Olmstead, Alan L. and Paul W. Rhode. 2011. "Adapting North American wheat production to climatic challenges, 1839–2009." *Proceedings of the National Academy of Sciences* 108 (2):480–485. URL https://www.pnas.org/content/108/2/480.

Patt, Anthony G., Detlef P. van Vuuren, Frans Berkhout et al. 2010. "Adaptation in integrated assessment modeling: where do we stand?" *Climatic Change* 99 (3):383–402. URL https://doi.org/10.1007/s10584-009-9687-y.

Ramankutty, Navin, Amato T. Evan, Chad Monfreda, and Jonathan A. Foley. 2008. "Farming the planet: 1. Geographic distribution of global agricultural lands in the year 2000." *Global Biogeochemical Cycles* 22 (1). URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2007GB002952.

Ramankutty, Navin and Jonathan A. Foley. 1998. "Characterizing patterns of global land use: An analysis of global croplands data." *Global Biogeochemical Cycles* 12 (4):667–685. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/98GB02512.

Sacks, William J., Delphine Deryng, Jonathan A. Foley, and Navin Ramankutty. 2010. "Crop planting dates: an analysis of global patterns." *Global Ecology and Biogeography* 19 (5):607–620. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1466-8238.2010.00551.x.

Schlenker, Wolfram and David B Lobell. 2010. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5 (1):014010. URL https://iopscience.iop.org/article/10.1088/1748-9326/5/1/014010/meta.

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

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

Watson, Geoffrey S. 1964. "Smooth Regression Analysis." *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* 26 (4):359–372. URL http://www.jstor.org/stable/25049340.

Welch, Jarrod R., Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. 2010. "Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures." *Proceedings of the National Academy of Sciences* 107 (33):14562–14567. URL https://www.pnas.org/content/107/33/14562.

Appendix

3.A Global patters in crop suitability and area planted

The importance of non-climatological factors in determining land is easily observed in the global cross-section. In Figure 3.A.1 we plot high-resolution (5 minute $\times$ 5 minute) global maps of the suitability of land for agriculture, based on local soil and climate (Ramankutty et al., 2008), and the fraction of land allocated to agriculture (Ramankutty and Foley, 1998). Visual inspection suggests that the fraction of land dedicated to agriculture is correlated with the suitability of that land, however many regions with high suitability are not heavily farmed, and visa versa, suggesting the presence of sometimes substantial frictions. Figure 3.1 of the main text makes this comparison more precise.

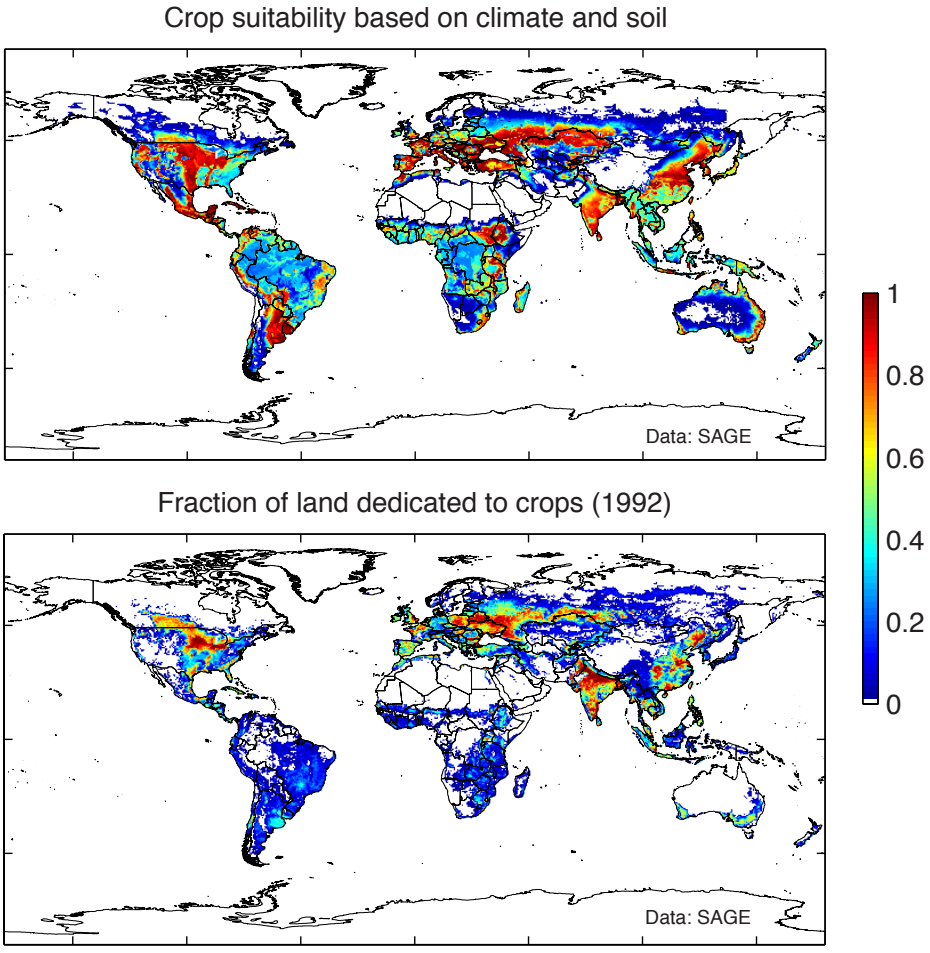

Figure 3.A.1: **Global patterns in crop suitability and area planted**. Global high-resolution (5" $\times$ 5") datasets of a crop suitability index based on climate and soil (1 is best) in 2000 (Ramankutty et al., 2008) and the fraction of land dedicated to any crop (Ramankutty and Foley, 1998). The latter is used to compute cropland area weights for weather variables.

3.B Theoretical framework

To clarify our analysis, we develop a simple framework for understanding how land is allocated to agricultural production. For simplicity, we focus on the producer's problem and all losses to climate are losses to producer surplus. This focus on producer welfare is consistent with our important simplifying assumption that prices remain fixed. If prices are fixed, consumers face no welfare changes and all climatological welfare losses are borne by producers. The assumption that prices are unaffected by local climatic changes is reasonable if local changes are idiosyncratic, however it becomes less tenable in scenarios where climate changes are spatially correlated on a global scale [Brunner, 2002, Hsiang et al., 2011]. However, for the sake of this analysis, we argue that assuming fixed prices is both reasonable and clarifying.

We assume one crop. Let yields Y by uniform over the location of interest. Let the total cost of cultivation, pecuniary cost plus opportunity cost, be R. Define the function $R(A)$ such that R is increasing in $A \in [0, 1]$, the fractional area planted. Note that we cannot observe $R(A)$ directly. For notational parsimony, let the price of crops be fixed equal to one. This framework is illustrated in Figure 3.

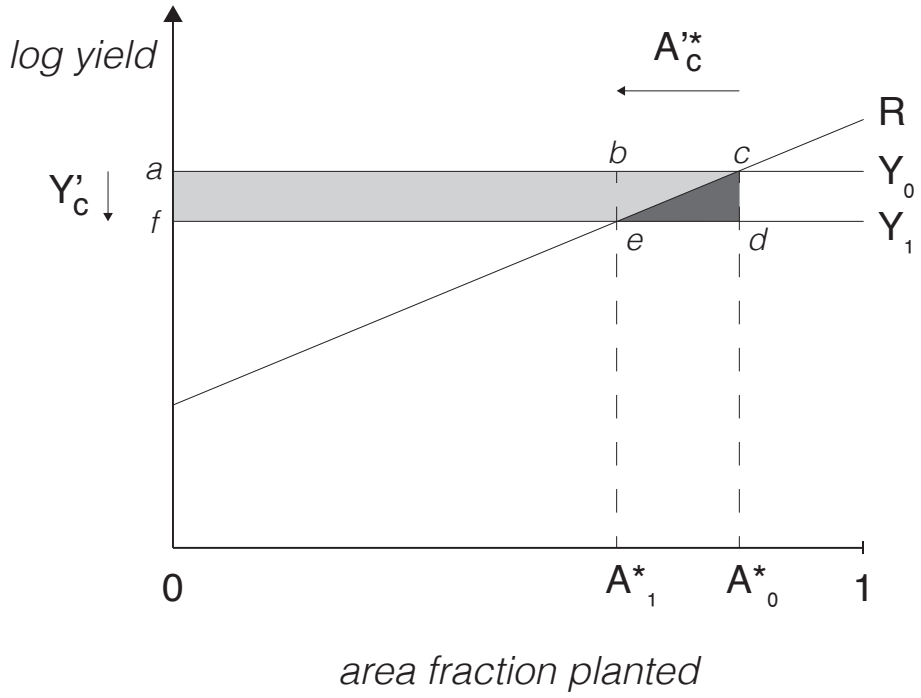

Figure 3.B.1: **Theoretical framework for evaluating land use changes in response to yield changes.** The shift in yields $Y_0 \to Y_1$ causes the adjustment in area planted $A_0^* \to A_1^*$ because the value of planting falls relative to the cost (pecuniary plus opportunity) of planting R.

Profit maximizing land use means that land will be cultivated only if the value of farming a hectare of land exceeds the cost of farming that hectare, that is $Y > R(A)$. Because production of the marginal hectare at A^* generates no profit, we have $Y = R(A^*)$ which can be inverted if we assume R is strictly increasing in A. Thus optimal land use requires

$$A^* = R^{-1}(Y). \tag{3.B.1}$$

This tells us that if yields change, then farmers will adjust A^*, the amount of land they plant. Therefore it is useful to consider the yield perturbation Y' which we decompose into a climatologically-driven component Y'_C and a residual component Y'_T that is driven by non-climatological changes, such as technology:

$$Y' = Y'_C + Y'_T. \tag{3.B.2}$$

Assuming that we can isolate Y'_C, we can define a land-use response that is driven only by climatological changes as

$$A'^*_C = R^{-1}(Y_0 + Y'_C) - R^{-1}(Y_0). \tag{3.B.3}$$

Similarly, we can define A'^*_T, the response in area planted to all non-climatological yield changes.

Because the pecuniary and opportunity cost of farming land may vary by location, the function R may vary across space. This means that even if yields are higher in one region, the area planted is not necessarily larger since R may also be higher. Thus, as we observed in the global cross-section, yields alone will not predict area planted. To predict the area planted we must also know the structure of R, the costs (and constraints) to farming.

The central exercise of this paper is to evaluate the structure of R near the current equilibrium $R(A^*)$. Essentially, we try to estimate the slope of R (actually R^{-1}) near $R(A^*)$. We do this by examining the response in area planted, A'^*_C, to variations in yields Y'_C that are driven by exogenous climatic changes. This allows us to infer the local structure of R because

$$\frac{\partial R^{-1}(Y_0)}{\partial Y_C} = \frac{A'^*_C}{Y'_C} = A'^*_C \tag{3.B.4}$$

where the last equality holds when we normalize the yield perturbation Y'_C to unity.

3.B.1 The effectiveness of land use change as adaptation to climate changes

The adjustment A'^*_C represents a type of adaptation to climatological changes that alter yields. This adaptive adjustment generates value relative to a counterfactual world without this adjustment. Assuming Y'_C is small so that R is approximately linear, the value generated by this adjustment is

$$V_{adapt} = \frac{Y'_C \times A'^*_C}{2}. \tag{3.B.5}$$

This can be seen in Figure 3, where the shift in yields $Y0 \rightarrow Y1$ causes the adjustment $A^*_0 \rightarrow A^*_1$. Following the change in yields, land area between A^*_0 and A^*_1 had costs of production R that exceeded the value of yields. So by stopping production over this area

and employing it for its next best use, the value in the triangle $c-d-e$ is recovered (V_{adapt}). Note that $V_{adapt} > 0$ if farmers always increase area planted when yields rise and visa versa.

Had there been no adjustment in area planted, all of the value in the rectangle $a-c-d-f$ would have been lost. We denote the value of this climatological change V_C:

$$V_C = Y'_C \times A'^*_0 \tag{3.B.6}$$

Where A^*_0 is the area planted prior to any climate changes. Since area planted is always positive, V_C always has the same sign as Y'_C, which may be either positive or negative.

By comparing V_{adapt} to V_C we obtain α, the fraction of climate change losses that are recovered through land use change:

$$\alpha = \frac{V_{adapt}}{V_C} = \frac{A'^*_C}{2A^*_0}. \tag{3.B.7}$$

Thus, for any climatological change, we can infer the fraction of value that is preserved by observing the land use response A'^*_C and comparing it to twice the original area planted. Note that α is a particularly convenient measure for the effectiveness of adaptation because it is normalized to gross climate change losses and because it is equal to one-half the elasticity of area planted to climate change.

Also note that if yields were to increase because the climate improved, say $Y_1 \rightarrow Y_0$ in Figure 3, then these results would continue to hold with this notation. However, the triangle of recovered value would be $b-c-e$ since land would be reallocated from non-farming uses to farming.

Chapter 4: Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits

Chapter Summary

This chapter develops the first globally comprehensive and empirically grounded estimates of mortality risk due to future temperature increases caused by climate change.[1] Using 40 countries' subnational data, we estimate age-specific mortality-temperature relationships that enable both extrapolation to countries without data and projection into future years while accounting for adaptation. We uncover a U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for the elderly, that is flattened by both higher incomes and adaptation to local climate (e.g., robust heating systems in cold climates and cooling systems in hot climates). Further, we develop a revealed preference approach to recover unobserved adaptation costs. We combine these components with 33 high-resolution climate simulations that together capture scientific uncertainty about the degree of future temperature change. Under a high emissions scenario, we estimate the mean increase in mortality risk is valued at roughly 3.2% of global GDP in 2100, with today's cold locations benefiting and damages being especially large in today's poor and/or hot locations. Finally, we estimate that the release of an additional ton of CO2 today will cause mean [interquartile range] damages of $36.6 [-$7.8, $73.0] under a high emissions scenario and $17.1 [-$24.7, $53.6] under a moderate scenario, using a 2% discount rate that is justified by US Treasury rates over the last two decades. Globally, these empirically grounded estimates substantially exceed the previous literature's estimates that lacked similar empirical grounding, suggesting that revision of the estimated economic damage from climate change is warranted.

[1]This material first appeared as a working paper of the same title, with authors Tamma Carleton, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert Kopp, Kelly E. McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang. This project is an output of the Climate Impact Lab that gratefully acknowledges funding from the Energy Policy Institute of Chicago (EPIC), International Growth Centre, National Science Foundation, Sloan Foundation, Carnegie Corporation, and Tata Center for Development. Tamma Carleton acknowledges funding from the US Environmental Protection Agency Science To Achieve Results Fellowship (#FP91780401). We thank Trinetta Chong, Greg Dobbels, Diana Gergel, Radhika Goyal, Simon Greenhill, Hannah Hess, Dylan Hogan, Azhar Hussain, Theodor Kulczycki, Brewster Malevich, Sébastien Annan Phan, Justin Simcock, Emile Tenezakis, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and Megan Landín, Terin Mayer, and Jack Chang for excellent project management. We thank David Anthoff, Max Auffhammer, Olivier Deschênes, Avi Ebenstein, Nolan Miller, Wolfram Schlenker, and and numerous workshop participants at University of Chicago, Stanford, Princeton, UC Berkeley, UC San Diego, UC Santa Barbara, University of Pennsylvania, University of San Francisco, University of Virginia, University of Wisconsin-Madison, University of Minnesota Twin Cities, NBER Summer Institute, LSE, PIK, Oslo University, University of British Columbia, Gothenburg University, the European Center for Advanced Research in Economics and Statistics, the National Academies of Sciences, and the Econometric Society for comments, suggestions, and help with data. Corresponding Author: Michael Greenstone; 1126 E. 59th Street, Chicago, IL 60637; Telephone: 773-702-8250; Email: mgreenst@uchicago.edu. Word count: 21,495.

4.1 Introduction

Understanding the likely global economic impacts of climate change is of tremendous practical value to both policymakers and researchers. On the policy side, decisions are currently made with incomplete and inconsistent information on the societal benefits of greenhouse gas emissions reductions. These inconsistencies are reflected in global climate policy, which is at once both lenient and wildly inconsistent. To date, the economics literature has struggled to mitigate this uncertainty, lacking empirically founded estimates of the economic damages from climate change. This problem is made all the more difficult because emissions today influence the global climate for hundreds of years, as Figure 4.1 illustrates. Thus, any reliable estimate of the damage from climate change must include long-run projections of economic impacts at global scale.

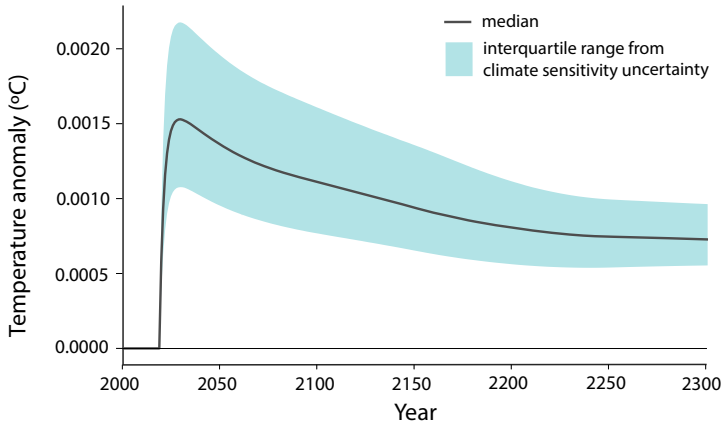

FIGURE 4.1

Temperature change due to a marginal emissions pulse in 2020 persists for centuries. The impact of a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) pulse of CO_2 in 2020 on temperature is shown. Median change is the difference in temperature of the "pulse" scenario relative to a high-emissions baseline scenario. The levels are anomalies in global mean surface temperature (GMST) in Celsius using our modification of the FAIR climate model. The shaded area indicates the inter-quartile range due to uncertainty in the climate's sensitivity to CO_2 (see Section 4.6 for details).

Decades of study have accumulated numerous insights and important findings regarding the economics of climate change, both theoretically and empirically, but a fundamental gulf persists between the two main types of analyses pursued. On the one hand, there are stylized models able to capture the global and multi-century nature of problem, such as "integrated assessment models" (IAMs) (e.g., Nordhaus, 1992; Tol, 1997; Stern, 2006), whose great appeal is that they provide an answer to the question of what the global costs of climate change will be. However, the many necessary assumptions of IAMs weaken the authority of these answers. On the other hand, there has been an explosion of highly resolved empirical analyses whose credibility lie in their use of real world data and careful econometric measurement (e.g., Schlenker and Roberts, 2009; Deschênes and Greenstone, 2007).[2] Yet these analyses tend to be limited in scope and rely on short-run changes in weather that might not fully account for adaptation to gradual climate change (Hsiang, 2016). At its

[2]For a comprehensive review, see Dell, Jones, and Olken (2014); Carleton and Hsiang (2016); Auffhammer (2018b).

core, this dichotomy persists because researchers have traded off between being complete in scale and scope or investing heavily in data collection and analysis. The result is that no study has delivered estimated effects of climate change that are at the scale of IAMs, while simultaneously being grounded in detailed econometric analyses using high-resolution globally representative micro-data.

The paper sets out to accomplish two major goals that require resolving the tension between these approaches in the context of mortality risk due to climate change. Specifically, it strives to provide the temporal and global scale of IAMs, but transparently built upon highly resolved econometric foundations. In so doing, it aims to account for both the benefits and costs of adaptation. The first goal is to produce local and global estimates of the mortality risk of climate change and its monetized value. The spatial resolution, based on dividing the world into 24,378 regions,[3] marks a substantial improvement upon existing IAMs, which represent (at most) 16 heterogeneous global regions (Tol, 1997). Using these calculations, we are able to accomplish the second goal, which is to estimate marginal willingness-to-pay (MWTP) to avoid the alteration of mortality risk associated with the release of an additional metric ton of CO_2. We call this the excess mortality "partial" social cost of carbon (SCC); a "full" SCC would encompass impacts across all affected outcomes.

In order to make these contributions, the analysis overcomes four fundamental challenges that have prevented the construction of empirically-derived and complete estimates of the costs of climate change to date. The first two of these challenges are due to the global extent and long timescale of both the causes and the impacts of climate change. For the first challenge, we note that CO_2 is a global pollutant, meaning that the costs of climate change must necessarily be considered at a global scale; anything less will lead to an incomplete estimate of the costs. The second challenge is that populations exhibit various levels of adaptation to current climate across space and adaptation levels are likely to be different in the future as populations become exposed to changes in their local climate. The extent to which investments in adaptation can limit the impacts of climate change is a critical component of cost estimates; ignoring this would lead to overstating costs.

We address both of these challenges simultaneously with a combination of extensive data and an econometric approach that models heterogeneity in the mortality-temperature relationship. We estimate this relationship using the most comprehensive dataset ever collected on annual, subnational mortality statistics from 41 countries that cover 55% of the global population. These data allow us to estimate age-specific mortality-temperature relationships with substantially greater resolution and coverage of the human population than previous studies; the most comprehensive econometric analyses to date have been for a few countries within a single region or individual cities from several countries. The analysis relies on inter-annual variation in temperature and uncovers a plausibly causal U-shaped relationship where extreme cold and hot temperatures increase mortality rates, especially for the elderly (those aged 65 and older).

We quantify the benefits of adaptation to gradual climate change and the benefits of projected future income growth by jointly modeling heterogeneity in the mortality-temperature response function with respect to the long-run climate (e.g., Auffhammer, 2018a) and income

[3]In the U.S., these impact regions map onto a county.

per capita (e.g., Fetzer, 2014). This cross-sectional modeling of heterogeneity allows us to predict the structure of the mortality-temperature relationship across locations where we lack data, as well as into the future, both of which are necessary to assess the global impacts of climate change. Such out-of-sample extrapolation of temperature-mortality relationships delivers the first empirically-based approach to including these populations in global climate impacts analysis, although a causal interpretation requires stronger econometric assumptions. Using readily available data and projections for current and future climate, income, and population projections, we estimate that the effect of an additional very hot day (35°C / 95°F) on elderly mortality is ∼50% larger in regions of the world where mortality data are unavailable. This finding underlines that current estimates may understate climate change impacts because they disproportionately rely on data from wealthy economies and temperate climates. Furthermore, accounting for changing mortality-temperature relationships is crucial to projecting the effect of warming in the future, as we expect the mortality consequences of heat to decline over time due to adaptations that individuals are predicted to undertake in response to warmer climates and higher incomes. Consistent with this intuition we find that climate adaptation and income growth have substantial benefits, marking a departure from previous literature that has often assumed that the mortality-temperature relationship was constant over space and time (e.g., Deschênes and Greenstone, 2011).

The third challenge is that the adaptation responses discussed in the previous paragraph are costly, and these costs, along with the direct mortality impacts, must be accounted for in a full assessment of climate change impacts. We develop a general revealed preference method to estimate the costs incurred to achieve the benefits from adapting to climate change, even though these costs cannot be directly observed. This is a critical step because a full accounting of the mortality-related costs of climate change necessarily accounts for the direct mortality impacts, the benefits of adaptation, and the opportunity costs of all resources deployed in order to achieve those adaptations. This is an advance on the previous literature that has either quantified adaptation benefits without estimating costs (e.g., Heutel, Miller, and Molitor, 2017) or tried to measure the costs of individual adaptations (e.g., Barreca et al., 2016). The latter approach is informative of individual costs, but poorly equipped to measure total adaptation costs, because the range of potential responses to warming – whether defensive investments (e.g., building cooling centers) or compensatory behaviors (e.g. exercising earlier in the morning) – is enormous, making a complete enumeration of their costs extraordinarily challenging.

The revealed preference approach is based on the assumption that individuals undertake adaptation investments when the expected benefits exceed the costs and that for the marginal investment, benefits and costs are equal. Because we can empirically observe adaptation benefits by measuring reduced mortality sensitivities to temperature, we can therefore infer their marginal cost. Then by integrating marginal costs, we can compute total costs for non-marginal climate changes. A simplified but illustrative example comes from comparing Seattle, WA and Houston, TX, which have similar income levels, but very different climates: on average, Seattle has less than 1 day per year where the daily average temperature exceeds 30°C (85°F), while Houston experiences over 8 of these days annually.[4] Our empirical analysis

[4]These values of *average* daily temperature are calculated from the GMFD dataset, described in Section 4.3.

below finds that Houston is relatively adapted to this hotter climate, with an individual hot day leading to much lower mortality than in temperate Seattle. By revealed preference, it must be the case that the costs required to achieve Houston-like adaptation exceeds the value of the lives Seattle would save by adopting similar practices. Similarly, the costs of adapting in Houston must be less than or equal to the value of the additional deaths that they would otherwise have to endure. These bounds shrink for smaller differences in climate (e.g. Seattle vs. Tacoma), such that we can show that for infinitesimally small differences in climate, these bounds collapse to a single value where the estimable marginal benefits and unobserved marginal costs are equal. Using this result, we are able to reconstruct non-marginal adaptation costs for each location, relying only on empirically recovered reduced-form estimates.

Together, addressing these three challenges allows us to achieve the first goal of this analysis: to develop measures of the full mortality-related costs of future climate change for the entire world, reflecting both the direct mortality costs accounting for adaptation and all adaptation costs. This exercise is done using 33 global climate models that together reflect current scientific uncertainty about the degree of temperature change[5] and results are expressed in "death equivalents", i.e., the number of deaths plus the adaptation costs incurred expressed in avoided deaths. We find that under a high emissions scenario (i.e., Representative Concentration Pathway (RCP) 8.5, in which CO_2 emissions growth is sustained) and a socioeconomic scenario with future global income and population growth rates approximately matching recent observations (i.e., Shared Socioeconomic Pathway (SSP) 3), the mean estimate of the total mortality burden of climate change is projected to be worth 85 death equivalents per 100,000, at the end of the century. Accounting for econometric and climate uncertainty leads to an interquartile range of [16, 121].[6] This is equal to roughly 3.2% of global GDP at the end of the century when death equivalents are valued using an age-varying value of a statistical life (VSL). Further, failing to account for income and climate adaptation would overstate the mortality costs of climate change by a factor of about 2.6, on average.

The analysis uncovers substantial heterogeneity in the full mortality costs of climate change around the globe. For example, mortality risk in Accra, Ghana is projected to increase by 19% of its current annual mortality rate at the end of the century under a high emissions scenario, while Oslo, Norway is projected to experience a decrease in mortality risk due to milder winters that is equal to 28% of its current annual mortality rate today (United Nations, 2017). Further, the share of the full mortality-related costs of climate change that are due to deaths, rather than adaptation costs, is 86% globally but varies greatly, with poor countries disproportionately experiencing impacts through deaths and wealthy countries disproportionately experiencing impacts through adaptation costs.

The last of the four challenges facing the literature is the calculation of an SCC that is based upon empirical evidence, the calculation of which is the second goal of our analysis. We develop and implement a framework to estimate the excess mortality partial SCC using empirically-based projections. The mortality partial SCC is defined as the marginal

[5]See Burke et al. (2015) for a discussion of combining physical climate uncertainty with econometric estimates.

[6]For reference, all cancers are responsible for approximately 125 deaths per 100,000 globally today (WHO, 2018). Of course, the full costs of cancer, including all adaptations incurred to avoid risk, would be even larger if expressed in death equivalents.

willingness-to-pay to avoid an additional ton of CO_2. A central element of this procedure is the construction of empirically grounded "damage functions," (Hsiang et al., 2017), each of which describes the costs of excess mortality risk in a given year as a function of the overall level of global climate change. Such damage functions have played a central role in the analysis of climate change as an economic problem since seminal work by Nordhaus (1992), but existing estimates have been criticized for having little or no empirical foundation (Pindyck, 2013). To our knowledge, ours is the first globally representative and empirically grounded partial damage function, enabling us to calculate a partial SCC when combined with a climate model.

Our estimates imply that the excess mortality partial SCC is roughly $36.6 (in 2019 USD) with a high emissions scenario (RCP8.5) under a 2% discount rate and using an age-varying VSL. This value falls to $17.1 with a moderate emissions scenario (i.e., Representative Concentration Pathway (RCP) 4.5, in which CO_2 emissions are stable through 2050 and then decline), due to the nonlinearity of estimated damage functions. The respective interquartile ranges are [-$7.8, $73.0] for RCP8.5 and [-$24.7, $53.6] for RCP4.5. These assumptions regarding discount rate and VSL are justified by US Treasury rates over the last two decades[7] and by standard economic reasoning regarding mortality risk valuation (Jones and Klenow, 2016; Murphy and Topel, 2006), respectively. However, under a higher discount rate of 3% and an age-invariant VSL, valuation assumptions used by the United States Government to set the SCC in 2009, the mortality partial SCC is approximately $22.1 [-$5.6, $53.4] under RCP8.5 and $6.7 [-$15.7, $32.1] under RCP4.5.

These empirically grounded estimates of the costs of climate-induced mortality risks substantially exceed available estimates from leading IAMs. For example, the total mortality-related damages from climate change under RCP8.5 in 2100 amount to about 49-135% of the comparable damages reported for *all sectors of the economy* in the IAMs currently determining the U.S. SCC (Interagency Working Group on Social Cost of Carbon, 2010). When considering the full discounted stream of damages from the release of an additional metric ton of CO_2, this paper's excess mortality partial SCC with a high emissions scenario amounts to ∼73% of the Obama Administration's *full* SCC (under a 2% discount rate and age-varying VSL); this value falls to ∼44% when using a 3% discount rate and age-invariant VSL.

The rest of this paper is organized as follows: Section 4.2 outlines a conceptual framework for the two key problems of the paper: projecting climate damages into the future, accounting for adaptation and its cost, and estimating a mortality partial SCC; Section 4.3 describes the data used in the estimation of impacts and in the climate change projected impacts; Section 4.4 details the econometric approach and explains how we extrapolate mortality impacts across space and project them over time while computing adaptation costs and benefits; Section 4.5 describes the results of the econometric analysis and presents global results from projections that use high-resolution global climate models; Section 4.6 details the calculation of a damage function based on these projections and computes a mortality partial SCC; Section 4.7 discusses limitations of the analysis; and Section 4.8 concludes.

[7]The average 10-year Treasury Inflation-Indexed Security value over the available record of this index (2003-present) is 1.01% (Board of Governors of the US Federal Reserve System, 2020).

4.2 Conceptual framework

Climate change is projected to have a wide variety of impacts on well-being, including altering the risk of mortality due to extreme temperatures. The ultimate effect on particular outcomes like mortality rates will be determined by the adaptations that are undertaken. Specifically, as the climate changes, individuals and societies will weigh the costs and benefits of undertaking actions that allow them to exploit new opportunities (e.g., converting land to new uses) and protect themselves against new risks (e.g., investments in air conditioning to mitigate mortality risks). The full cost of climate change, and hence the social cost of carbon, will thus reflect both the realized *direct impacts* (e.g., changes in mortality rates), which depend on the benefits of these adaptations, and *the costs of these adaptations in terms of foregone consumption*. However, to date it has proven challenging to develop a theoretically founded and empirically credible approach to explicitly recover the full costs of climate change and to incorporate such costs into an SCC.[8]

This section develops a simple framework to define (i) the full value of mortality risk due to climate change, and (ii) the mortality partial SCC, such that each reflects both the costs and benefits of adaptation. In both cases, we are able to derive expressions for these objects that are composed of terms that can be estimated with data.

4.2.1 Setup

We define the *climate* as the joint probability distribution over a vector of possible conditions that can be expected to occur over a specific interval of time. Following the notation of Hsiang (2016), let $\boldsymbol{C}_t$ be a vector of parameters describing the entire joint probability distribution over all relevant climatic variables in time period t (e.g., $\boldsymbol{C}$ might contain the mean and variance of daily average temperature and rainfall, among other parameters). The climate is determined both by natural variations in the climate system, and by the history of anthropogenic emissions. Thus, we write $\boldsymbol{C}_t = \varphi(E_0, E_1, E_2, ...E_t)$ where E_t represents total global greenhouse gas emissions in period t and $\varphi(\cdot)$ is a general function determined by the climate system that links past emissions to present climate.[9]

Define weather realizations as a random vector $\boldsymbol{c}_t$ drawn from a distribution characterized by $\boldsymbol{C}_t$. Emissions therefore influence realized weather by shifting the probability distribution, $\boldsymbol{C}_t$. Mortality risk is a function of both weather $\boldsymbol{c}$ and a composite good $\boldsymbol{b} = \xi(b_1, ..., b_K)$ comprising all b_k, where each b_k is an endogenous economic variable that influences adaptation. The composite $\boldsymbol{b}$ thus captures all adaptive behaviors and investments that interact with individuals' exposure to a warming climate, such as installation of air conditioning and time allocated to indoor activities. Mortality risk is captured by the probability of death during a unit interval of time $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$.

Consider a single representative global agent who derives utility in all time periods t from consumption of the numeraire good x_t and who faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$.[10] Because

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

[9] For more discussion, see Hsiang and Kopp (2018).

[10] Note that our empirical analysis relies on heterogeneous agents exposed to different climates, realizing different incomes,

weather realizations $\boldsymbol{c}_t$ are a random vector, this 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, conditional on the climate.[11] 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, where $\boldsymbol{c}(\boldsymbol{C})$ is a random vector $\boldsymbol{c}$ drawn from a distribution characterized by $\boldsymbol{C}$. This agent chooses their adaptations by solving:

$$\max_{\boldsymbol{b}_t, x_t} \; u(x_t) \left[1 - \tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t)\right] \;\; s.t. \; Y_t \geq x_t + A(\boldsymbol{b}_t), \tag{4.1}$$

where climate is determined by the history of emissions, where $A(\boldsymbol{b}_t)$ represents expenditures for all adaptive investments, and Y is an income we take to be exogenous.[12]

The following framework relies on a key set of assumptions. First, we assume that adaptation costs are a function of technology and do not depend on the climate. Additionally, we assume that $\tilde{f}(\cdot)$ is continuous and differentiable, that markets clear for all technologies and investments represented by the composite good $\boldsymbol{b}$, as well as for the numeraire good x, and that all choices $\boldsymbol{b}$ and x can be treated as continuous. Importantly, Equation 4.1 is static, because we assume that there is a competitive and frictionless rental market for all capital goods (e.g., air conditioners), so that fixed costs of capital can be ignored, and that all rental decisions are contained in $\boldsymbol{b}$. As long as agents have accurate expectations over their current climate, markets will clear efficiently in each period. Under these assumptions, the first order conditions of Equation 4.1 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.

4.2.2 The full value of mortality risk due to climate change

We first use the representative agent's problem in Equation 4.1 to derive an empirically tractable expression for the full value of mortality risk due to climate change. Before doing so, we highlight how this expression builds upon prior work quantifying the impacts of climate change.

Climate change will influence mortality risk $f = f(\boldsymbol{b}, \boldsymbol{c})$ through two pathways. First, a change in $\boldsymbol{C}$ will directly alter realized weather draws, changing $\boldsymbol{c}$. Second, as implied by Equation 4.1, a change in $\boldsymbol{C}$ will alter individuals' beliefs about their distribution of potential weather realizations, shifting how they act, and ultimately changing the optimal endogenous

and exhibiting different demographics. However, for expositional simplicity here we derive the full mortality risk of climate change and the mortality partial SCC using a globally representative agent.

[11] The assumption of exogenous emissions implies that we rule out the possibility that the agent will choose an optimal E^*. This is unrealistic with a single agent. But in practice, the world is comprises a continuum of agents or countries and the absence of coordinated global climate policy today is consistent with agents failing to choose optimal $\boldsymbol{b}^*$.

[12] The specification in Equation 4.1 imposes the assumption that there are no direct utility benefits or costs of adaptation behaviors or investments $\boldsymbol{b}$. In an alternative specification detailed in Appendix 4.A.4, we allow agents to derive utility both from x and from the choice variables in $\boldsymbol{b}$; for example, air conditioning may increase utility directly, in addition to lowering mortality risk. We show that under this alternative framework, the costs of adapting to climate change that we can empirically recover include pecuniary expenditures on adaptation, $A(\boldsymbol{b})$, *net* of any changes in direct utility benefits or costs. All other aspects of the framework presented here are unaffected. Additionally, in a variant of this model in which agents derive utility directly from the climate, the interpretation of empirically recovered adaptation costs is modified to include an additional component representing changes in utility derived directly from the changing climate. However, this change of interpretation to include another term that is "netted out" in estimated adaptation costs is the only implication of adding climate directly to the utility function.

choice of $\boldsymbol{b}^*$. Therefore, since the climate $\boldsymbol{C}$ influences both $\boldsymbol{c}$ and $\boldsymbol{b}^*$, the probability of death at the initial time period $t = 1$ is:

$$\Pr(\text{death} \mid \boldsymbol{C}_1) = f(\boldsymbol{b}^*(Y_1, \boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) \tag{4.2}$$

Many previous empirical estimates of the effects of climate assume no adaptation takes place (e.g., Deschênes and Greenstone, 2007; Hsiang et al., 2017), such that projections of future impacts are computed assuming economic decisions embodied by $\boldsymbol{b}$ do not change. In reality, optimizing agents will update their behaviors and technologies $\boldsymbol{b}$ to attenuate climate-induced increases in mortality risk. Several analyses have empirically confirmed that accounting for endogenous changes to technology, behavior, and investment mitigates the direct effects of climate in a variety of contexts (e.g., Barreca et al., 2016; Park et al., 2020).[13] However, existing climate change projections accounting for these adaptation benefits do not account for the *costs* of adaptation, i.e., $A(\boldsymbol{b})$.

A full measure of the economic burden of climate change must account not only for the benefits generated by adaptive reactions to these changes but also their cost. Thus, the total cost of changing mortality risks that result from climate change between time periods $t = 1$ and $t = 2$ is:

$$\begin{aligned} &\textit{full value of mortality risk due to climate change} = \\ &\qquad VSL_2 \underbrace{[f(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2), \boldsymbol{c}(\boldsymbol{C}_2)) - f(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1))]}_{\text{observable change in mortality rate}} + \\ &\qquad\qquad \underbrace{A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1))}_{\text{adaptation costs}}, \end{aligned} \tag{4.3}$$

where VSL_2 is the value of a statistical life in time period 2, or the willingness to pay for a marginal increase in the probability of survival, and is used to convert mortality risk to dollar value (Becker, 2007). Importantly, this definition includes changes in mortality risk and adaptation costs due *only* to changes in the climate, as income Y is held fixed at its $t = 2$ level. This ensures, for example, that increases in air conditioning prevalence due to rising incomes are not included in adaptation benefits or costs of climate change. Note that the omission of the costs of adaptation, $A(\boldsymbol{b})$, would underestimate the overall economic burden of warming.

The first key objective of this paper is to empirically quantify the total costs of climate change impacts on mortality risk, following Equation 4.3. However, the changes in adaptation costs between time periods (second term in Equation 4.3) are unobservable, practically speaking. In principle, data on each adaptive action could be gathered and modeled (Deschênes and Greenstone, 2011, e.g.,), but since there exists an enormous number of possible adaptive margins that together make up the composite good $\boldsymbol{b}$, computing the full cost of climate change using such an enumerative approach quickly becomes intractable.

To circumvent this challenge, we use a revealed preference approach derived from the first order conditions of the agents' problem (Equation 4.1) to construct empirical estimates

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

of changes in adaptation costs due to climate change. We begin by rearranging the agent's first order conditions and using the conventional definition of the VSL (i.e., $\frac{u(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x}$ (Becker, 2007)) to show that in any time period t,

$$\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}} = \frac{-u(x_t^*)}{\partial u/\partial x[1-\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = -VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \tag{4.4}$$

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 infer estimates of marginal adaptation costs.

To make the expression in Equation 4.4 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:

$$\frac{d\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} = \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \tag{4.5}$$

The first term on the righthand side of Equation 4.5 represents the expected impacts on mortality of all changes in adaptive investments induced by the change in climate; as discussed, this is of limited practical value because of data and estimation limitations.[14] The second term is the direct effect that the climate would have if individuals did not adapt (i.e., the partial derivative).[15] For example, if climate change produces an increase in the frequency of heat events that threaten human health, it would be natural to expect the first term to be negative, as people make adjustments that save lives, and the second term to be positive, reflecting the impacts of heat on fatalities absent those adjustment. Equation 4.5 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.

We use this fact in combination with Equation 4.4 to develop an expression for the *total* adaptation costs incurred as the climate changes gradually from $t = 1$ to $t = 2$, which is composed of elements which can be estimated:[16]

$$\begin{aligned} A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1)) &= \int_1^2 \frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_t}{dt} dt \\ &= -\int_1^2 VSL_t \left[\frac{d\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_t}{dt} dt \end{aligned} \tag{4.6}$$

The practical value of Equation 4.6 is that it outlines how we can use estimates of the total and partial derivatives of mortality risk—with respect to the climate—to infer net adaptation costs, even though adaptation itself is not directly observable. In the following sections,

[14] 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.

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

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

we develop an empirical panel model exploiting both short-run and long-run variation in temperature through which the total derivative $\frac{d\tilde{f}}{d\boldsymbol{C}}$ can be separated from the partial derivative $\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$. The details of implementing Equation 4.6 are discussed in Section 4.4.5 and we empirically quantify these values globally in Section 4.5.3.

Before proceeding, a few details are worth underscoring. First, while we integrate over changes in climate in Equation 4.6, we hold income fixed at its endpoint value. This is because the goal is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate only. In contrast, changes in expenditures due to rising income will alter mortality risk under climate change, but are not a consequence of the changing climate; therefore not included in our calculation of the total mortality-related costs of climate change.

Second, the revealed preference approach for recovering adaptation costs relies on the first order condition that guarantees that *marginal* costs of adaptation are equal to *marginal* benefits at the optimal choice $\{x^*, \boldsymbol{b}^*\}$. Since we can estimate marginal benefits, we can back out marginal costs.

Third, the *total* adaptation costs associated with the climate shifting from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ are calculated by integrating marginal benefits of adaptation for a series of infinitesimal changes in climate (Equation 4.6), where marginal benefits continually evolve with the changing climate $\boldsymbol{C}$. Thus, total adaptation costs in a given period, relative to a base period, are the sum of the adaptation costs induced by a series of small changes in climate in the preceding periods (see Appendix 4.A.1 for a visual description).

Finally, the *total* adaptation benefits associated with the climate shifting from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ are defined as the dollar value of the difference between the effects of climate change with optimal adaptation and without any adaptation: $-VSL_2[\tilde{f}(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2), \boldsymbol{C}_2) - \tilde{f}(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1), \boldsymbol{C}_2)]$. In contrast to total adaptation costs, this expression relies on the relationship between mortality and temperature that holds *only* at the final climate, $\boldsymbol{C}_2$. Therefore, when the marginal benefits of adaptation are greater at the final climate than at previous climates, the total benefits of adaptation will exceed total adaptation costs, generating an adaptation "surplus". For example, at a climate between $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$, the marginal unit of air conditioning (a key form of adaptation) purchased will have benefits that are exactly equal to its costs. However, at the warmer climate $\boldsymbol{C}_2$, this same unit of air conditioning becomes inframarginal, and may have benefits that exceed its costs. Appendix 4.A.2 derives a formal expression for this adaptation surplus.

4.2.3 The mortality partial social cost of carbon

The second objective of this paper is to quantify the mortality partial SCC. Here, we use the representative agent's problem in Equation 4.1 to derive an expression for the partial SCC, which we empirically estimate using a procedure outlined in Section 4.6.

Given the agent's expectations of the weather, the indirect utility function for the problem in Equation 4.1 in each period t is:

$$v_t(Y_t, A(\boldsymbol{b}_t^*), \boldsymbol{C}_t) = u(x_t^*)\left[1 - \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)\right] + \lambda_t[Y_t - x_t^* - A(\boldsymbol{b}_t^*)], \tag{4.7}$$

where, as above, climate is determined by the prior evolution of global emissions through $\boldsymbol{C}_t = \varphi(E_0, E_1, E_2, ...E_t)$ and λ_t is the marginal utility of income. The mortality partial SCC

is defined as the marginal willingness-to-pay (MWTP) in period t to avoid the mortality consequences from a marginal increase in emissions. Because emissions released in period t influence the trajectory of global emissions for hundreds of years (see Figure 4.1), this MWTP includes impacts of carbon emissions on utility in future time periods $s > t$. Thus, we derive the MWTP to avoid a marginal change in emissions in period t by differentiating the indirect utility function in each future period s with respect to emissions E_t, and integrating over time:[17]

$$\begin{aligned} \textit{Mortality partial } SCC_t \text{ (utils)} &= \int_t^\infty e^{-\delta(s-t)} \frac{-dv_s}{dE_t} ds \\ &= \underbrace{\int_t^\infty e^{-\delta(s-t)} u(x_s^*) \left(\frac{\partial \tilde{f}_s}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}_s}{\partial \boldsymbol{C}} \right) \frac{\partial \boldsymbol{C}_s}{\partial E_t} ds}_{\substack{\text{discounted damages of emissions from} \\ \text{change in mortality rates}}} + \\ &\quad \underbrace{\int_t^\infty e^{-\delta(s-t)} \lambda_s \frac{\partial A_s}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \frac{\partial \boldsymbol{C}_s}{\partial E_t} ds}_{\substack{\text{discounted damages of emissions} \\ \text{from adaptation costs}}}, \end{aligned} \tag{4.8}$$

where δ indicates the discount rate.[18]

The two terms in Equation 4.8 show that increases in emissions cause damages through two channels. First, emissions change mortality rates, net of optimal adaptation, for all future time periods. Second, emissions change the expenditures that the agent must incur in order to update her optimal adaptation.

Some manipulation allows us to convert Equation 4.8 into dollars. Using the standard definition of the VSL and the first order conditions from Equation 4.1, we divide Equation 4.8 by $[1 - \tilde{f}]\partial u/\partial x$ to rewrite the mortality partial SCC in units of dollars:

$$\begin{aligned} \textit{Mortality partial } SCC_t \text{ (dollars)} &= \int_t^\infty e^{-\delta(s-t)} \underbrace{\left[VSL_s \left(\frac{\partial \tilde{f}_s}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} + \frac{\partial \tilde{f}_s}{\partial \boldsymbol{C}} \right) + \frac{\partial A_s}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_s^*}{\partial \boldsymbol{C}} \right]}_{\substack{\text{total monetized mortality-related damages} \\ \text{from a marginal change in climate}}} \frac{\partial \boldsymbol{C}_s}{\partial E_t} ds \\ &= \int_t^\infty e^{-\delta(s-t)} \frac{dD(\boldsymbol{C}_s, s)}{d\boldsymbol{C}} \frac{\partial \boldsymbol{C}_s}{\partial E_t} ds, \end{aligned} \tag{4.9}$$

where $D(\boldsymbol{C}_s, s)$ represents a damage function describing *total* global economic losses due to changes in mortality risk in year s, as a function of the global climate $\boldsymbol{C}$. Critically, this damage function is inclusive of adaptation benefits and costs.

In practice, we approximate Equation 4.9 by combining empirically grounded estimated damage functions $D(\cdot)$ with climate model simulations of the impact of a small change in emissions on the global climate, i.e., $\frac{\partial \boldsymbol{C}_s}{\partial E_t}$. Expressing the mortality partial SCC using a

[17] For display purposes only we have omitted the arguments of $\tilde{f}(\cdot)$ in Equations 4.8 and 4.9.

[18] Equation 4.8 assumes a constant discount rate δ. This approach is taken because it is standard in policy applications of the SCC (Interagency Working Group on Social Cost of Carbon, 2010), although future work should explore the implications of more complex discounting procedures, such as declining discount rates (e.g., Newell and Pizer, 2004; Millner and Heal, 2018).

damage function has three key practical advantages. First, the damage function represents a parsimonious, reduced-form description of the otherwise complex dependence of global economic damage on the global climate. Second, as we demonstrate in Section 4.6, it is possible to empirically estimate damage functions from the climate change projections described in Section 4.2.2. Finally, because they are fully differentiable, empirical damage functions can be used to compute *marginal* costs of an emissions impulse released in year t by differentiation. The construction of these damage functions, as well as the implementation of the mortality partial SCC, are detailed in Section 4.6.

4.3 Data

We believe that we have collected the most comprehensive data file ever compiled on mortality, historical climate data, and climate, population, and income projections. Section 4.3.1 describes the data necessary to estimate the relationship between mortality and temperature. Section 4.3.2 outlines the data we use to predict the mortality-temperature relationship across the entire planet today and project its evolution into the future as populations adapt to climate change. Appendix 4.B provides a more extensive description of all of these datasets.

4.3.1 Data to estimate the mortality-temperature relationship

4.3.1 Mortality Data.

Our mortality data are collected independently from 41 countries.[19] Combined, this dataset covers mortality outcomes for 55% of the global population, representing a substantial increase in coverage relative to existing literature; prior studies investigate an individual country (e.g., Burgess et al., 2017) or region (e.g., Deschenes, 2018), or combine small nonrandom samples from across multiple countries (e.g., Gasparrini et al., 2015). Spatial coverage, resolution, and temporal coverage are shown in Figure 4.2A, and each dataset is summarized in Table 4.1 and detailed in Appendix 4.B.1. We harmonize all records into a single multi-country panel dataset of age-specific annual mortality rates, using three age categories: <5, 5-64, and >64, where the unit of observation is ADM2 (e.g., a county in the U.S.) by year. Note that the India mortality data lacks age-specific rates, and therefore is used for out-of-sample tests rather than in the main analysis.

4.3.2 Historical Climate Data.

We perform analyses with two separate groups of historical data on precipitation and temperature. First, we use the Global Meteorological Forcing Dataset (GMFD) (Sheffield, Goteti, and Wood, 2006), which relies on a weather model in combination with observational data. Second, we repeat our analysis with climate datasets that strictly interpolate observational data across space onto grids, combining temperature data from the daily Berkeley Earth

[19] Our main analysis uses age-specific mortality rates from 40 of these countries. We use data from India as cross-validation of our main results, as the India data do not have records of age-specific mortality rates. The omission of India from our main regressions lowers our data coverage to 38% of the global population.

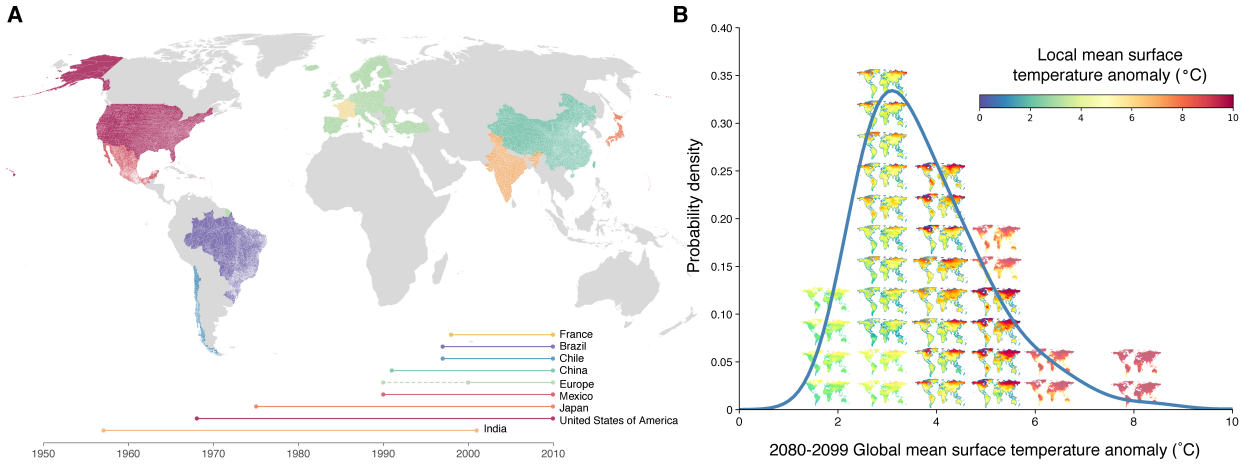

FIGURE 4.2

Mortality statistics and future climate projections used in generating empirically-based climate change mortality impact projections. Panel A shows the spatial distribution and resolution of mortality statistics from all countries used to generate regression estimates of the temperature-mortality relationship. Temporal coverage for each country is shown under the map (the dotted line for the European Union (EU) time series indicates that start dates vary for a small subset of countries). Panel B shows the 21 climate models (outlined maps) and 12 model surrogates (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.

Surface Temperature dataset (BEST) (Rohde et al., 2013) with precipitation data from the monthly University of Delaware dataset (UDEL) (Matsuura and Willmott, 2007). Table 4.1 summarizes these data; full data descriptions are provided in Appendix 4.B.2. We link climate and mortality data by aggregating gridded daily temperature data to the annual measures at the same administrative level as the mortality records using a procedure detailed in Appendix 4.B.2.4 that preserves potential nonlinearities in the mortality-temperature relationship.

4.3.3 Covariate Data.

Our analysis allows for heterogeneity in the age-specific mortality-temperature relationship as a function of two long-run covariates: a measure of climate (in our main specification, long-run average temperature) and income per capita. We assemble time-invariant measures of both these variables at the ADM1 unit (e.g., state) level using GMFD climate data and a combination of the Penn World Tables (PWT), Gennaioli et al. (2014), and Eurostat (2013). The construction of the income variable requires some estimation to downscale to ADM1 level; details on this procedure are provided in Appendix 4.B.3.

Mortality records

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

Historical climate datasets

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

TABLE 4.1
Historical mortality & climate data

*In units of deaths per 100,000 population.
†To remove outliers, particularly in low-population regions, we winsorize the mortality rate at the 1% level at high end of the distribution across administrative regions, separately for each country.
□ All covariate values shown are averages over the years in each country sample.
× ADM2 refers to the second administrative level (e.g., county), while ADM1 refers to the first administrative level (e.g., state). NUTS2 refers to the Nomenclature of Territorial Units for Statistics 2^{nd} (NUTS2) level, which is specific to the European Union (EU) and falls between first and second administrative levels.
◇ Global population share for each country in our sample is shown for the year 2010.
⊗ GDP per capita values shown are in constant 2005 dollars purchasing power parity (PPP).
⊘ Average daily temperature and annual average of the number of days above 28°C are both population weighted, using population values from 2010.
‡ EU data for 33 countries were obtained from a single source. Detailed description of the countries within this region is presented in Appendix 4.B.1.
▷ Most countries in the EU data have records beginning in the year 1990, but start dates vary for a small subset of countries. See Appendix 4.B.1 and Table 4.B.1 for details.
⊕ We separate France from the rest of the EU, as higher resolution mortality data are publicly available for France.

^ It is widely believed that data from India understate mortality rates due to incomplete registration of deaths.

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

4.3.1 Unit of Analysis for Projections.

We partition the global land surface into a set of 24,378 regions onto which we generate location-specific projected damages of climate change. These regions (hereafter, *impact regions*) are constructed such that they are either identical to, or are a union of, existing administrative regions. They (i) respect national borders, (ii) are roughly equal in population across regions, and (iii) display approximately homogenous within-region climatic conditions. Appendix 4.C details the algorithm used to create impact regions.

4.3.2 Climate Projections.

We use a set of 21 high-resolution bias-corrected, downscaled global climate projections produced by NASA Earth Exchange (NEX) (Thrasher et al., 2012)[20] that provide daily temperature and precipitation through the year 2100.[21] We obtain climate projections based on two standardized emissions scenarios: Representative Concentration Pathways 4.5 (RCP4.5, an emissions stabilization scenario) and 8.5 (RCP8.5, a scenario with intensive growth in fossil fuel emissions) (Van Vuuren et al., 2011; Thomson et al., 2011)).

These 21 climate models systematically underestimate tail risks of future climate change (Tebaldi and Knutti, 2007; Rasmussen, Meinshausen, and Kopp, 2016).[22] To correct for this, we follow Hsiang et al. (2017) by assigning probabilistic weights to climate projections and use 12 surrogate models that describe local climate outcomes in the tails of the climate sensitivity distribution (Rasmussen, Meinshausen, and Kopp, 2016). Figure 4.2B shows the resulting weighted climate model distribution. The 21 models and 12 surrogate models are treated identically in our calculations and we describe them collectively as the surrogate/model mixed ensemble (SMME). Gridded output from these projections are aggregated to impact regions; full details on the climate projection data are in Appendix 4.B.2.

4.3.3 Socioeconomic Projections.

Projections of population and income are a critical ingredient in our analysis, and for these we rely on the Shared Socioeconomic Pathways (SSPs), which describe a set of plausible scenarios of socioeconomic development over the 21^{st} century (see Hsiang and Kopp (2018) for a description of these scenarios). We use SSP2, SSP3, and SSP4, which yield emissions in the absence of mitigation policy that fall between RCP4.5 and RCP8.5 in integrated assessment modeling exercises (Riahi et al., 2017). For population, we use the International Institute for Applied Systems Analysis (IIASA) SSP population projections, which provide estimates of population by age cohort at country-level in five-year increments (IIASA Energy Program, 2016). National population projections are allocated to impact regions based on current satellite-based within-country population distributions from Bright et al. (2012) (see Appendix 4.B.3.3). Projections of national income per capita are similarly derived from the SSP scenarios, using both the IIASA projections and the Organization for Economic Cooperation and Development (OECD) Env-Growth model (Dellink et al., 2015) projections. We allocate national income per capita to impact regions using current nighttime light satellite imagery from the NOAA Defense Meteorological Satellite Program (DSMP). Appendix 4.B.3.2 provides details on this calculation.

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

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

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

4.4 Methods

Here we describe a set of methods designed to generate future projections of the impacts of climate change on mortality across the globe, relying on empirically estimated historical relationships. In the first subsection, we detail the estimating equation used to recover the average treatment effect of temperature on mortality rates across all administrative regions in our sample. This gives us the casual estimate of temperature's impact upon mortality using historical data. In the second subsection, we describe a model of heterogeneous treatment effects that allows us to capture differences in temperature sensitivity across distinct populations in our sample, and thus to quantify the benefits of adaptation as observed in historical data. The remaining sections detail how we combine this empirical information with the theoretical framework from Section 4.2 to generate global projections of mortality risk under climate change, accounting for both benefits and costs of adaptation, in addition to how we account for uncertainty in these projections.

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

We begin by estimating a pooled, multi-country mortality-temperature response function. The model exploits year-to-year variation in the distribution of daily weather to identify the response of all-cause mortality to temperature, following, for example, Deschênes and Greenstone (2011). Specifically, we estimate the following equation on the pooled mortality sample from 40 countries,[23]

$$M_{ait} = g_a(\boldsymbol{T}_{it}) + q_{ac}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{4.10}$$

where a indicates age category with $a \in \{< 5, 5\text{-}64, > 64\}$, i denotes the second administrative level (ADM2, e.g., county),[24] c denotes country, and t indicates years. Thus, M_{ait} is the age-specific all-cause mortality rate in ADM2 unit i in year t. α_{ai} is a fixed effect for $age \times ADM2$, and δ_{act} a vector of fixed effects that allow for shocks to mortality that vary at the $age \times country \times year$ level.

Our focus in Equation 4.10 is the effect of temperature on mortality, represented by the response function $g_a(\cdot)$, which varies by age. Before describing the functional form of this response, we note that our climate data are provided at the grid-cell-by-day level. To align gridded daily temperatures with annual administrative mortality records, we first take nonlinear functions of grid-level daily average temperature and sum these values across the year. This is done before the data are spatially averaged in order to accurately represent the distributions at grid cell level. We then collapse annual observations across grid cells within each ADM2 using population weights in order to represent temperature exposure for the average person within an administrative unit (see Appendix 4.B.2.4 for details). This process results in the annual, ADM2-level vector $\boldsymbol{T}_{it}$. We then choose $g_a(\cdot)$ to be a *linear* function

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

[24] This is usually the case. However, as shown in Table 4.1, 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.

of the *nonlinear* elements of $\boldsymbol{T}_{it}$. This construction allows us to estimate a linear regression model while preserving the nonlinear relationship between mortality and temperature that takes place at the grid-cell-by-day level (Hsiang, 2016). The nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ determine, through their linear combination in $g_a(\cdot)$, the functional form of the mortality-temperature response function.

In our main specification, $\boldsymbol{T}_{it}$ contains polynomials of daily average temperatures (up to fourth order), summed across the year. We emphasize results from the polynomial model because it strikes a balance between providing sufficient flexibility to capture important nonlinearities, parsimony, and limiting demands on the data when covariate interactions are introduced (see Section 4.4.2). Results for alternative functional form specifications are very similar to the fourth-order polynomial and are provided in Appendices 4.D.1 and 4.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 $q_{ac}(\cdot)$.

The core appeal of Equation 4.10 is that the mortality-temperature response function is identified from the plausibly random year-to-year variation in temperature within a geographic unit (Deschênes and Greenstone, 2007). Specifically, the $age \times ADM2$ fixed effects α_{ai} ensure that we isolate within-location year-to-year variation in temperature and rainfall exposure, which is as good as randomly assigned. The $age \times country \times year$ fixed effects δ_{act} account for any time-varying trends or shocks to age-specific mortality rates which are unrelated to the climate.

We fit the multi-country pooled model in Equation 4.10 using weighted least squares, weighting by age-specific population so that the coefficients correspond to the average person in the relevant age category and to account for the greater precision associated with mortality estimates from larger populations.[25] 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 Section 4.5, and to alternative climate datasets in Appendix 4.D.1.

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

The average treatment effect identified through Equation 4.10 is likely to mask important differences in the sensitivity of mortality rates to changes in temperature across the diverse populations included in our sample. These differences in sensitivity reflect differential investments in adaptation – i.e., different levels of $\boldsymbol{b}^*$. We cannot observe the level of $\boldsymbol{b}$ directly, but we can observe those factors that influence how poplations select an optimal $\boldsymbol{b}^*$ and condition on those directly to model heterogeneity in the temperature-mortality relationship. We develop a simple two-factor interaction model using average temperature

[25] 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.

(i.e., long-run climate) and average per capita incomes to explain cross-sectional variation in the estimated mortality-temperature relationship. This approach provides provides separate estimates for the effect of climate-driven adaptation and income growth on shape of the temperature-mortality relationship, as they are observed in the historical record.

The two factors defining this interaction model come directly from the theoretical framework in Section 4.2. First, a higher average temperature incentivizes investment in heat-related adaptive behaviors, as the return to any given adaptive mechanism is higher the more frequently the population experiences days with life-threatening temperatures. In Section 4.2, this was represented by $\boldsymbol{b}^*$ being a function of climate $\boldsymbol{C}$. In our empirical specification, we use a parsimonious parameterization of the climate, interacting our nonlinear temperature response function with the location-specific long-run average temperature.[26] Second, higher incomes relax agents' budget constraints and hence facilitate adaptive behavior. In Section 4.2, this was captured by optimal adaptation $\boldsymbol{b}^*$ being an implicit function of income Y. To capture this effect, we interact the temperature polynomial with location-specific per capita income.

In addition to these theoretical arguments, there is a practical reason to restrict ourselves to these two covariates when estimating this interaction model. In order to predict responses around the world and inform projections of damages in the future, it is necessary for all key covariates in the specification to be available globally today, at high spatial resolution, *and* that credible projections of their future evolution are available. Unlike other covariates that may be of interest, average incomes and climate can be extracted from the SSPs and the climate simulations, respectively. These two factors have been the focus of studies modeling heterogeneity across the broader climate-economy literature.[27]

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

$$M_{ait} = g_a(\boldsymbol{T}_{it} \mid TMEAN_s, \log(GDPpc)_s) + q_{ca}(\boldsymbol{R}_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{4.11}$$

where s refers to ADM1-level (e.g., state or province), $TMEAN$ is the sample-period average annual temperature, $GDPpc$ is the sample-period average of annual GDP per capita, and all other variables are defined as in Equation 4.10. We implement a form of $g_a(\cdot)$ that exploits linear interactions between each ADM1-level covariate and all nonlinear elements of the temperature vector $\boldsymbol{T}_{it}$. The model does not include uninteracted terms for $TMEAN$ and $GDPpc$ because they are collinear with α_{ai}. In contrast to the uninteracted models in Equation 4.10, we estimate Equation 4.11 without any regression weights since we are explicitly modeling heterogeneity in treatment effects rather than integrating over it (Solon, Haider, and Wooldridge, 2015). This specification allows for the same flexibility in the functional form of temperature as in Equation 4.10, it is just conditional on income and climate. More details on implementation of this regression are given in Appendix 4.D.4.[28]

[26] In Appendix 4.D.5, we show robustness of this parsimonious characterization of the long-run climate to a more complex specification.

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

[28] To see how we implement Equation 4.11 in practice, note that in Equation 4.10, 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

Equation 4.11 relies on both plausibly random year-to-year fluctuations in temperature within locations and cross-sectional variation in climate and income between administrative units. We rely on cross-sectional variation to identify the interaction effects, because a representative sample of modern populations have not experienced an alternative climate that could be exploited to identify these terms. The consequence is that the case for causally interpreting the coefficients capturing interactions in Equation 4.11 is weaker than for other coefficients in Equation 4.11 and those in Equation 4.10.

We nonetheless view the resulting estimates as informative for at least two reasons. First, the objects of interest are the interactions, not the level of mortality, so while unobserved factors like institutions undoubtedly affect the overall mortality rate (Acemoglu, Johnson, and Robinson, 2001), their potential influence on the mortality sensitivity of temperature is less direct, particularly after adjustment for income and climate. Second, we probed the reliability of the interaction coefficients in several ways and found them to be robust. For example, we found that the estimation off Equation 4.11 in the main sample provides reliable estimates of the mortality temperature sensitivity in India (see Appendix section 4.D.6), providing an out-of-sample test of Equation 4.11. Additionally, the coefficients are qualitatively unchanged when we use alternative characterizations of the climate (see Appendix 4.D.5) and weather (see Appendix 4.F).

4.4.3 Spatial extrapolation: Constructing a globally representative response

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

To achieve the global coverage essential to understanding the costs of climate change, we use the results from the estimation of Equation 4.11 on the observed 38% global sample to estimate the sensitivity of mortality to temperature everywhere, including the unobserved 62% of the world's population. Specifically, the results from this model enable us to use

$p = 1, ..., P$ and aggregated across grid cells. The coefficients $\boldsymbol{\beta}_a$ therefore fully describe the age-specific nonlinear response function. In Equation 4.11, 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. 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 $\boldsymbol{T}_{it}$.

two observable characteristics – average temperature and income – to predict the mortality-temperature response function for each of our 24,378 impact regions. Importantly, it is not necessary to recover the overall mortality rate for these purposes.

To see how this is done, we note that the projected response function for any impact region r requires three ingredients. The first are the estimated coefficients $\hat{g}_a(\cdot)$ from Equation 4.11. The second are estimates of GDP per capita at the impact region level.[29] And third is the average annual temperature (i.e., a measure of the long-run climate) for each impact region, where we use the same temperature data that were assembled for the regressions in Equations 4.10 and 4.11.

We then predict the shape of the response function for each age group a, impact region r, and year t, up to a constant: $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt})$ for $t = 2015$. The various fixed effects in Equation 4.11 are unknown and omitted, since they were nuisance parameters in the original regression. This results in a unique, spatially heterogeneous, and globally comprehensive set of predicted response functions for each location on Earth.

The accuracy of the predicted response functions will depend, in part, on its ability to capture responses in regions where mortality data are unavailable. An imperfect but helpful exercise when considering whether our model is representative is to evaluate the extent of common overlap between the two samples. Figure 4.3A shows this overlap in 2015, where the grey squares reflect the joint distribution of GDP and climate in the full global partition of 24,378 impact regions and orange squares represent the analogous distribution only for the impact regions in the sample used to estimate Equations 4.10 and 4.11. It is evident that temperatures in the global sample are generally well-covered by our data, although we lack coverage for the poorer end of the global income distribution due to the absence of mortality data in poorer countries. We explore this extrapolation to lower incomes with a set of robustness checks in Appendix 4.D; we find the model to perform well in an out-of-sample test and to be robust to alternative functional form assumptions. We do a similar type of prediction when we project temperature-mortality relationships into the future, discussed in the next section, and thus make a similar comparison of samples. We find that, at the end of the century, the overlap is generally better, although unsurprisingly the support of our historical data does not extend to the highest projected temperatures and incomes. Thus, in our projections of the future, in some location and year combinations, we must make out-of-sample predictions about how temperature sensitivity will diminish beyond that observed anywhere in the world today, as temperatures and incomes rise outside of the support in existing global cross-section. We assess the robustness of our results to different assumptions regarding out-of-sample temperatures in Section 4.5.4 and Appendix 4.F.3.

4.4.4 Temporal projection: Accounting for future adaptation benefits

As discussed in Section 4.2, a measure of the full mortality risk of climate change must account for the benefits that populations realize from optimally adapting to a gradually warming climate, as well as from income growth relaxing the budget constraint and enabling

[29] The procedure is described in Section 4.3.2 and Appendix 4.B.3.2

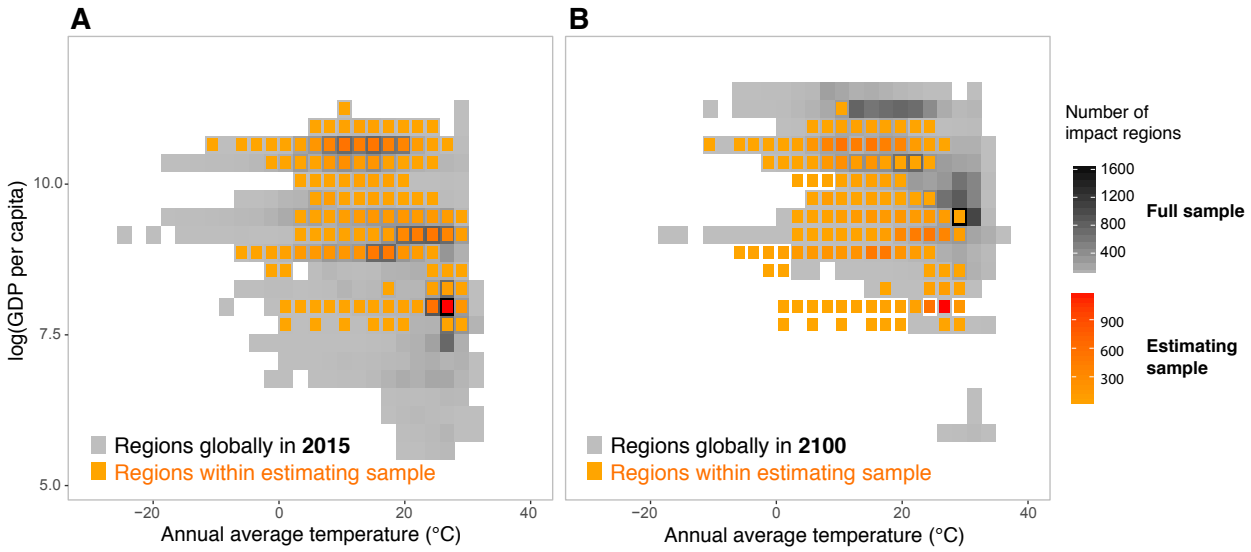

FIGURE 4.3
Joint coverage of income and long-run average temperature for estimating and full samples.

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

compensatory investments. Thus, we allow each impact region's mortality-temperature response function to evolve over time, reflecting how we might plausibly expect climate and incomes to change—as described in a set of internationally standardized and widely used scenarios. We model the evolution of response functions based on projected changes to average climate and GDP per capita, again using the estimation results from fitting Equation 4.11.

We allow the response function in region r and in year t to evolve over time as follows. First, a 13-year moving average of income per capita in region r is calculated using national forecasts from the Shared Socioeconomics Pathways (SSP), combined with a within-country allocation of income based on present-day nighttime lights (see Appendix 4.B.3.2), to generate a new value of $\log(GDPpc)_{rt}$. The length of this time window is chosen based on a goodness-of-fit test across alternative window lengths (see Appendix 4.E.1). Second, a 30-year moving average of temperatures for region r is updated in each year t to generate a new level of $TMEAN_{rt}$. Finally, the response curves $\hat{g}_{art} = \hat{g}_a(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt})$ are calculated for each region for each age group in each year with these updated values of $TMEAN_{rt}$ and $\log(GDPpc)_{rt}$.

The calculation of future mortality-temperature response functions is conceptually straightforward and mirrors the procedure used to extrapolate response functions across locations that do not have historical data. However, as we are generating projections decades into the future, we must impose a set of reasonable constraints on this calculation in order to ensure plausible out-of-sample projections. The following two constraints, guided by economic theory and by the physiological literature, ensure that future response functions are consistent with the fundamental characteristics of mortality-temperature responses that we

observe in the historical record and demonstrate plausible out-of-sample projections.[30] First, we impose the constraint that the response function must be weakly monotonic around an empirically estimated, location-specific, optimal mortality temperature, called the *minimum mortality temperature* (MMT). That is, we assume that temperatures farther from the MMT (either colder or hotter) must be at least as harmful as temperatures closer to the MMT. This assumption is important because Equation 4.11 uses within-sample variation to parameterize how the U-shaped response function flattens; with extrapolation beyond the support of historically observed income and climate, this behavior could go "beyond flat" and the response function would invert (Figure 4.E.1). In fact, this is guaranteed to occur mechanically if enough time elapses, because our main specification only allowed income and climate to interact with the response functions linearly. However, such behavior, in which extreme temperatures are less damaging to mortality rates than more moderate temperatures, is inconsistent with a large body of epidemiological and econometric literature recovering U-shaped response functions for mortality-temperature relationships under a wide range of functional form assumptions and across diverse locations globally (Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011), as well as what we observe in our data. 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.[31,32]

Second, we assume that rising income cannot make individuals worse off, in the sense of increasing the temperature sensitivity of mortality. Because increased income per capita strictly expands the choice set of individuals considering whether to make adaptive investments, it should not increase the effect of temperature on mortality rates. We place no restrictions on the cross-sectional effect of income on the temperature sensitivity when estimating Equation 4.11, but we constrain the marginal effect of income on temperature sensitivity to be weakly negative in future projections. This assumption never binds for temperature sensitivity to hot days (>35°C).[33]

Under these two constraints, we estimate projected impacts separately for each impact region and age group for each year from 2015 to 2100 by applying projected changes in the climate to these spatially and temporally heterogeneous response functions. We compute the nonlinear transformations of daily average temperature that are used in the function $g_a(\boldsymbol{T}_{rt})$ under both the RCP4.5 and RCP8.5 emissions scenarios for all 33 climate projections in the SMME (as described in Section 4.3.2). This distribution of climate models captures uncertainties in the climate system through 2100.

[30] See Appendix 4.E.2 for details on these assumptions and their implementation.

[31] 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.

[32] In imposing this constraint, we hold the MMT fixed over time at its baseline level in 2015 (Figure 4.E.1D). We do so because the use of spatial and temporal fixed effects in Equation 4.11 implies that response function levels are not identified; thus, while we allow the *shape* of response functions to evolve over time as incomes and climate change, we must hold fixed their *level* by centering each response function at its time-invariant MMT. Note that these fixed effects are by definition not affected by a changing weather distribution. Thus, their omission does not influence estimates of climate change impacts.

[33] The assumption that rising income cannot increase the temperature sensitivity of mortality does not bind for hot days because our estimated marginal effects of income are negative for high temperatures (see Table 4.D.3). However, it does bind for the >64 age category under realized temperatures in 30% and 24% of impact regions days in 2050 and 2100, respectively.

4.4.5 Computing adaptation costs using empirical estimates

As shown in Section 4.2, the full cost of the mortality risk due to climate change is the sum of the observable change in mortality and adaptation costs (Equation 4.3). The latter cannot be observed directly; however, as derived in Section 4.2.2, we can recover an expression for adaptation costs that is, in principle, empirically tractable. Specifically, these costs can be computed by taking the difference between the total and partial derivative of expected mortality risk with respect to changes in the climate, and integrating this difference (Equation 4.6). Here, we describe a practical implementation for this calculation.

Our empirical approximation of the adaptation costs incurred as the climate changes gradually from $t = 1$ to $t = 2$ is:

$$\begin{aligned} &\widehat{A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - A}(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1)) \\ &\qquad \approx -\int_1^2 VSL_t \left[\frac{d\hat{\hat{f}}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial \hat{\hat{f}}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_t}{dt} dt \\ &\qquad \approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \underbrace{\left(\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\bigg|_{\boldsymbol{C}_\tau, Y_2} \right)}_{\hat{\boldsymbol{\gamma}}_{\mathbf{1}} \mathrm{E}[\boldsymbol{T}]_\tau} (TMEAN_\tau - TMEAN_{\tau-1}), \end{aligned} \tag{4.12}$$

where the first line of Equation 4.12 is identical to Equation 4.6, except that we use "hat" notation to indicate that $\hat{\hat{f}}(\cdot)$ is an empirical estimate of expected mortality risk. The second (approximate) equality follows from (i) taking the total and partial derivative of our estimating equation (Equation 4.11) with respect to climate — where the total derivative accounts for adaptation while the partial does not, (ii) substituting terms and simplifying the expression, and (iii) implementing a discrete-time approximation for the continuous integral (see Appendix 4.A.3 for a full derivation). The under-braced object, $\hat{\boldsymbol{\gamma}}_{\mathbf{1}} \mathrm{E}[\boldsymbol{T}]_\tau$, is the product of the expectation of temperature and the coefficient associated with the interaction between temperature and climate from estimating equation 4.11: it represents our estimate of marginal adaptation benefits.[34] This derivative is then multiplied by the change in average temperature between each period.[35]

In implementation of Equation 4.12, we treat the VSL as a function of income, which evolves with time, but as invariant to changes in the climate (see Section 4.6). Note that these adaptation cost estimates are calculated for each impact region, age group, and year, using $t = 2015$ as the baseline year, for each of our 33 high-resolution climate model projections.

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

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

4.4.6 Accounting for uncertainty in projected mortality effects of climate change

An important feature of the analysis is to develop estimates of the mortality impacts of climate change that reflect the inherent uncertainty in these future projections.[36] As discussed in Sections 4.4.4 and 4.4.5, we construct estimates of the mortality risk of climate change for each of 33 distinct climate projections in the SMME, capturing uncertainty in the climate system.[37] Additionally, there exists an important second source of uncertainty in our projected impacts that is independent of physical uncertainty, arising from the econometric estimates of response functions; i.e., uncertainty in the estimate of $\hat{g}_a(\cdot)$.

In order to account for both of these sources of uncertainty, we execute a Monte Carlo simulation following the procedure in Hsiang et al. (2017). First, for each age category, we randomly draw a set of parameters, corresponding to the terms composing $\hat{g}_a(\cdot)$, from an empirical multivariate normal distribution characterized by the covariance between all of the parameters from the estimation of Equation 4.11.[38] Second, using these parameters in combination with location- and time-specific values of income and average climate provided by a given SSP scenario and RCP-specific climate projection from each of the 33 climate projections in the SMME, we construct a predicted response function for each of our 24,378 impact regions. Third, with these response functions in hand, we use daily weather realizations for each impact region from the corresponding simulation to predict an annual mortality impact. Finally, this process is repeated until approximately 1,000 projection estimates are complete for each impact region, age group, and RCP-SSP combination.

With these $\sim$1,000 response functions, we calculate the full mortality risk (i.e., inclusive of adaptation benefits and costs) for each impact region for each year between 1981 and 2100. The resulting calculation is computationally intensive (requiring $\sim$94,000 hours of CPU time across all scenarios reported in the main text and Appendix) but incorporates important uncertainty from climate and econometric sources. When reporting projected impacts in any given year, we report summary statistics (e.g., mean, median) of this entire distribution.

4.5 Results I: Full mortality risk of climate change

This section presents results describing temperature's impact on mortality, heterogeneity in that impact, and projections of the full mortality risk of climate change into the future. Projections of global climate change impacts rely on extrapolation of mortality-temperature responses to parts of the world where historical mortality data are unavailable and to future

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

[37] Note that while the SMME fully represents the tails of the climate sensitivity distribution as defined by a probabilistic simple climate model (see Appendix 4.B.2.3), there remain important sources of climate uncertainty that are not captured in our projections, due to the limitations of both the simple climate model and the GCMs. These include some climate feedbacks that may amplify the increase of global mean surface temperature, as well as some factors affecting local climate that are poorly simulated by GCMs.

[38] Note that coefficients for all age groups are estimated jointly in Equation 4.11, such that across-age-group covariances are accounted for in this multivariate distribution.

time periods; applying the approaches described in Sections 4.4.3 and 4.4.4. Using the approach outlined in Sections 4.4.5 and 4.4.6, we then calculate the full global mortality risk of climate change, accounting for the benefits and costs of adaptation and for climate model and econometric uncertainty.

4.5.1 The mortality-temperature relationship: Pooled multi-country results

4.5.1 Pooled Results.

Pooling subnational mortality records across 40 countries and all age groups, we first estimate a version of Equation 4.10 in which $g(\boldsymbol{T}_{it})$ does not depend on age, showing results for the all-age mortality-temperature response function obtained with a fourth-order polynomial in daily average temperature. Table 4.2 displays this result, showing marginal effects at various temperatures. These estimates can be interpreted as the change in the number of deaths per 100,000 per year resulting from one additional day at each temperature, compared to the reference day of 20°C (68°F). Columns (1)-(3) increase the saturation of temporal controls in the model specification, ranging from country-year fixed effects in column (1) to country-year-age fixed effects in column (2) and adding age-specific state-level linear trends in column (3). Our preferred specification is column (2), as column (1) does not account for differential temporal shocks to mortality rates by age group, while in column (3) we cannot reject the null of equal age-specific, ADM1-level trends.

The data reveal the U-shaped response that is common in the prior literature across all specifications. This is noteworthy because the previous literature has relied on samples with much more restrictive geographic and population coverage. Examining column (2), we find that a day at 35°C (95°F) leads to an increase in the all-age mortality rate of approximately 0.4 deaths per 100,000, relative to a day at 20°C. A day at -5°C (23°F) similarly increases the all-age mortality rate by 0.3 deaths per 100,000. This result is robust to alternative functional form assumptions (i.e., different nonlinear functions of $\boldsymbol{T}_{it}$), including a non-parametric binned regression, as well as to the use of alternative, independently-sourced, climate datasets (Figure 4.D.1).

4.5.2 Age Group Heterogeneity.

As prior work has shown that age cohorts respond differently to temperature, and because we expect considerable demographic transitions in the future, we test for heterogeneity across age groups using Equation 4.10. Specifically, we allow for separate mortality-temperature response functions $g_a(\boldsymbol{T}_{it})$ for each of three age categories. Figure 4.4 displays the mortality-temperature responses for each of our three age categories (<5, 5-64, >64) estimated from Equation 4.10 and using the pooled 41-country sample. These curves correspond with our primary specification (equivalent to column (2) in Table 4.D.1).[39] This reveals substantial heterogeneity across age groups within our multi-country sample. In our preferred specification, people over the age of 64 experience approximately 4.7 extra deaths per 100,000 for a

[39]Regression results for all specifications shown in Table 4.2 are also shown for the age-specific model in Table 4.D.1.

TABLE 4.2

TEMPERATURE-MORTALITY RESPONSE FUNCTION ESTIMATED USING POOLED SUBNATIONAL DATA ACROSS 40 COUNTRIES.

	All-age mortality rate (per 100,000)				
	(1)	(2)	(3)	(4)	(5)
35° C	0.410***	0.446***	0.210*	0.675**	0.470***
	(0.128)	(0.168)	(0.118)	(0.274)	(0.163)
30° C	0.305***	0.307***	0.130**	0.338***	0.323***
	(0.068)	(0.080)	(0.065)	(0.106)	(0.075)
25° C	0.147***	0.141***	0.054	0.119***	0.150***
	(0.037)	(0.035)	(0.039)	(0.031)	(0.034)
20° C	–	–	–	–	–
	–	–	–	–	–
0° C	0.121	0.114	0.108	0.144**	0.070
	(0.125)	(0.128)	(0.100)	(0.067)	(0.122)
-5° C	0.307**	0.275*	0.193*	0.244**	0.212
	(0.152)	(0.158)	(0.105)	(0.096)	(0.148)
Adj R-squared	0.983	0.989	0.991	0.999	0.989
N	820697	820237	820237	820237	820237
Age × ADM2 FE	Yes	Yes	Yes	Yes	Yes
Country × Year FE	Yes				
AGE × Country × Year FE		Yes	Yes	Yes	Yes
Age × ADM1 linear trend			Yes		
Precision weighting (FGLS)				Yes	
13-month exposure					Yes

Notes: This table shows coefficient estimates (standard errors) for a temperature-mortality response function estimated using pooled subnational data across 40 countries and 38% of the global population. Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD temperature data with a sample that was winsorized at the top 1% level. Point estimates indicate the effect of a single day at each daily average temperature value shown, relative to a day with an average temperature of 20°C (68°F). 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).

day at 35°C compared to a day at 20°C, a substantially larger effect than that for younger cohorts, which exhibit little response. This age group is also more severely affected by cold days; estimates suggest that people over the age of 64 experience 3.4 deaths per 100,000 for a day at −5°C 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, causing it to differ from the average treatment effect in Table 4.2.

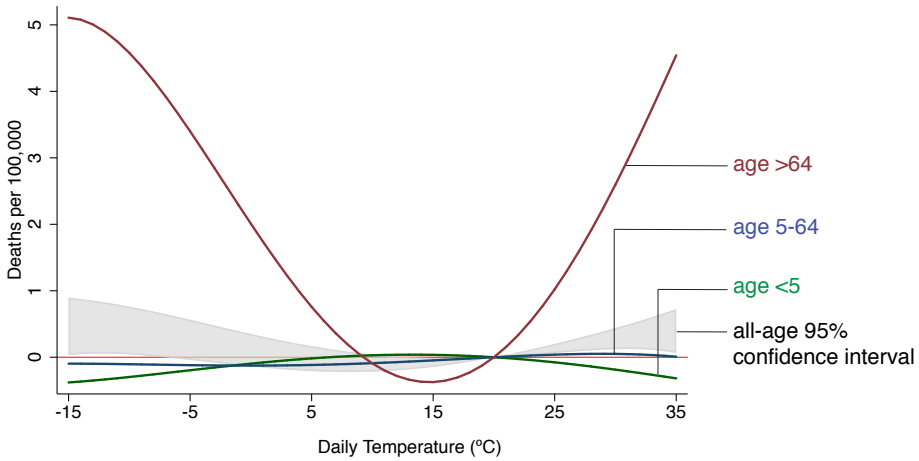

FIGURE 4.4

Mortality-temperature response function with demographic heterogeneity. Mortality-temperature response functions are estimated for populations <5 years of age (green), between 5 and 64 years of age (blue), and >64 years of age (red). Regression estimates shown are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level. All response functions are estimated jointly in a stacked regression model that is fully saturated with age-specific fixed effects (Equation 4.10).

4.5.3 Alternative Specifications.

In both Tables 4.2 and 4.D.1, columns (4) and (5) provide results from alternative specifications. 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 is 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.[40] In column (5), we address the possibility that temperatures can exhibit lagged effects on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). Lagged effects within and across months in the same calendar year are accounted for in the net annual mortality totals used in all specifications. However, it is possible that temperature exposure in December of each year affects mortality in January of the following year. To account for this, in column (5) we define a 13-month exposure window to additionally account for temperatures previous December.[41] Tables 4.2 and 4.D.1 show that the results for both of these alternative specifications are similar in sign and magnitude to those from column (2).

[40] To do this, we estimate the model in Equation 4.10 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.

[41] 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 $t-1$.

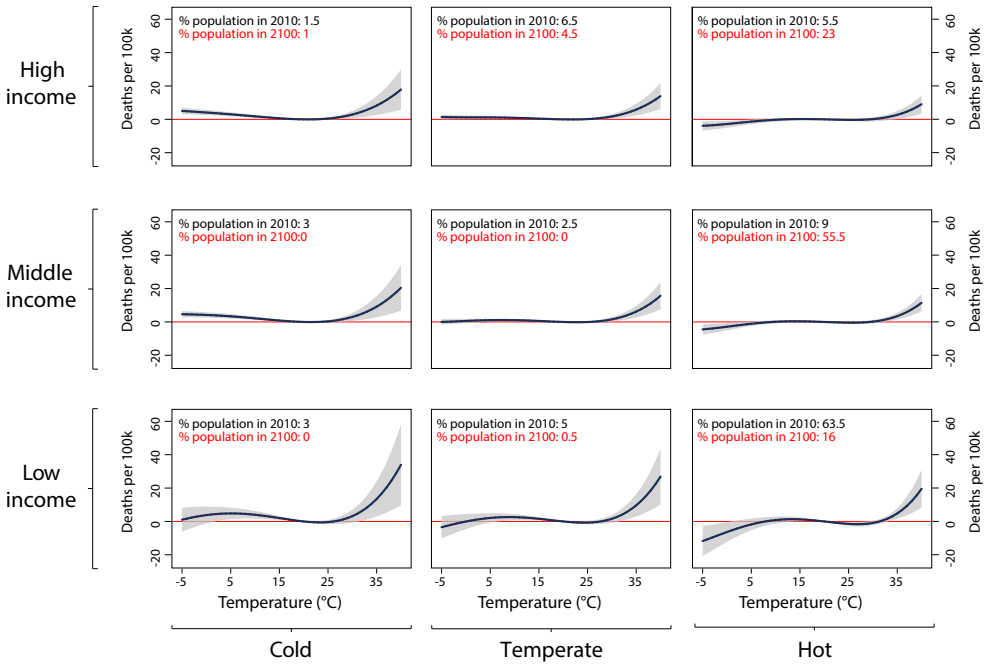

FIGURE 4.5
Heterogeneity in the mortality-temperature relationship (age >64 mortality rate).

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

4.5.2 Subnational heterogeneity in the mortality-temperature response

It is likely that Equation 4.10 obscures heterogeneity in the mortality-temperature response function; this subsection evaluates whether mortality sensitivity to temperature varies with average climate and average income through estimation of Equation 4.11. Tabular results are reported in Table 4.D.3 for each of three age groups. As these terms are difficult to interpret, we visualize this heterogeneity by dividing the sample into terciles of income and climate (i.e., the two interaction terms), creating nine discrete bins that partition the $\log(GDPpc) \times TMEAN$ space. We plot predicted response functions at the mean value of climate and income within each of these nine bins, using the coefficients in Table 4.D.3. This results in a set of predicted response functions that vary across the joint distribution of income and average temperature within our sample data. The resulting response functions are shown in Figure 4.5 for the >64 age category (other age groups are shown in Appendix 4.D.4), where average incomes are increasing across bins vertically and average temperatures are increasing across bins horizontally.

The Figure 4.5 results are broadly consistent with the predictions from the theoretical framework in Section 4.2. Recall that we expect increased frequency of exposure to higher

temperatures to incentivize investment in adaptive behaviors or technologies, as the marginal mortality benefit of adaptation is higher in hotter locations. This would lead to lower temperature sensitivities to heat in places which are warmer. A striking visual finding is that within each income tercile, the effect of hot days (e.g., days >35°C) declines as one moves from left (cold climates) to right (hot climates). Similarly, a loosening of the budget constraint, as proxied by increasing GDP per capita, should enable individuals to invest further in adaptation. Indeed, Figure 4.5 reveals that moving from the bottom (low income) to top (high income) within a climate tercile causes a substantial flattening of the response function, especially at high temperatures.

Two statistics help to summarize the findings from Figure 4.5. First, in the >64 age category across all income values, moving from the coldest to the hottest tercile saves on average 7.9 (p-value=0.06) deaths per 100,000 at 35°C. Second, moving from the poorest to the richest tercile across all climate values in the sample saves approximately 5.0 (p-value=0.1) deaths per 100,000 at 35°C for the > 64 age category.

4.5.3 Spatial extrapolation of temperature sensitivity

Figure 4.6 reports on our extrapolation of mortality-temperature response functions to the entire globe. In panel A, these predicted mortality-temperature responses are plotted for each impact region for 2015 values of income and climate for the oldest age category and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Despite a shared overall shape, panel A reveals substantial heterogeneity across regions in this temperature response. Panel B shows an analogous figure for the youngest age category. 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 day, relative to a day at a location-specific minimum mortality temperature. Grey areas are locations where mortality data are unavailable.

Figure 4.6E–H show analogous plots, but now extrapolated to the entire globe. We can fill in the estimated mortality effect of a 35°C day for regions without mortality data by using location-specific information on income and climate during 2015. The predicted responses at the global scale imply that a 35°C day increases the average mortality rate across the globe for the oldest age category by 10.1 deaths per 100,000 relative to a location-specific minimum mortality temperature.[42] It is important to note that the effect in locations without mortality data is 11.7 deaths per 100,000, versus 7.8 within the sample of countries for which mortality data are available, largely driven by the fact that our sample represents wealthier populations where temperature responses are more muted. Overall, there is substantial heterogeneity across the planet and it is evident that the effects of temperature on human well-being are quite different in places where we are and are not able to obtain subnational mortality data.

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

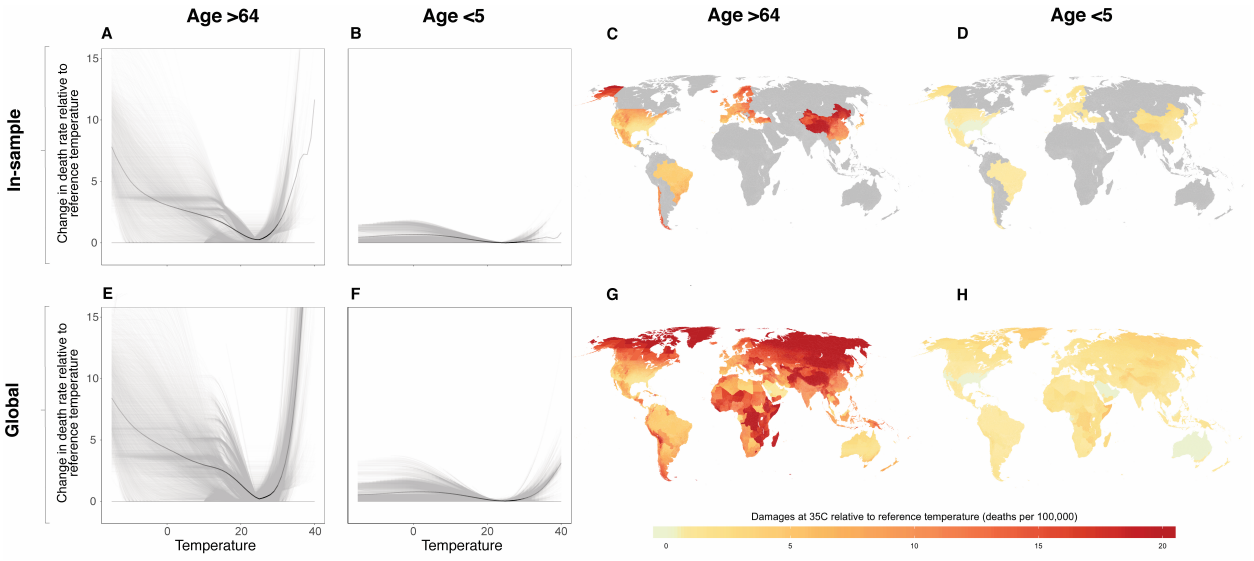

FIGURE 4.6
Using income and climate to predict current response functions globally.

In panels A, B, E and F, grey lines are predicted response functions for impact regions, each representing a population of 276,000 on average. Solid black lines are the unweighted average of the grey lines, where the opacity indicates the density of realized temperatures (Hsiang, 2013). Panels C, D, G and H show each impact region's mortality sensitivity to a day at 35°C, relative to a location-specific minimum mortality temperature. The top row shows all impact regions in the sample of locations with historical mortality data (included in main regression tables), and the bottom row shows extrapolation to all impact regions globally. Column titles indicate corresponding age categories. Predictions shown are averages over the period 2001-2015.

4.5.4 Projection of future climate change impacts and adaptation

The previous subsection demonstrated that the model of heterogeneity outlined in Equation 4.11 allows us to extrapolate mortality-temperature relationships to regions of the world without mortality data today. However, to calculate the full global mortality risks of climate change, it is also necessary to allow these response functions to change through time to capture the benefits of adaptation and the effects of income growth. We use our model of heterogeneity and downscaled projections of income and climate to predict impact region-level response functions for each age group and year, as described in Section 4.4.4. Uncertainty in these estimated response functions is accounted for through Monte Carlo simulation, as described in Section 4.4.6. Throughout this subsection, we show results relying on income and population projections from the socioeconomic scenario SSP3; see Appendix 4.F for results using SSP2 and SSP4. 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 show results relying on SSP3 throughout the main text 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 (see Table 4.B.3).

4.5.1 Projected Adaptation Benefits.

Figure 4.7 provides an initial look into changes in the mortality-temperature relationship over time, which is a key ingredient for projections of future damages and adaptation. In

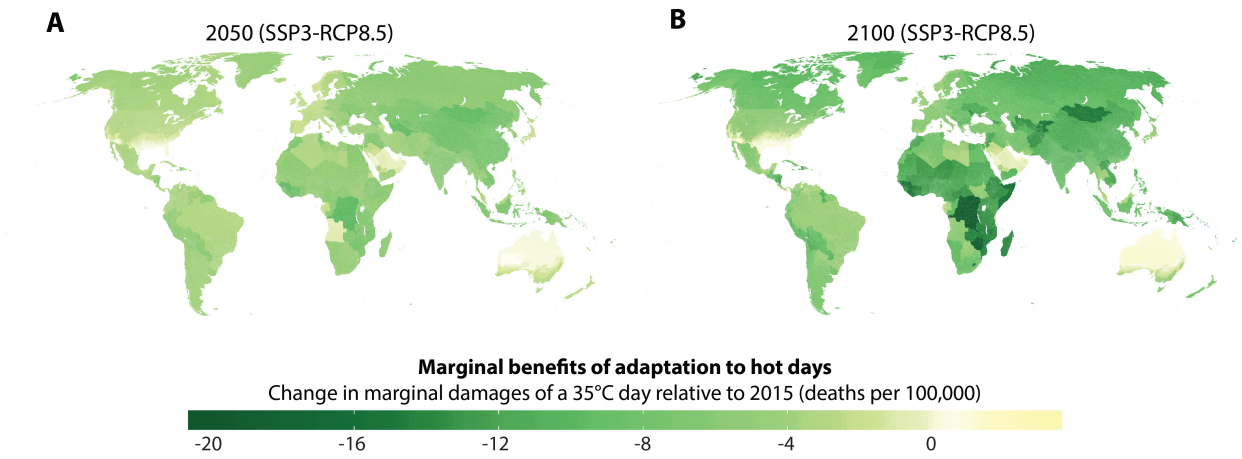

FIGURE 4.7
Spatial and temporal heterogeneity in temperature sensitivity.

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

particular, 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.[43] 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. Increasing incomes account for 77% of the decline in marginal damages for the >64 age category with adaptation to climate explaining the remainder; income gains account for 89% and 82% of the decline for the <5 and 5-64 categories, respectively.[44]

4.5.2 Defining Four Measures of Expected Climate Change Impacts.

We now use our estimates of adaptation benefits, adaptation costs, and changes in climate exposure to develop measures of the expected costs of climate change induced mortality risk. In so doing, we separate the role of income growth from that of adaptation to warming. While the central welfare metric of concern is the full mortality-related costs of climate change, i.e., the sum of the increase in deaths and adaptation costs (Equation 4.3), we also derive three other measures of climate change impacts that elucidate the roles of adaptation and income growth in determining the full mortality-related costs. The empirical estimation of each of these measures is first reported in units of deaths per 100,000, although it is straightforward to monetize these measures using estimates of the value of a statistical life (VSL), and we

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

[44]These values apply to socioeconomic scenario SSP3.

do so at the end of this section.

Note that in all expressions of climate change impacts below, the first term represents the predicted mortality rate under a future warming climate. The second term represents a counterfactual predicted mortality rate that would be realized under current temperatures, but in a population that benefits from rising incomes over the coming century. These counterfactuals thus include the prediction, for example, that air conditioning will become much more prevalent in a country like India as the economy grows, regardless of whether climate change unfolds or not. One exception is expression (i), where the counterfactual is identical to that under current temperatures and incomes, as in this case the climate change projection ignores income growth and adaptation entirely. By subtracting off these counterfactuals, our predicted changes in mortality rates isolate the additional cost of climate change on a population experiencing economic growth.

The first measure is the *mortality effects of climate change with neither adaptation nor income growth*, which provides an estimate of the increases in mortality rates when each impact region's response function in each year t is a function of their 2015 level of income and average climate. In other words, mortality sensitivity to temperature is assumed not to change with future income or temperature. This is a benchmark model often employed in previous work. Specifically, the expected climate induced mortality risk that we estimate for an impact region and age group in a future year t under this measure are (omitting subscripts for impact regions and age groups for clarity):[45]

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

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

The second measure is the *mortality effects of climate change with benefits of income growth*, which allows response functions to change with future incomes. This measure captures the change in mortality rates that would be expected from climate change if populations became richer, but they did not respond optimally to warming by adapting above and beyond how they would otherwise cope with their historical climate. This measure is defined as:

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

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

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

[45]Note that in all estimates of climate change impacts in (i)–(iv), population growth is accounted for as an exogenous projection that does not depend on the climate.

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

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

The final measure is the most complete, as it captures both the benefits and costs of adaptation. Recall that adaptation costs cannot be observed directly, but that we construct estimates using the revealed preference methodology detailed in Section 4.2. We call this measure the *full mortality risk of climate change*, and it captures the opportunity costs of investing in the adaptation benefits described by (iii):

(iv) *Full mortality risk of climate change* (including adaptation costs, recall Equation 4.3):

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

In all of these measures, year $t = 2015$ is treated as the baseline year, meaning that climate change impacts are defined to be zero in this year. These four measures are all reported below in units of human lives. Using human lives serves as a natural numeraire in this revealed preference framework since we estimate adaptation costs based on lives that could be saved via adaptation, but are not. We refer to these as "death equivalents", i.e., the number of avoided deaths equal in value to the adaptation costs incurred. Note that the use of these units is why adaptation costs in expression (iv) are multiplied by $\frac{1}{VSL}$, as the definition of adaptation costs $A(\cdot)$ in Equation 4.6 is given in dollars.

4.5.3 The Full Mortality Risk of Climate Change for 24,378 Global Regions.

Figure 4.8 shows the spatial distribution of the full mortality risk of climate change (expression (iv)) in 2100 under the emissions scenario RCP8.5, expressed in death-equivalents per 100,000. All other measures of climate change impacts (expressions (i)–(iii)) are mapped in Appendix Figure 4.F.1. To construct these estimates, we generate impact-region specific predictions of mortality damages from climate change for all years between 2015 and 2100 (equal to expression (iii)), separately for each age group. Following the approach outlined in Section 4.4.3, we simultaneously compute associated measures of adaptation costs for each location and age at each point in time and add them to expression (iii). The map displays the spatial distribution of the climate-induced death equivalent (expression (iv)), depicting the mean estimate across our ensemble of Monte Carlo simulations, accounting for both climate and statistical uncertainty and pooling across all age groups.[46] The density plots for select cities show the full distribution of impacts across all Monte Carlo simulations, with the white line equal to the mean estimate displayed on the map.

[46] When calculating mean values across estimates generated for each of the 33 climate models that form our ensemble, we use model-specific weights. These weights are constructed as described in Appendix 4.B.2.3 in order to accurately reflect the full probability distribution of temperature responses to changes in greenhouse gas concentrations.

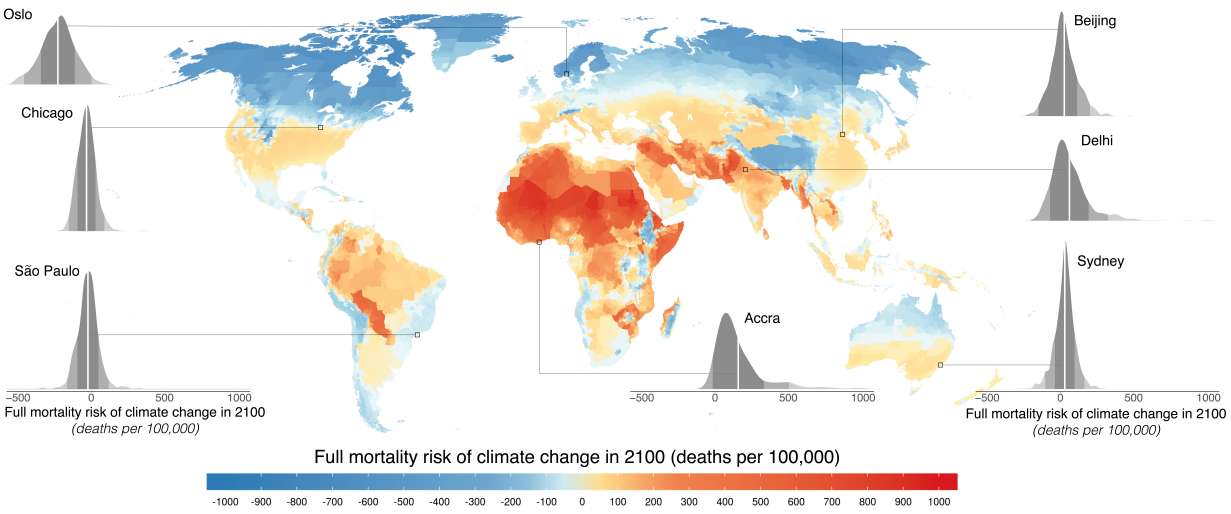

FIGURE 4.8
The mortality risk of future climate change.

The map indicates the full mortality risk of climate change, measured in units of deaths per 100,000 population, in the year 2100. Estimates come from a model accounting for both the costs and the benefits of adaptation, and the map shows the climate model weighted mean estimate across Monte Carlo simulations conducted on 33 climate models; density plots for select regions indicate the full distribution of estimated impacts across all Monte Carlo simulations. In each density plot, solid white lines indicate the mean estimate shown on the map, while shading indicates one, two, and three standard deviations from the mean. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario. See Figure 4.F.4 for a comparison of impacts to RCP4.5 and SSP3

Figure 4.8 makes clear that the costs of climate change-induced mortality risks are distributed unevenly around the world. Despite the gains from adaptation shown in Figure 4.7, there are large increases in mortality risk in the global south. For example, in Accra, Ghana, climate change is predicted to cause damages equivalent to approximately 160 additional deaths per 100,000 annually under RCP8.5 in 2100. In contrast, there are gains in many impact regions in the global north, including in Oslo, Norway, where we predict that the equivalent of approximately 230 lives per 100,000 are saved annually. These changes are equal to an 18% increase in Accra's annual mortality rate and a 28% decline in Oslo's.

4.5.4 Aggregate Global Mortality Consequences of Climate Change.

Figure 4.9 plots predictions of global increases in the mortality rate (deaths per 100,000, including death equivalents for adaptation costs) for all four measures of climate change impacts, under emissions scenario RCP8.5. The measures are calculated for each of the 24,378 impact regions and then aggregated to the global level. In panel A, each line shows a mean estimate for the corresponding climate change impact measure and year. Averages are taken across the full set of Monte Carlo simulation results from all 33 climate models, and all draws from the empirical distribution of estimated regression parameters, as described in Section 4.4.6. In panel B, the 25^{th}-75^{th} and 10^{th}-90^{th} percentile ranges of the Monte Carlo simulation distribution are shown for the full mortality risk of climate change (expression (iv)); the black line represents the same average value in both panels. Boxplots to the right summarize the distribution of mortality impacts for both RCP8.5 and the moderate emissions scenario of RCP4.5, and Figure 4.F.5 replicates the entire figure for RCP4.5.

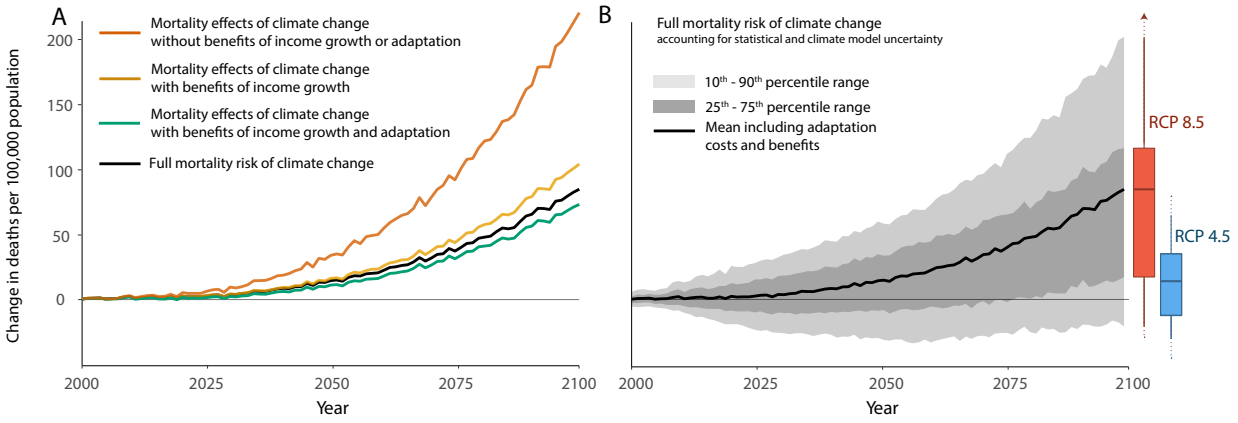

FIGURE 4.9
Time series of projected mortality risk of climate change.

All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. In panel A, each colored line represents a partial mortality effect, while the black line shows the full mortality risk due to climate change, accounting for both adaptation costs and benefits. Orange (expression (i)): mortality effects without adaptation. Yellow (expression (ii)): mortality effects with benefits of income growth. Green (expression (iii)): mortality effects with benefits of income growth and adaptation. Black (expression (iv)): full mortality risk calculated as the sum of mortality effects with adaptation and income growth benefits plus estimates of costs incurred to achieve adaptation, measured in units of death equivalents. Panel B shows the 10^{th}-90^{th} percentile range of the Monte Carlo simulations for the full mortality risk of climate change (black line in panel A), as well as the mean and interquartile range. The boxplots show the distribution of full mortality risk impacts in 2100 under both RCPs. All line estimates shown refer to the RCP8.5 emissions scenario and all line and boxplot estimates refer to the SSP3 socioeconomic scenario. Figure 4.F.5 shows the equivalent for SSP3 and RCP4.5.

Figure 4.9A illustrates that the mortality cost of climate change would be 221 deaths per 100,000 by 2100, on average across simulation runs (orange line), if the beneficial impacts of adaptation and income are shut down. This is an enormously large estimate; if it were correct, the mortality costs of climate change would be roughly equivalent in magnitude to all global deaths from cardiovascular disease today (WHO, 2018). However, we estimate that future income growth and adaptation to climate substantially reduce these impacts. Higher incomes lower the mortality effect of climate change to an average of 104 deaths per 100,000 in 2100 (yellow line); climate adaptation reduces this further to 73 deaths per 100,000 (green line). Although substantially lower than the *no adaptation* projection, these smaller counts of direct mortality remain economically meaningful—for comparison, the 2017 mortality rate from automobile accidents in the United States was 11.4 per 100,000.

Figure 4.9A also demonstrates that climate adaptation is projected to be costly. Our estimates of climate adaptation costs are valued at 12 death-equivalents per 100,000 in 2100. The net result is that the full mortality risk due to climate change (i.e., expression (iv)) is projected to equal to 85 deaths per 100,000 by the end of the century under RCP8.5. Had we accounted for the benefits of adaptation but failed to account for their costs, we would have underestimated the total aggregate impact of these changes, particularly in regions of the world where adaptation costs compose a substantial share of total damages.[47] Nonetheless, our estimate for the global average benefits of adaptation (31 deaths per 100,000) outweighs

[47] We previously noted considerable heterogeneity across age-groups in our results. We will take this into account in our approach to valuing mortality damages monetarily in subsequent sections, and we display the underlying age group heterogeneity of these projections in Appendix 4.F.

the costs of these adjustments (12 death-equivalents per 100,000), demonstrating that the adaptation surplus discussed in Section 4.2 and detailed in Appendix 4.A.2 is substantial.

The values in Figure 4.9A are mean values aggregated across results from 33 high-resolution climate models and all Monte Carlo simulation runs. However, the full distribution of our estimated damages across climate models (panel B of Figure 4.9) is right-skewed with a tail of potential mortality risk far higher than our central estimate. As evidence of this, the *median* value of the full mortality risk of climate change under RCP8.5 at end of century is 56 deaths per 100,000, as compared to a *mean* value of 85 and a 10^{th} to 90^{th} percentile range of [-21, 202].

Figure 4.9B and Appendix Figure 4.F.3 show the expected implications of emissions mitigation. The average estimate of the full mortality risk of climate change of 85 deaths per 100,000 by the end of the century under RCP8.5 falls to 14 under the emissions stabilization scenario of RCP4.5 (where emissions decline after 2050). For RCP4.5, the median end-of-century estimate is 9, and the 10^{th} to 90^{th} percentile range is [-30, 65]. As a point of comparison to the prior literature, we compare these results to the FUND model, which is unique among the IAMs for calculating separate mortality impacts as a component of its SCC calculation. Although it is difficult to make a direct comparison due to differences in scenarios and lack of adaptation costs in the FUND estimates, the closest analog is to compare our estimates including adaptation benefits but not adaptation costs, a change of 73 deaths per 100,000 by 2100, to FUND's change of -2 deaths per 100,000 in the same year (Tol, 1997).[48]

A limit of our empirical approach is that we must sometimes extrapolate response functions to temperatures outside of those historically observed within our data. To address the concern that out-of-sample behavior is disproportionately influencing our results, we repeat the projections of mortality risk changes with two extra sets of restrictions imposed upon our empirically-estimated response functions. These two restrictions, described in Appendix 4.F.3, involve either forcing the response function to be flat for all temperatures outside the observed range, *or* setting the marginal effect to be linearly increasing in the out-of-sample regions with a slope equal to the slope at the edge of the observed range. Figure 4.F.9 reveals that these two restrictions on out-of-sample behavior have negligible effects on our overall impacts. The value of the mortality impact of climate change including benefits of income growth and adaptation is approximately 1 death per 100,000 smaller by 2100 under RCP 8.5 in the case of the flat out-of-sample restriction.

4.5.5 Unequal Distribution of Mortality Risk from Climate Change.

Whether the full mortality risk caused by climate change is realized through actual deaths (first term in expression (iv)), as opposed to costly compensatory investments (second term in expression (iv)), differs substantially across the globe. While some locations suffer large increases in mortality rates, others avoid excess mortality through expensive adaptation. Figures 4.10A-C demonstrate that present day income is strongly correlated with the composition of future damages. In panel A, the negative correlation indicates that today's poor locations tend to suffer large increases in mortality rates by end of century, while mortality

[48] This value was calculated by running the FUND model and extracting quantities from each sector of the model separately.

rates tend to decline due to climate change in today's rich locations. However, there is large variance across impact regions within each income decile, implying that some poor regions are projected to experience mortality rate declines, and some wealthy regions mortality rate increases. In panel B, the positive correlation indicates that wealthier locations are predicted to pay for future adaptive investments, while such costs are predicted to be much smaller in poor parts of the globe. Panel C shows that the full mortality risk of climate change, the sum of both deaths and adaptation costs measured in death-equivalents, is still borne disproportionately by regions that are poor today. On average, we find that in the poorest decile of today's income distribution, just 9% of the total burden of climate change induced mortality risk is borne as adaptation costs. In contrast, in the richest decile, on average approximately 3 lives are saved per 100,000 in 2100 due to climate change, while adaptation costs are three times larger than they are in today's poorest regions. It is also apparent that poorer regions face higher uncertainty in the magnitude of their projected impacts (Figure 4.F.6). Similar figures in panels D–F demonstrate that the hottest locations today suffer the largest increase in death rates, while the coldest pay the highest adaptation costs. The impacts in the top decile of the current long-run climate distribution are noteworthy and raise questions about the habitability of these locations at the end of the century.

4.5.6 Climate Change Projection Scenarios.

The results in this section illustrate a single benchmark emissions and socioeconomic scenario (RCP8.5, SSP3). In Appendix 4.F we report on the sensitivity of the results to alternative choices about the economic and population scenario, the emissions scenario, and assumptions regarding the rate of adaptation. These exercises underscore that the projected impacts of climate change over the remainder of the 21^{st} century will depend greatly on difficult-to-predict factors such as policy, technology, and demographics. However, we note that under both emissions scenarios RCP8.5 and RPC4.5, under all SSP scenarios, and under an alternative projection in which the rate of adaptation is deterministically slowed, the average estimate of the full mortality risk due to climate change is positive (both RCPs) and steadily increasing (RCP8.5) throughout the 21^{st} century.

4.5.7 Monetized Value of the Full Mortality Risk of Climate Change

To monetize the full mortality risk of climate change, we use the value of a statistical life (VSL) to convert changes in mortality rates into dollars. Our primary approach relies on the U.S. EPA's VSL estimate of $10.95 million (2019 USD).[49] We transform the VSL into a value per life-year lost using a method described in Appendix 4.H.1, which allows us to compute the total value of expected life-years lost due to climate change, accounting for the different mortality- temperature relationships among the three age groups documented above. We allow the VSL to vary with income, as the level of consumption affects the relative marginal utilities of a small increment of consumption and a small reduction in the

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

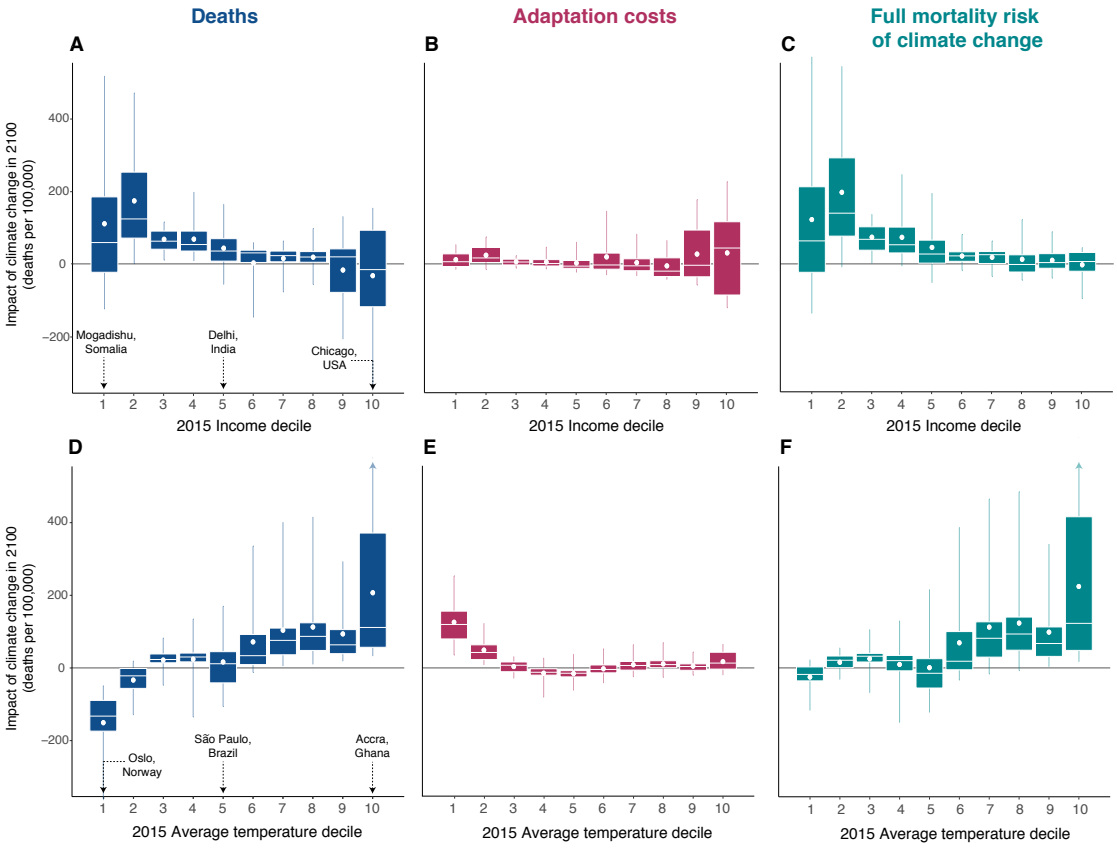

FIGURE 4.10

Climate change impacts and adaptation costs are correlated with present-day income and climate.

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

probability of death. Consistent with existing literature (e.g., Viscusi, 2015), in our primary estimate we use an income elasticity of unity to adjust the U.S. estimates of the VSL to different income levels across the world and over time.[50] When computing the mortality partial SCC in Section 4.6, we provide multiple alternative valuation scenarios in addition to this benchmark case.

Using the SSP3 socioeconomic scenario, we find that under RCP8.5, our benchmark monetized value of the full mortality risk of climate change in 2100 amounts to 3.2% of global GDP on average, with an interquartile range of [-5.4%, 9.1%]. In contrast, under RCP4.5, this value falls to just 0.6% [-3.9%, 4.6%] of global GDP in 2100. While a direct comparison to leading IAMs is made difficult by the use of distinct socioeconomic scenarios and climate models, among many other factors, these mortality-related damages amount to ∼49-135% of the damages reported for *all sectors of the economy* in FUND, PAGE, and DICE, when the damage functions from each model are evaluated at the mean end-of-century warming observed in our multi-model ensemble under RCP8.5. Under RCP4.5, our mortality-related damages amount to 32-61% of the damages from DICE and PAGE, while damages from FUND are negative at RCP4.5 levels of warming.[51] It is apparent that this paper's results indicate that the mortality risks from climate change are much greater than had previously been understood. Some plausible scenarios suggest that, by themselves, they are larger than previous understanding of climate change's *full* impacts in 2100. The uncertainty around these estimates is also meaningful and while we leave explicit pricing of this uncertainty to future work, accounting for it would only increase the estimated welfare loss.

4.6 Results II: Partial mortality social cost of carbon

Section 4.5.3 provided results for the first key goal of this analysis: empirical estimates of the full mortality risk of climate change at high resolution across the globe, for individual years and two emissions scenarios. Here, we take on the second goal of the analysis: to monetize the full mortality-related social cost generated by emitting a marginal ton of CO_2 (defined in Equation 4.9). This represents the component of the *total* SCC that is mediated through excess mortality, but it leaves out adverse impacts in other sectors of the economy, such as reduced labor productivity or changing food prices. Hence, it is a mortality *partial* SCC.

One of two key building blocks of this exercise is an empirically-derived "damage function" (Nordhaus, 1992; Hsiang et al., 2017), which describes economic losses to an economy as a function of the change in the global climate. It is necessary to develop time-varying damage functions in our context, because the mortality sensitivity of temperature and total monetized impacts of climate change evolve over time due to changes in per capita income and the underlying population. Thus, the damages from a marginal change in emissions

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

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

will vary depending on the year in which they are evaluated. Importantly, these damage functions are fully differentiable, so when they are combined with climate model output determining the change in warming arising from any emissions trajectory, it is possible to determine the losses from a marginal increase in emissions. The second key building block is a climate model that is capable of simulating climate trajectories far into the future, and can capture the global climate response of a marginal emission today. Combined, these two elements characterize a trajectory of global impacts, which can then be valued and discounted to inform decisions today.

In this section, we transform monetized projections of deaths due to climate change from Section 4.5.3 into damage functions for excess mortality risk. We then use these time-varying empirical damage functions to compute marginal costs from a marginal CO_2 emission, which is the mortality partial SCC.

4.6.1 Constructing damage functions for excess mortality risk

For this paper, damage functions describe the total monetized losses due to changes in mortality risk, inclusive of adaptation benefits and costs, as a function of the change in *global mean surface temperature* (ΔGMST), our empirical estimate of the change in the global climate.[52] Due to differences in the character of climate projections pre- and post-2100, and lack of available socioeconomic projections after 2100, this subsection details some important differences in the approach for calculation of damage functions before and after 2100. Additionally, it explains the approach to account for damage function uncertainty.

4.6.1 Computing Damage Functions through 2100.

We estimate a set of time-varying quadratic damage functions that relate the total global value of mortality-related climate change damages (D) to the magnitude of global warming (ΔGMST):

$$D(\Delta GMST, t)_{irmt} = \alpha^t + \psi_1^t \Delta GMST_{rmt} + \psi_2^t \Delta GMST_{rmt}^2 + \varepsilon_{irmt} \tag{4.13}$$

To construct the data necessary to estimate Equation 4.13, damages (D_{irmt}) are computed in each year (t) using many climate models (m), two emissions scenarios (r), and a resampling of the econometric parameters recovered from estimation of the mortality-temperature relationship in Equation 4.6 (i). These multiple simulations lead to an empirically-derived distribution of potential economic outcomes that are conditional on the ΔGMST value determined by each climate model in each year for each emissions scenario.[53] To estimate Equation 4.13 for year t, we fit a quadratic function through these simulated outcomes, using all 9,750 Monte Carlo simulation runs within a 5-year window of t, thereby allowing

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

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

$D(\Delta GMST, t)_{irmt}$ to evolve flexibly over the century. We note that pre-2100 damage functions are indistinguishable if we use a third-, fourth- or fifth-order polynomial, and we show robustness of our mortality partial SCC estimates to functional form choice in Appendix 4.H.4.

Figure 4.11A illustrates the procedure for $t = 2097$, with D_{irmt} estimates from all Monte Carlo simulations shown as points[54] located along the horizontal axis based on their corresponding ΔGMST_{irmt}. The median end-of-century warming relative to 2001-2010 under RCP8.5 (red points) across our climate models is +3.7°C, while under RCP4.5 (blue points) it is +1.6°C. The black line is the quadratic damage function estimated for the year 2097, the latest year for which a full 5-year window of damage estimates can be constructed. The estimated damage function in 2097 recovers total (undiscounted) damages with an age-varying VSL at 3.7°C and 1.6°C of $7.8 and $1.2 trillion USD, respectively. Analogous curves are constructed for all years, 2015 to 2100.

4.6.2 Computing Post-2100 Damage Functions.

Even with standard discount rates, a meaningful fraction of the present discounted value of damages from the release of CO_2 today will occur after 2100 (Kopp and Mignone, 2012), so it is important to develop post-2100 damage functions. The pre-2100 approach cannot be used for these later years because only 6 of the 21 GCMs that we use to build our SMME ensemble (see Section 4.3.2) simulate the climate after 2100 for both RCP scenarios. Similarly, the SSPs needed to project the benefits of income growth and changes in demographic compositions also end in 2100.

To estimate post 2100-damages, we develop a method to extrapolate changes in the damage function beyond 2100 using the observed evolution of damages near the end of the 21^{st} century. The motivating principle of 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. This reduced-form approach allows our empirical results to constrain and guide a projection to years beyond 2100. To execute this extrapolation, we pool values D_{irmt} from 2085-2100 and estimate a quadratic model similar to Equation 4.13, but interacting each term linearly with year t.[55] This allows estimation of a damage surface as a parametric function of year, which can then be used to predict extrapolated damage functions for all years after 2100, smoothly transitioning from our climate model-based damage functions prior to 2100. Appendix 4.G provides a detailed explanation of the approach.

The approach produces post-2100 damage functions that become more convex over time, suggesting larger damages for a given level of warming. This finding comes directly from the estimation of Equation 4.13 that found that in the latter half of the 21^{st} century the full mortality damages are larger when they occur later, holding constant the degree of

[54] This scatterplot includes realizations under all RCP4.5 and RCP8.5 scenarios for all projections in our 33-member ensemble under our benchmark method of valuation – the age-invariant EPA VSL with an income elasticity of one applied to all impact regions – in the end-of-century years 2095-2100. See Appendix 4.H for results across different valuation assumptions. Due to the dependence of damages D on GDP per capita and on demographics, we estimate separate damage functions for every SSP scenario. Results across different scenarios are shown in Appendix 4.H.

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

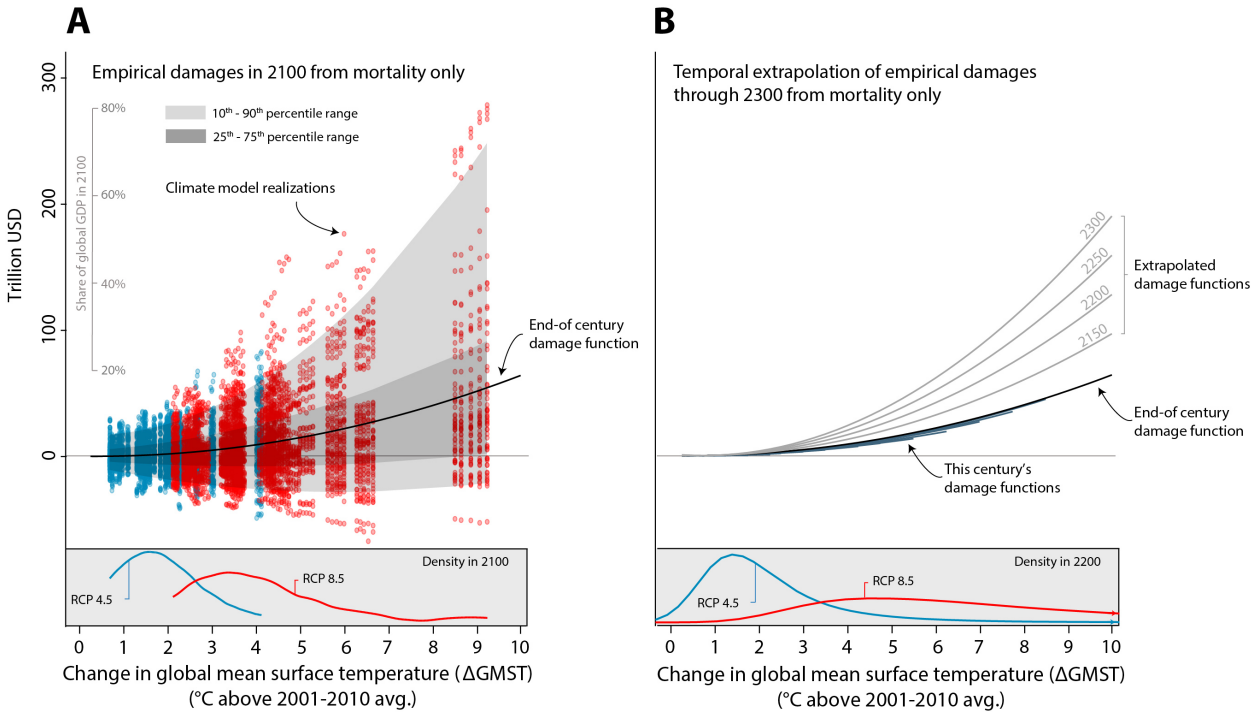

FIGURE 4.11
Empirically-derived mortality-only damage functions.

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

warming. This finding that mortality costs rise over time is the net result of countervailing forces. On the one hand, damages are larger in later years because there are larger and older populations[56] with higher VSLs due to rising incomes. On the other hand, damages are smaller in later years because populations are better adapted due to higher incomes and a slower rate of warming projected in later years. Our results suggest the former dominates by end of century, causing damages to be trending upward when high-resolution simulations end in 2100.

Panel B of Figure 4.11 illustrates damages functions every 10 years prior to 2100, as well as extrapolated damage functions for the years 2150, 2200, 2250, and 2300. In dollar terms, these extrapolated damages continue to rise post-2100 and become steeper, as they did pre-2100. In Appendix 4.H, we explore the importance of this extrapolation by using an alternative approach to estimating post-2100 damages, instead calculating partial SCC estimates using a damage function frozen at its 2100 shape for all years 2101-2300. With this alternative approach, our central estimate of the mortality partial SCC falls 21%, indicating that extrapolation of the damage function has a modest impact on our partial SCC estimates, due in part to the important role of discounting (Table 4.H.6).

4.6.3 Accounting for Uncertainty in Damage Function Estimation.

As discussed, there is substantial uncertainty in projected mortality effects of climate change due to statistical uncertainty in the estimation of mortality-temperature response functions. The approach described above details the estimation of a damage function using the conditional expectation function through the full distribution of simulation results. In addition to reporting the predicted damages resulting from this damage function describing (conditional) expected values, we also estimate a set of quantile regressions to capture the full distribution of simulated mortality impacts.[57] Just as above for the mean damage function, extrapolation past the year 2100 is accomplished using a linear time interaction, estimated separately for each quantile. Central estimates of the mortality partial SCC reported below use the mean regression, while ranges incorporating damage uncertainty use the full set of time-varying quantile regressions.

4.6.2 Computing marginal damages from a marginal carbon dioxide emissions pulse

We empirically approximate the mortality partial SCC for emissions that occur in the year 2020 (Equation 4.9) as:

$$\textit{Mortality partial } SCC_{2020} \approx \sum_{t=2020}^{2300} e^{-\delta(t-2020)} \frac{\partial \hat{D}(\Delta GMST, t)}{\partial \Delta GMST_t} \frac{d\Delta GMST_t}{dCO2_{2020}}, \tag{4.14}$$

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

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

where ΔGMST approximates the multi-dimensional climate vector $\boldsymbol{C}$, we use changes in CO_2 to represent changes in global emissions E,[58] we assume that discounted damages from an emissions pulse in year 2020 become negligible after 2300, and we approximate the integral in Equation 4.9 with a discrete sum using time steps of one year. The values $\frac{\partial \hat{D}(\Delta GMST,t)}{\partial \Delta GMST_t}$ are the marginal damages at each moment in time that occur as a result of this small change in all future global temperatures; they are computed using the damage functions described in the last subsection.

The term $\frac{d\Delta GMST_t}{dCO2_{2020}}$ is the increase in ΔGMST that occurs at each moment in time along a baseline climate trajectory as a result of a marginal unit of emissions in 2020, which we approximate with an "infinitesimally small" pulse of CO_2 emissions. Its estimation requires a climate model capable of estimating the global temperature response in each year to a single pulse of CO_2 emissions. Because we are interested in computing this value for a large number of scenarios, including ones that reflect scientific uncertainty about the physical magnitude and timing of warming, referred to as the *global climate sensitivity*, we use a "simple climate model" to estimate $\frac{d\Delta GMST_t}{dCO2_{2020}}$.

4.6.1 Applying a Simple Climate Model to the Damage Function.

To calculate the change in ΔGMST_t due to a marginal pulse of CO_2 in 2020, we use the Finite Amplitude Impulse Response (FAIR) simple climate model that has been developed especially for this type of calculation (Millar et al., 2017).[59] We use FAIR to calculate ΔGMST_t trajectories for emissions scenarios RCP4.5 and RCP8.5, both with and without an exogenous impulse of 1 gigaton C (equivalent to 3.66Gt CO_2) in the year 2020, an approximation of an infinitesimal emission for which the model numerics are stable. In FAIR, this emissions impulse perturbs the trajectory of atmospheric CO_2 concentrations and ΔGMST for 2020-2300, with dynamics that are influenced by the baseline RCP scenario. In each scenario, the trajectory of ΔGMST_t in the "RCP + pulse" simulation is differenced from the baseline "RCP only" simulation to compute $\frac{d\Delta GMST_t}{dCO2_t}$, and the resulting damages are converted into USD per one ton CO_2. There is naturally uncertainty in these trajectories, and our approach accounts for uncertainty associated with four key parameters of the FAIR model (i.e., the transient climate response, equilibrium climate sensitivity, the short thermal adjustment time, and the time scale of rapid carbon uptake by the ocean mixed layer). This approach, detailed in Section 4.G, ensures that the distribution of transient warming responses we use to generate partial SCC values matches the corresponding distributions from the IPCC Assessment Report 5 (AR5).

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

[59] FAIR is a zero-dimensional structural representation of the global climate designed to capture the temporal dynamics and equilibrium response of ΔGMST to greenhouse gas forcing. Appendix 4.G shows that our simulation runs with FAIR create ΔGMST distributions that match those from the climate projections in the high-resolution models in the SMME.

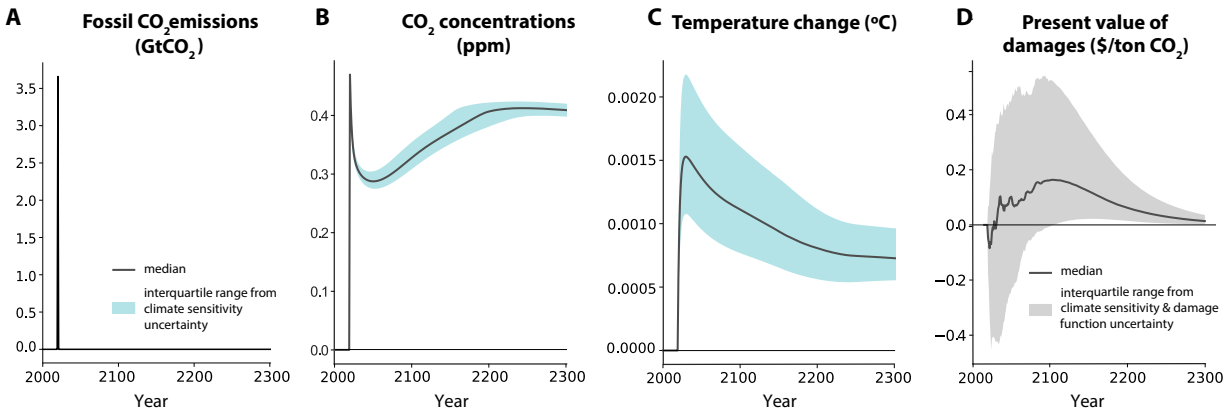

FIGURE 4.12

Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020.

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

4.6.2 Summarizing the Impacts of a Marginal Increase in CO_2 Emissions.

Figure 4.12 graphically depicts the difference between the "RCP + pulse" and baseline RCP trajectories for four key outcomes.[60] The pulse in emissions is shown in panel A. Its influence on CO_2 concentrations is reported in panel B; the immediate decline followed by a century-long increase is largely due to dynamics involving the ocean's storage and release of emissions. Panel C displays the resulting change in temperature, which makes clear that a pulse today will influence temperatures for over three centuries. The shaded green area in panels A-C represents the inter-quartile range of each year's outcome, and reflects uncertainty in the climate system (see Appendix 4.G for details).

Panel D plots the discounted (2% discount rate) stream of damages due to this marginal pulse of emissions. The temporal pattern of the present value of mortality damages reflects the nonlinearity of the damage function (e.g., Figure 4.11), which itself depends on nonlinearities in location-specific mortality-temperature relationships (e.g., Figure 4.5), as well as the discount rate and the dynamic temperature response to emissions (panel C). The peak damages from a ton of CO_2 emissions are \$0.16 in year 2104; by year 2277, annual damages are always less than \$0.02. It is noteworthy that 67.9% of the present value of damages occur after the year 2100. The shaded grey area represents the inter-quartile range of each year's outcome, and reflects uncertainty in the climate system as well as uncertainty in the damage function. RCP 4.5 results are shown in Figure 4.G.3 and additional details are in Appendix 4.G.

[60] Using the trajectories in Figure 4.12 is consistent with the "SCC experiment" that is used in IAMs to calculate an SCC (National Academies of Sciences, Engineering, and Medicine, 2017). We discuss uncertainties in FAIR configuration parameters below and in Appendix 4.G. The median values of parameter-specific distributions used for the central mortality partial SCC estimate include a transient climate response (TCR) of 1.6 and an equilibrium climate sensitivity (ECS) of 2.7.

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

Table 4.3 reports mortality partial SCC estimates. The columns apply four different annual discount rates – three used in prior estimates of the SCC (2.5%, 3%, and 5%) (Interagency Working Group on Social Cost of Carbon, 2010; National Academies of Sciences, Engineering, and Medicine, 2017), and one lower rate that aligns more closely with recent global capital markets (2%) (Board of Governors of the US Federal Reserve System, 2020). Panel A uses the U.S. EPA's VSL of $10.95 million (2019 USD), transformed into value per life-year lost (see Appendix 4.H.1 for details). This accounts for the different mortality-temperature relationships among the three age groups documented above.[61] Panel B uses age-invariant values of the VSL. Both valuations adjust for cross sectional variation in incomes among contemporaries and global income growth. Appendix 4.H presents results under a wide range of additional valuation scenarios, including an alternative and lower Ashenfelter and Greenstone (2004) VSL of $2.39 million (2019 USD),[62] and an approach where the VSL is adjusted only based on global average income such that the lives of contemporaries are valued equally, regardless of their relative incomes.

The estimates in Table 4.3 utilize the median values of FAIR's four key parameter distributions and the mean global damage function. Interquartile ranges (IQRs) are reported, reflecting uncertainty in climate sensitivity (uncertainty in the simple climate model FAIR) and in the damage function (which includes uncertainty arising from econometric estimation and from climate model uncertainty regarding the spatial distribution of warming). All values represent the global sum of each impact region's MWTP today (2019 USD) to avoid the release of an additional metric ton of CO_2 in 2020, including both the costs and benefits of adaptation.

In column 1 and panel A of Table 4.3, we report partial SCC values under a $\delta = 2\%$ discount rate and using an age-varying VSL. Under this valuation approach, the mortality partial SCC is $17.1 [-$24.7, $53.6] for the low to moderate emissions scenario and $36.6 [-$7.8, $73.0] for the high emissions scenario. We highlight a 2% discount rate because it conservatively reflects changes in global capital markets over the last several decades: while the Interagency Working Group on Social Cost of Greenhouse Gases (2016) recommends a discount rate of 3%, the average 10-year Treasury Inflation-Indexed Security value over the available record of this index (2003-present) is just 1.01% (Board of Governors of the US Federal Reserve System, 2020). We show results for a discount rate of 1.5% in Appendix Table 4.H.5. We emphasize the age-varying VSL approach because standard economic reasoning implies that valuation of life lost should vary by age (Jones and Klenow, 2016; Murphy and Topel, 2006).

When following the Interagency Working Group on Social Cost of Greenhouse Gases (2016) preference for a discount rate of $\delta = 3\%$ and the use of an age-invariant VSL, the

[61] In the main text, a simple life-years calculation that assigns each life-year lost the same economic value is used. In Appendix 4.H, we also show calculations that adjust the value of remaining life by age at death using the estimates of age-specific value of remaining life from Murphy and Topel (2006), which produces results that differ only slightly from those under the primary approach.

[62] See Appendix Table 4.H.1 for a comparison of these VSL values with values from the OECD, which are higher than Ashenfelter and Greenstone (2004), but lower than the U.S. EPA's VSL.

TABLE 4.3

ESTIMATES OF A PARTIAL SOCIAL COST OF CARBON FOR EXCESS MORTALITY RISK INCORPORATING THE COSTS AND BENEFITS OF ADAPTATION

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

Notes: In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value (alternative values using the VSL estimate from (Ashenfelter and Greenstone, 2004) are shown in Appendix 4.H). All regions thus have heterogeneous valuation, based on local income. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3 (alternative values using other SSP scenarios are shown in Appendix sec:appsccrobustness). 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 4.H.1 for details and Appendix 4.H.2 for alternative calculations that allow the value of a life-year to vary with age, based on Murphy and Topel (2006)). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Point estimates rely on the median values of the four key input parameters into the climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] including both climate sensitivity uncertainty and damage function uncertainty (see Appendix 4.G for details).

central estimate of the mortality partial SCC is \$6.7 per metric ton of CO_2 for the low to moderate emissions scenario (RCP4.5), with an IQR of [-\$15.7, \$32.1], and \$22.1 [-\$5.6, \$53.4] per metric ton for the high emissions scenario (RCP8.5).

Some other features of these results are worth underscoring. First, mortality partial SCC estimates for RCP4.5 are systematically lower than RCP8.5 primarily because the damage function is convex, so marginal damages increase in the high emissions scenario. Second, the combination of this convexity, which itself is accentuated at higher quantiles of the damage function, and the skewness of the climate sensitivity distribution causes the distribution of partial SCCs to also be right skewed with a long right tail. For example, the 95^{th} and 99^{th} percentiles of the partial SCC for $\delta = 2\%$ and an age-varying VSL for RCP8.5 are \$290.3 and \$704.1, respectively; with $\delta = 3\%$ and an age-invariant VSL these values are \$175.3 and \$391.9. Third, in Appendix 4.G we show that uncertainty in the partial SCC is consistently dominated by uncertainty in the damage function, as opposed

to uncertainty in climate sensitivity. Finally, while Table 4.3 reflects what we believe to be mainstream valuation and socioeconomic scenarios, Appendix Tables 4.H.2-4.H.7 report estimates based on multiple alternative approaches. Naturally, the resulting SCC estimates vary under different valuation assumptions and baseline socioeconomic trajectories, and we point readers to these specifications for a more comprehensive set of results.

4.7 Limitations

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

4.7.1 Migration Responses.

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

4.7.2 Humidity.

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

4.7.3 Technological Change.

Third, the paper's projections incorporate advancements in technology that enhance adaptive ability, even though we have not explicitly modeled technological change. In particular,

we allow mortality-temperature response functions to evolve in accordance with rising incomes and temperatures and do not restrict them to stay within the bounds of the current observed distribution of temperature responses. However, while our estimates reflect technical advancement as it historically relates to incomes and climate, they do not reflect the likely passing of climate-biased technical change that lowers the *relative* costs of goods which reduce the health risks of high temperatures. Therefore, the paper's results will overstate the mortality risk of climate change if directed technological innovations lower the relative costs of adapting to high temperatures.

4.7.4 Uncertainty.

Throughout this paper, we have emphasized central estimates of the mortality risk damages of climate change and the partial SCC, but also noted the meaningful uncertainty associated with these estimates. However, this does not reflect efforts to price this uncertainty. With risk averse agents, this uncertainty would undoubtedly increase the estimates of the climate change mortality risk damages and the mortality partial SCC (Weitzman, 2011; Traeger, 2014).

4.8 Conclusion

This paper has outlined a new method for empirically estimating the costs of climate change for a single sector of the economy and implemented it in the context of mortality risks associated with temperature change. Specifically, we highlight two primary contributions: a framework for estimating the total impact of climate change on mortality risk both globally and for more than 24,000 regions that comprise the planet, and a framework for estimating a mortality partial SCC. There are several noteworthy methodological innovations and intermediate findings.

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

Second, the costs of climate change-induced mortality risks are distributed unevenly across the 24,378 regions that we use to create local-level projections. For example, by 2100, we project that climate change will cause annual damages equivalent to approximately 160 additional deaths per 100,000 in Accra, Ghana, but will also generate annual benefits equivalent to approximately 230 lives per 100,000 in Oslo, Norway. To put these numbers in perspective, Figure 4.13 reveals that the projected impact of climate change on mortality will be comparable globally to leading causes of death today, such as cancer and infectious disease. It also demonstrates the benefits from mitigation, as the end of century estimate of mortality risk of climate change falls from 85 deaths per 100,000 under RCP 8.5 to just 14 per 100,000 under RCP 4.5, though much of the inequality in impacts evidence here remains under RCP4.5 (see Figure 4.F.7). Today's poor bear a disproportionately high share of

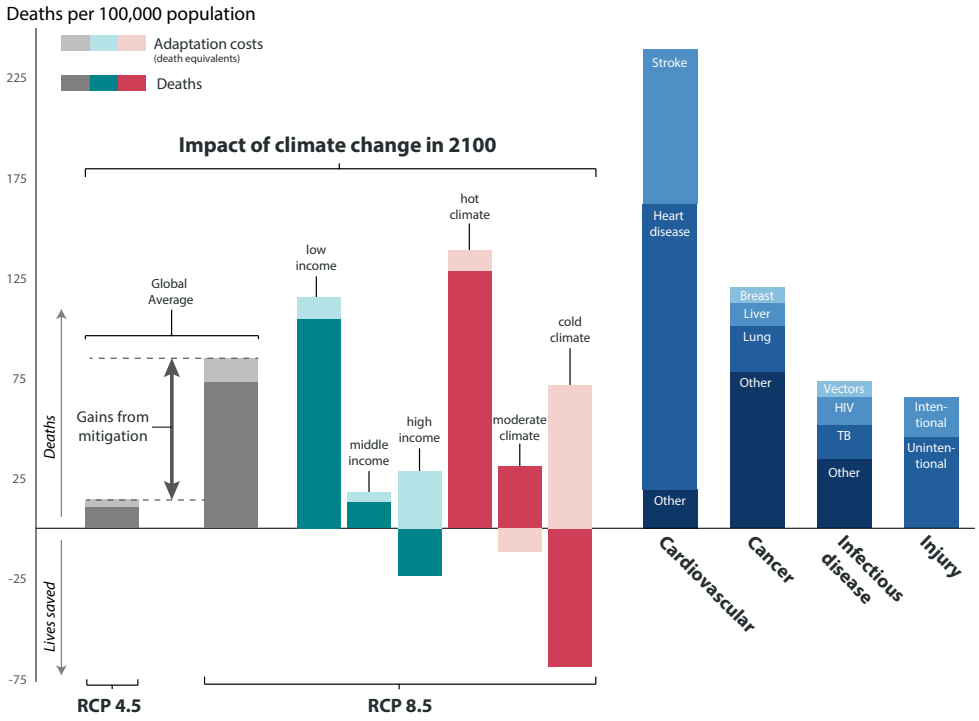

FIGURE 4.13

The impact of climate change in 2100 is comparable to contemporary leading causes of death.

Impacts of climate change (grey, teal, and coral) are calculated for the year 2100 for SSP3 and include changes in death rates (solid colors) and changes in adaptation costs, measured in death equivalents (light shading). Global averages for RCP 8.5 and RCP 4.5 are shown in grey, demonstrating the gains from mitigation. Income and average climate groups under RCP8.5 are separated by tercile of the 2015 global distribution across all 24,378 impact regions. Blue bars on the right indicate average mortality rates globally in 2018, with values from WHO (2018). Figure 4.F.7 provides an RCP4.5 version of this figure.

the global mortality risks of climate change, as current incomes (as well as current average temperatures) are strongly correlated with future climate change impacts.

Third, the heterogeneous impacts appear to reflect investments that individuals and societies make based on the costs and benefits of responding to differences in climate. Estimated mortality impacts that do not account for adaptation overstate the mortality impacts of climate change in 2100 by more than a factor of 2.6. We outline and implement a revealed preference method to recover the costs of these adaptation investments, even though they cannot be directly observed. Notably, the degree to which the full mortality risk of climate change is realized through actual deaths, as opposed to costly adaptation, varies widely across space and time. For example, Figure 4.13 shows that today's poor locations tend to bear a larger share of the projected burden in the form of direct mortality impacts, while today's rich face large increases in projected adaptation costs.

Fourth, there is substantial climate and statistical uncertainty around these estimates and we find that the distribution of projected losses is right skewed; for example the mean loss in 2100 is about 51% larger than the median loss. Although we do not account for risk aversion, it is evident that doing so would increase the valuations of these impacts.

The paper's ultimate goal is to develop the first empirically grounded estimates of climate damages and partial SCC that reflects the consequences of climate change on mortality and investments in adaptation. The estimates suggest that the mortality risk damages account for 3.2% of global GDP in 2100 under a high emissions scenario (RCP8.5) and 0.6% under a moderate one (RCP4.5). Our estimate of the mortality risk partial SCC with a 2% discount rate and age-varying VSL, implies that the present value of excess mortality risk due to climate change imposed by a marginal metric ton of CO_2 emissions in 2020 is roughly \$36.6 [-\$7.8, \$73.0] (in 2019 USD) with a high emissions scenario and \$17.1 [-\$24.7, \$53.6] with a moderate one, where brackets indicate interquartile ranges accounting for climate sensitivity and damage function uncertainty. Under a 3% discount rate and age-invariant VSL, assumptions preferred by the Interagency Working Group on Social Cost of Carbon (2010), these values are \$22.1 [-\$5.6, \$53.4] and \$6.7 [-\$15.7, \$32.1], respectively.

These findings suggest that previous research has significantly understated climate change damages due to mortality. For example, the mortality damages we estimate in 2100 account for 49% to 135% of *total* damages across all sectors of the economy according to leading IAMs. Moreover, the mortality partial SCC reported here, under comparable valuation assumptions, is more than 10 times larger than the total health impacts embedded in the FUND IAM (Diaz, 2014).[63] Further, the high emissions and 3% discount rate partial SCC that is most similar to the scenario underlying the Obama Administration SCC produces an excess mortality partial SCC that is 44% of the Administration's *full* SCC.

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

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

in access to data and computing render these modeling choices unnecessary.

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

Chapter Review and Looking Ahead

In this chapter, we presented the first globally comprehensive and empirically grounded estimates of mortality risk due to future temperature increases caused by climate change. Under a high emissions scenario, we estimate the mean increase in mortality risk is valued at roughly 3.2% of global GDP in 2100, with today's cold locations benefiting and damages being especially large in today's poor and/or hot locations. We estimate that the release of an additional ton of CO2 today will cause mean [interquartile range] damages of $36.6 [-$7.8, $73.0] under a high emissions scenario. Globally, these empirically grounded estimates substantially exceed the previous literature's estimates that lacked similar empirical grounding, suggesting that revision of the estimated economic damage from climate change is warranted.

In the next chapter, we will examine another channel through which climate change is expected to impact society – energy consumption – where we will find a sector of the economy for which gains from reduced cold days offset losses from increased hot days.

References

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2001. "The Colonial Origins of Comparative Development: An Empirical Investigation." *American Economic Review* 91 (5):1369–1401. URL https://www.aeaweb.org/articles?id=10.1257/aer.91.5.1369.

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

———. 2018b. "Quantifying economic damages from climate change." *Journal of Economic Perspectives* 32 (4):33–52.

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

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

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

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

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

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

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

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

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

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

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

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2014. "What do we learn from the weather? The new climate–economy literature." *Journal of Economic Literature* 52 (3):740–798.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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* .

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

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

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

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

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

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

Millner, Antony and Geoffrey Heal. 2018. "Time consistency and time invariance in collective intertemporal choice." *Journal of Economic Theory* 176:158–169.

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.

Newell, Richard G and William A Pizer. 2004. "Uncertain discount rates in climate policy analysis." *Energy Policy* 32 (4):519–529.

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

Park, R Jisung, Joshua Goodman, Michael Hurwitz, and Jonathan Smith. 2020. "Heat and learning." *American Economic Journal: Economic Policy* 12 (2):306–39.

Pindyck, Robert S. 2013. "Climate change policy: What do the models tell us?" *Journal of Economic Literature* 51 (3):860–872.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Appendix

4.A Using revealed preference to estimate adaptation costs

4.A.1 Graphical solution to inferring unobserved adaptation cost

In Section 4.2, 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 4.A.1 depicts this optimal adaptation problem faced by individuals and illustrates how we overcome two key empirical challenges to measuring adaptation costs: (1) the universe of adaptation adjustments and their costs are not directly observable and (2) adaptive adjustments are continuous for continuous changes in climate. The problem must be displayed in three dimensions because it involves at least three orthogonal subspaces: climate ($\boldsymbol{C}$), adaptive adjustments to climate ($\boldsymbol{b}$), and an outcome (expressed in dollars of WTP). For illustrative simplicity, here we assume income is held fixed, and we consider a simplified example with univariate climate and univariate adaptation. Further, for this example, higher $\boldsymbol{C} = C$ indicates higher temperatures and higher $\boldsymbol{b} = b$ indicates greater adaptation (i.e., greater protection) from high temperatures, where these terms are unbolded to indicate that they are scalars.

In the lower left panel of Figure 4.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[64] 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 4.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_1 and C_2, where slices of the benefits surface V are drawn overlaid in red and are tangent to $A(b)$ at the blue circles. Corresponding slices of the adaptation cost surface A are overlaid in green on the benefits surface in the lower right panel. The blue line traces out the

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

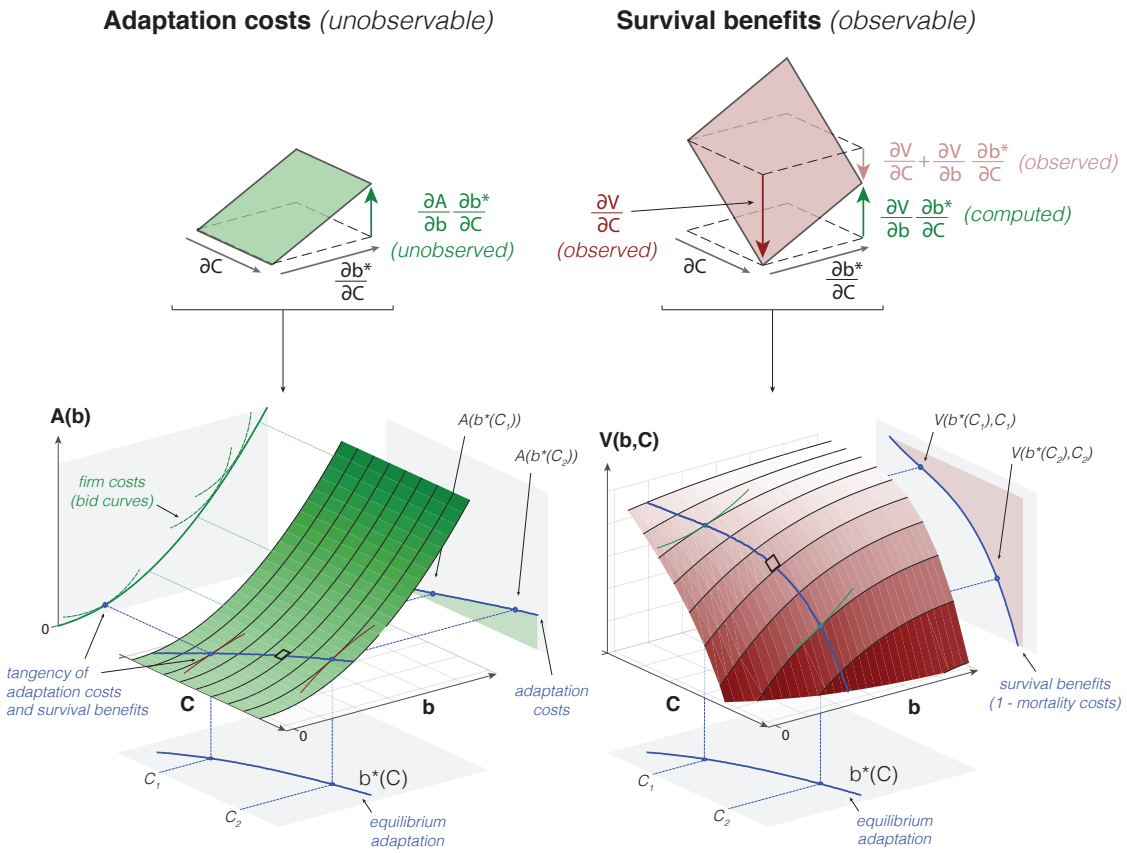

FIGURE 4.A.1

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.

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

A key innovation to our analysis is fully accounting for adaptation costs $A(b)$ even though neither $A(.)$ nor b is observed. Indeed, there may be a very large, even infinite, number of ways that populations adapt to climate that cannot be feasibly enumerated by the econometrician. All the econometrician can observe are the effects of adaptation on survival probability $1 - \tilde{f}$. If a climate were gradually warmed from C_1 to C_2, individuals would continuously respond by adapting along $b^*(C)$ and traveling up the cost surface in the lower left panel, eventually incurring costs $A(b^*(C_2))$ rather than the initial costs $A(b^*(C_1))$ that

they incurred prior to warming. We point out that the change in this total adaptation cost $A(b^*(C_2)) - A(b^*(C_1))$ can be inferred based only on the shape of the benefits surface along the equilibrium, information that is recoverable by the econometrician.

To show this, at the top of Figure 4.A.1 we draw tangency planes for both the costs and benefits surfaces for a single location along the equilibrium adaptation locus between C_1 and C_2, indicated by black squares on the two surfaces in the lower left and lower right panels. Both tangency planes span an area $\partial C \times \frac{\partial b^*}{\partial C}$, indicating how much additional adaptation populations undertake ($\frac{\partial b^*}{\partial C}$) for an exogenous change in climate (∂C), changes that would cause them to traverse each of these planes from their respective left-most corner to their right-most corner. The corresponding change in survival benefits is $\frac{dV}{dC} = \frac{\partial V}{\partial C} + \frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$ (downward pink arrow on the right), which the econometrician can observe by computing the change in survival probability due to climate between two adjacent locations after allowing them both to fully adapt to their respective climates. If the cooler location is heated by ∂C but not permitted to adapt, its survival benefits change by $\frac{\partial V}{\partial C}$ (downward red arrow), a counterfactual outcome that the econometrician can compute by simulating a warmer environment without allowing for adaptation. The difference between these two changes is equal to the benefits of marginal adaptations $\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$ (upward green arrow, right panel). Along the equilibrium $b^*(C)$, these marginal benefits of adaptation must equal their marginal costs, thus we know the corresponding increase in unobserved adaptation costs $\frac{\partial A}{\partial b}\frac{\partial b^*}{\partial C}$ (upward green arrow, left panel) must be equal in magnitude to $\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$. By continuously computing and differencing the total and partial derivatives of V with respect to an incremental change in climate dC (i.e., $\frac{dV}{dC} - \frac{\partial V}{\partial C}$), we recover the marginal benefits of unobserved incremental adaptations ($\frac{\partial V}{\partial b}\frac{\partial b^*}{\partial C}$), which we know must also equal their marginal costs ($\frac{\partial A}{\partial b}\frac{\partial b^*}{\partial C}$). Then by integrating these marginal costs with respect to the climate (shown in the $A \times C$ plane of the lower left panel) we can compute the total change in adaptation costs $A(b^*(C_2)) - A(b^*(C_1))$ for the non-marginal change in climate from C_1 to C_2. This intuition holds for an unknown number of margins of adaptation and a climate of arbitrary dimension, which we allow for in the main text and in derivations below.

4.A.2 Surplus generated from compensatory investments

As discussed in the main text, the equivalence of *marginal* adaptation benefits and *marginal* adaptation costs at each point along the equilibrium pathway $\boldsymbol{b}^*(Y, \boldsymbol{C})$ (Equation 4.4) 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 4.5.3).

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

climate changing from time period $t = 1$ to $t = 2$ can be written as:[65]

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

$$= -\int_{\boldsymbol{b}^*(\boldsymbol{C}_1)}^{\boldsymbol{b}^*(\boldsymbol{C}_2)} VSL \frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_2)}{d\boldsymbol{b}} d\boldsymbol{b} - \int_{\boldsymbol{b}^*(\boldsymbol{C}_1)}^{\boldsymbol{b}^*(\boldsymbol{C}_2)} \frac{\partial A(\boldsymbol{b})}{\partial \boldsymbol{b}} d\boldsymbol{b}^* \tag{4.A.1}$$

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

The first term in the definition of adaptation surplus in Equation 4.A.1 is the total benefits of adaptation, defined as [minus] the mortality effects of climate $\boldsymbol{C}_2$ with optimal adaptation (i.e. $\boldsymbol{b}^*(\boldsymbol{C}_2)$) minus the mortality effects of that same climate, but with adaptation fixed at its initial level (i.e., $\boldsymbol{b}^*(\boldsymbol{C}_1)$). The second term is the total costs of adaptation, defined as the adaptation costs under optimal adaptation in climate $\boldsymbol{C}_2$ minus adaptation costs under optimal adaptation in the initial climate $\boldsymbol{C}_1$. Adaptation benefits (the first term) can be computed by integrating $\frac{d\tilde{f}(\boldsymbol{b}, \boldsymbol{C}_2)}{d\boldsymbol{b}}$, the marginal mortality effect of adaptation evaluated at fixed climate $\boldsymbol{C}_2$. Note that this integration is not computed over the optimal pathway, as the climate is fixed at $\boldsymbol{C}_2$ and any $\boldsymbol{b} \neq \boldsymbol{b}^*(\boldsymbol{C}_2)$ is thus off-equilibrium. Adaptation costs (the second term) can be computed by integrating marginal adaptation costs of $\boldsymbol{b}$ along the optimal pathway $\boldsymbol{b}^*(\boldsymbol{C})$.

The expression for adaptation surplus in Equation 4.A.1 is represented as the difference between two integrals, each computed over the unobserved choice vector $\boldsymbol{b}$. To empirically identify adaptation surplus, we aim to rewrite this expression as a difference between integrals which are computed over the multi-dimensional climate $\boldsymbol{C}$, which changes over time t. 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 4.2 in the main text), total adaptation costs, the second term in Equation 4.A.1, can be rewritten as an integral over time using a simple change of variables. However, rewriting total adaptation benefits, the first term in Equation 4.A.1, as an integral over time (and hence, climate $\boldsymbol{C}$) requires multiple steps, which we outline below.

To see how we construct an empirically tractable expression for total adaptation benefits (first term in Equation 4.A.1), we first consider a visual illustration. Figure 4.A.2 shows the construction of total adaptation benefits using the same notation and format as the lower right panel of Figure 4.A.1. As in Figure 4.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

[65] 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.

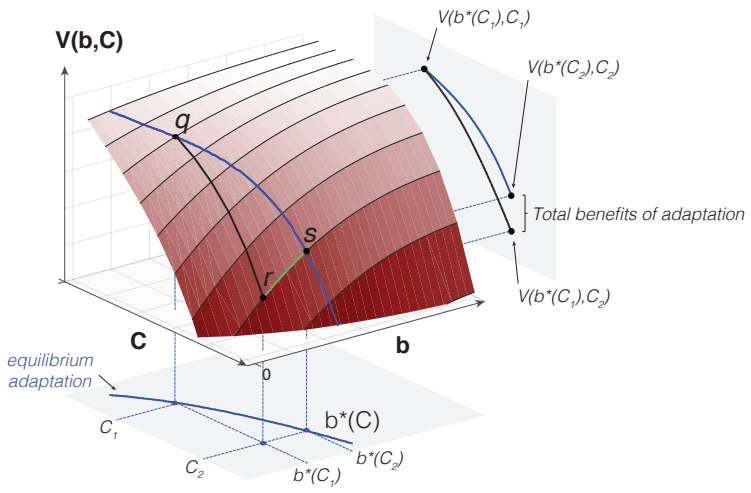

FIGURE 4.A.2

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 $\{b^*(C_1), C\}$ is drawn in black (line $q \to r$). To derive the total benefits of adaptation under a change in climate from C_1 to C_2, we integrate the surface along the green line (line $r \to s$), evaluating changes in survival benefits at a fixed climate C_2, as adaptation evolves from $b^*(C_1)$ to $b^*(C_2)$. The magnitude of total adaptation benefits is shown on the $V \times C$ plane on the right panel.

want to quantify the vertical difference between points s and r (i.e., $s - r$), which can be computed empirically as the vertical difference $q - r$ minus the difference $q - s$. To see why, note that the total benefits of adaptation incurred under a climate change from C_1 to C_2 are represented by the vertical difference between points s and r (shown on the $V \times C$ plane on the right panel), because this height measures the total mortality benefits realized from optimally investing in adaptation $b^*(C_2)$ when experiencing climate C_2, instead of holding adaptation fixed at its initial level $b^*(C_1)$. This difference can be computed in two ways. First, total benefits of adaptation can be computed by traversing along the off-equilibrium green line between points r and s; that is, by holding C fixed at C_2 and integrating $V(b, C)$ over b from $b^*(C_1)$ to $b^*(C_2)$. This integration along the green line represents the definition of total adaptation benefits written in Equation 4.A.1. However, this same vertical distance can alternatively be calculated by traversing along the off-equilibrium black line between points q and r (i.e., holding b fixed at $b^*(C_1)$ and integrating $V(b, C)$ over C from C_1 to C_2), and then subtracting off the value of the survival impacts of the optimal pathway from C_1 to C_2 (i.e., the height of the surface at point q minus point s). This integration over C (twice) is empirically identifiable, as changes in climate can, in principle, be observed.

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 4.A.1 writes total adaptation benefits using a path along the surface in the $\boldsymbol{b}$ dimension between the end points $\{\boldsymbol{b}^*(\boldsymbol{C}_2), \boldsymbol{C}_2\}$ and $\{\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_2\}$.[66] However, as

[66]Note that while Figure 4.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 4.A.1. For parsimony, we use the latter notation here and in the subsequent expressions.

discussed above, we cannot compute traversing of this path, as changes in $\boldsymbol{b}$ are unobservable. Thus, we need to define an alternative computable path between the same endpoints. If we can construct a loop on the surface that connects the two endpoints, the sum of the desired segment and the remaining segments defining that loop must equal zero, because the line integral over any closed loop L must, by construction, equal zero. We can then rearrange this identity to isolate the computable segments of the loop, allowing us to back out the unobserved segment defining the total benefits of adaptation.

We define such a loop that begins at $\{\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_1\}$ (analogous to point q in Figure 4.A.2) and traverses along the off-equilibrium path from $\boldsymbol{C}_1$ to $\boldsymbol{C}_2$ with adaptation fixed at $\boldsymbol{b}^*(\boldsymbol{C}_1)$ (analogous to the black line between q and r in Figure 4.A.2). In the second segment, it traverses in the $\boldsymbol{b}$ dimension, holding $\boldsymbol{C}$ fixed at $\boldsymbol{C}_2$, to arrive at $\{\boldsymbol{b}^*(\boldsymbol{C}_2), \boldsymbol{C}_2\}$ (analogous to the green line in Figure 4.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 4.A.2):

$$\begin{aligned}\oint_L \nabla[VSL\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\cdot\partial\boldsymbol{b}\partial\boldsymbol{C} &= \int_1^2 VSL\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C}_t)}{\partial\boldsymbol{C}}\frac{d\boldsymbol{C}_t}{dt}dt + \int_{\boldsymbol{b}^*(C_1)}^{\boldsymbol{b}^*(C_2)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_2)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &\quad + \int_2^1 VSL\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t),\boldsymbol{C}_t)}{d\boldsymbol{C}}\frac{d\boldsymbol{C}_t}{dt}dt \\ &= 0 \end{aligned} \tag{4.A.2}$$

By rearranging Equation 4.A.2 (including changing the direction of integration for the third segment), we can use this closed loop, which is composed of two computable segments and a third that is unobservable, to calculate the total benefits of adaptation:

$$\begin{aligned}\textit{Total adaptation benefits} &= -\int_{\boldsymbol{b}^*(C_1)}^{\boldsymbol{b}^*(C_2)} VSL\frac{d\tilde{f}(\boldsymbol{b},\boldsymbol{C}_2)}{d\boldsymbol{b}}d\boldsymbol{b} \\ &= -\int_1^2 VSL\left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t),\boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C}_t)}{\partial\boldsymbol{C}}\right]\frac{d\boldsymbol{C}_t}{dt}dt \end{aligned} \tag{4.A.3}$$

Using Equation 4.A.3 and a change of variables to rewrite the total costs of adaptation as an integral over $\boldsymbol{C}$, we can rewrite Equation 4.A.1 as:

$$\begin{aligned}\textit{Adaptation surplus}\,(\boldsymbol{C}_1\to\boldsymbol{C}_2) = -\int_1^2 VSL\left[\underbrace{\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t),\boldsymbol{C}_t)}{d\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/ adaptation}}} - \underbrace{\frac{\partial\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1),\boldsymbol{C}_t)}{\partial\boldsymbol{C}}}_{\substack{\text{mortality risk}\\ \text{w/o adaptation}}}\right]\frac{d\boldsymbol{C}_t}{dt}dt \\ - \int_1^2 \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_t))}{\partial\boldsymbol{b}}\frac{\partial\boldsymbol{b}_t^*}{\partial\boldsymbol{C}}\frac{d\boldsymbol{C}_t}{dt}dt \end{aligned} \tag{4.A.4}$$

While the total adaptation benefits term in Equation 4.A.4 (the first term) is composed of values that are, in principle, empirically identifiable, the adaptation cost expression (the

second term) remains unobservable because the net cost function $A(\boldsymbol{b}^*(\boldsymbol{C}))$ is unknown. Thus, we take a final step to rewrite the entire adaptation surplus expression in Equation 4.A.4 in terms of objects that are measurable, using Equation 4.6 from the main text to substitute for the object $\int_1^2 \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}_t))}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} \frac{d\boldsymbol{C}_t}{dt} dt$:

$$\begin{aligned} \textit{Adaptation surplus}\,(\boldsymbol{C}_1 \to \boldsymbol{C}_2) &= -\int_1^2 VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_t}{dt} dt \\ &\quad + \int_1^2 VSL \left[\frac{d\tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_t}{dt} dt \\ &= \int_1^2 VSL \left[\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \right] \frac{d\boldsymbol{C}_t}{dt} dt \end{aligned} \tag{4.A.5}$$

In Equation 4.A.5, the first term inside the integral represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, but holding adaptation actions fixed at the levels that were optimal under the original climate, $\boldsymbol{C}_1$. In contrast, the second term represents the marginal mortality effect of a change in climate evaluated at climate $\boldsymbol{C}$, allowing adaptation actions $\boldsymbol{b}^*(\boldsymbol{C})$ to evolve optimally with the changing climate. Note that because the second term is a partial derivative, its integral is *not* the total change in the mortality rate. While the two partial derivatives in Equation 4.A.5 will be identical when $\boldsymbol{C} = \boldsymbol{C}_1$, if they diverge at some point after $\boldsymbol{C}$ warms beyond $\boldsymbol{C}_1$, then surplus will be nonzero. Thus, a sufficient condition for positive surplus is:

$$\frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_1), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} > \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}_t), \boldsymbol{C}_t)}{\partial \boldsymbol{C}} \quad \forall t \in (1, 2] \tag{4.A.6}$$

This condition says that mortality risk must rise more with changes in the climate at lower levels of adaptation. If this condition holds, the difference between the two partial derivatives in Equation 4.A.6 is weakly positive, and the total adaptation surplus over the climate trajectory $\boldsymbol{C}_1 \to \boldsymbol{C}_2$ is positive.

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

In Section 4.4.5, we describe how we use econometric estimation of Equation 4.11 in combination with climate model projections to construct empirical estimates of changes in adaptation costs due to climate change, following the theoretical derivation in Section 4.2. Here, we provide some additional details on this implementation.

Theoretically, adaptation costs can be computed by taking the difference between the total and partial derivative of expected mortality risk with respect to changes in the climate (Equation 4.6), and integrating this difference. To empirically construct an estimate of these

costs, we begin by taking expectations of Equation 4.11 over weather realizations $\boldsymbol{T}$, to specify our empirically estimated *expected* mortality risk for an age group a in region r for year t:

$$\hat{\bar{f}}(.)_{art} \equiv \mathrm{E}[\hat{f}(.)_{art}] = \mathrm{E}[\underbrace{\hat{g}_a\left(\boldsymbol{T}_{rt} \mid TMEAN_{rt}, \log(GDPpc)_{rt}\right)}_{\hat{g}_{art}(\cdot)}] + ... \tag{4.A.7}$$

where we omit the various estimated terms orthogonal to temperature, which fall out after differentiation. Recall that the estimates $\hat{g}_{art}(\cdot)$ describe the shape of the annual response function in region r and year t for age group a, taking as inputs the summary climate parameter $TMEAN$ and log income per capita, where the coefficients used to construct $\hat{g}_{art}(\cdot)$ are recovered from the regression in Equation 4.11. 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:

$$\frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}} = \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} = \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} \tag{4.A.8}$$

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

In contrast, the total derivative of expected mortality risk with respect to a change in climate reflects endogenous adaptations through adjustments to $\boldsymbol{b}$, which in turn change the shape of the response function. Our econometric framework captures these effects through the $TMEAN$ interactions in $g(\cdot)$, which modify the shape of a region's response function based on long run average conditions. When we compute the total derivative of $\hat{\bar{f}}(.)$ 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:

$$\frac{d\hat{\bar{f}}_t}{d\boldsymbol{C}} = \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{b}} \frac{\partial \boldsymbol{b}_t^*}{\partial \boldsymbol{C}} = \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{T}} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \frac{\partial \hat{\bar{f}}_t}{\partial TMEAN} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}}$$
$$= \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial \boldsymbol{T}}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial \boldsymbol{T}_t}{\partial \boldsymbol{C}} + \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \tag{4.A.9}$$

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

As shown in Equation 4.12 in the main text, the *difference* between the total and partial derivatives of expected mortality risk with respect to the climate is thus the difference between Equations 4.A.9 and 4.A.8:

$$\frac{d\hat{\bar{f}}_t}{d\boldsymbol{C}} - \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}} = \left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_t, Y_t} \frac{\partial TMEAN_t}{\partial \boldsymbol{C}} \tag{4.A.10}$$

The righthand side of Equation 4.A.10 is fully computable for years in our projection using a combination of empirically estimated parameters, $\hat{g}(\cdot)$, and climate projections, $\{\boldsymbol{T}, TMEAN\}$. Substituting Equation 4.A.10 into Equation 4.6 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 $t = 1$ to $t = 2$:

$$\widehat{A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - A}(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1))$$
$$\approx -\int_1^2 VSL_t \left[\frac{d\hat{\bar{f}}_t}{d\boldsymbol{C}} - \frac{\partial \hat{\bar{f}}_t}{\partial \boldsymbol{C}}\right] \frac{d\boldsymbol{C}_t}{dt} dt$$
$$\approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \left(\left.\frac{\partial \mathrm{E}[\hat{g}]}{\partial TMEAN}\right|_{\boldsymbol{C}_\tau, Y_2}\right) (TMEAN_\tau - TMEAN_{\tau-1})$$
$$\approx -\sum_{\tau=t_1+1}^{t_2} VSL_\tau \, \hat{\boldsymbol{\gamma}}_{\mathbf{1}} \mathrm{E}[\boldsymbol{T}]_\tau \left(TMEAN_\tau - TMEAN_{\tau-1}\right), \tag{4.A.11}$$

where the second equality results from substitution of Equation 4.A.10 into Equation 4.6 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 4.A.11. This is because the goal of the calculation is to develop an estimate of the additional adaptation expenditures incurred due to the changing climate only. Changes in adaptation expenditures due to rising incomes may change mortality risk under climate change, but these changes are voluntary and are not the consequence of the changing climate, and are therefore not included in our calculation of

the total mortality-related costs of climate change. These income effects are accounted for econometrically in the estimation of Equation 4.11 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 $t_1 = 2015$ as the baseline year, for each of our 33 high-resolution climate model projections.

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

Throughout the main text, we construct estimates of adaptation costs derived from a representative agent's problem in which utility is a function only of a consumption good x. In this simple model (see Equation 4.1), 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 4.4.5 of the main text does not change.

As in Section 4.2 of the main text, we consider a single representative global agent who faces mortality risk $f_t = f(\boldsymbol{b}_t, \boldsymbol{c}_t)$ in each period t. We further assume there exists some numeraire good x_t for which utility $u(x_t, \boldsymbol{b}_t)$ is quasilinear. As above, this agent maximizes utility conditional on *expected* weather realizations, subject to an exogenous budget constraint and exogenously determined emissions. Letting $\tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t) = \mathbb{E}_{\boldsymbol{c}_t}[f(\boldsymbol{b}_t, \boldsymbol{c}(\boldsymbol{C}_t)) \mid \boldsymbol{C}_t]$ represent the expected probability of death, the agent solves:

$$\max_{\boldsymbol{b}_t, x_t} \; u(x_t, \boldsymbol{b}_t) \left[1 - \tilde{f}(\boldsymbol{b}_t, \boldsymbol{C}_t)\right] \;\; s.t. \; Y_t \geq x_t + A(\boldsymbol{b}_t), \tag{4.A.12}$$

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

Rearranging the agent's first order conditions and using the conventional definition of

the VSL,[67] we can write:

$$\underbrace{\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}_t} - \frac{\partial u/\partial \boldsymbol{b}}{\partial u/\partial x}}_{\text{net marginal cost of } \boldsymbol{b}} = \frac{-u(x_t^*, \boldsymbol{b}_t^*)}{\partial u/\partial x[1-\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)]} \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}} = \underbrace{-VSL_t \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{b}}}_{\substack{\text{marginal survival} \\ \text{benefit of } \boldsymbol{b}}} \tag{4.A.13}$$

This expression governs expenditures on adaptation. Its righthand side is the product of the negative of the VSL and the marginal change in expected mortality risk due to a change in adaptation, so it represents the expected marginal benefit (in dollar value) of adjusting $\boldsymbol{b}$ through its effect on mortality risk. This object is identical to its counterpart in Equation 4.4 in the main text. The lefthand side has two parts. The first term represents the marginal cost of all pecuniary expenditures incurred due to a marginal change in adaptation $\boldsymbol{b}$, such as spending on units of air conditioning. The second term represents [minus] the dollar value of all non-mortality marginal utility benefits or costs derived from a marginal change in $\boldsymbol{b}$, such as the utility of enjoying air conditioning or the disutility of exercising at midnight to avoid daytime heat (note that this object is expressed in dollars of WTP by dividing through by the marginal utility of consumption, $\partial u/\partial x$). Together, these two terms can be interpreted as the *net* marginal cost of all adaptive actions composing the composite $\boldsymbol{b}$, because non-mortality marginal benefits and costs are removed from the marginal pecuniary expenditures term $\partial A/\partial \boldsymbol{b}$.

Both terms composing net marginal costs in Equation 4.A.13 are unobservable. In contrast, the marginal survival benefit can be rewritten as the product of the negative of the VSL and the difference between the total and partial derivatives of mortality risk with respect to the climate – i.e., $\frac{d\tilde{f}}{d\boldsymbol{C}} - \frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$ (see Equation 4.5). As discussed in the main text, we develop an empirical model that allows us to estimate both the total and partial derivates, rendering the marginal survival benefits empirically tractable.

In the main text, we use this insight to develop an expression for the additional adaptation costs incurred as the climate changes gradually, which is composed of observable terms. This expression remains unchanged under the alternative model specification described here, with the exception that the adaptation costs recovered are *net* of utility benefits or costs incurred due to changes in optimal adaptation $\boldsymbol{b}^*$. Here, the additional net adaptation costs incurred as the climate changes gradually from period $t = 1$ to period $t = 2$ are:

$$\begin{aligned} A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - A(\boldsymbol{b}^*(Y_2, \boldsymbol{C}_1)) - \frac{1}{\partial u/\partial x}\left[u(x^*(Y_2, \boldsymbol{C}_2), \boldsymbol{b}^*(Y_2, \boldsymbol{C}_2)) - u(x^*(Y_2, \boldsymbol{C}_1), \boldsymbol{b}^*(Y_2, \boldsymbol{C}_1))\right] \\ = \int_1^2 \left[\frac{\partial A(\boldsymbol{b}_t^*)}{\partial \boldsymbol{b}} - \frac{\partial u(x_t^*, \boldsymbol{b}_t^*)/\partial \boldsymbol{b})}{\partial u(x_t^*, \boldsymbol{b}_t^*)/\partial x)}\right] \frac{d\boldsymbol{b}_t^*}{d\boldsymbol{C}} \frac{d\boldsymbol{C}_t}{dt} dt \\ = -\int_1^2 VSL_t \left[\frac{d\tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{d\boldsymbol{C}} - \frac{\partial \tilde{f}(\boldsymbol{b}_t^*, \boldsymbol{C}_t)}{\partial \boldsymbol{C}}\right] \frac{d\boldsymbol{C}_t}{dt} dt, \end{aligned} \tag{4.A.14}$$

[67] 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: $VSL = \frac{u(x)}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\partial u/\partial x}$.

where the last line relies on substitution from Equations 4.A.13 and 4.5. The righthand side of Equation 4.A.14 can be approximated empirically as shown in Section 4.4.5 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 derivation of the mortality partial social cost of carbon shown in main text Section 4.2, which relies on an empirically tractable estimate of adaptation costs, is unchanged under this alternative model specification. However, as in Equation 4.A.14, the mortality partial SCC should be interpreted here as the marginal willingness to pay to avoid a marginal increase in greenhouse gas emissions inclusive of adaptation benefits and *net* adaptation costs.

4.B Data appendix

4.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.

4.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).[68] We use data from 1997-2010 and aggregate the monthly data to annual frequency. Data were provided for both place of death and place of residence. As with all subsequent datasets, we assign weather exposure to deaths in our data at the place of residence, as this is provided for all sources. Data were downloaded in 5-year age groups which were then aggregated to the age groups used in the analysis. ADM2-level populations were obtained from the same source. Administrative boundary files were downloaded from GADM (Global Administrative Areas, 2012). Brazilian death data as downloaded contained a number of ADM2 units with missing values for deaths and no values of zero, implying that these are a mix of true zeros and missing values. To ascertain whether they are more likely to be the former, we examined the relationship between death counts and population in all ADM2 units, and then in only those ADM2 units that ever show a missing value in any year. We found that missing values are more likely to occur in low population ADM2 units, suggesting that these are places that should have recorded zero deaths. We consequently treat these missing values as zeros, but in robustness tests find that treating them as missing does not substantially change any of our results.

4.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).[69] 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

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

[69] Data are available here: http://www.deis.cl/bases-de-datos-defunciones/.

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

4.B.1.3 China

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

4.B.1.4 European Union

The EU maintains a centralized statistical database known as EuroStat (Eurostat, 2013)[71] 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.[72] 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 4.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.

[70] Data are available here: http://www.ine.cl/estadisticas/demograficas-y-vitales

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

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

4.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)[73] 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 Tables 4.2 and 4.D.1. 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).

4.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 4.4). However, these data are used to assess the external validity of our extrapolation methods, as discussed in Appendix 4.D.6.

4.B.1.7 Japan

Japanese data on mortality and population at the prefecture-year[74] level were obtained from the National Institute of Population and Social Security Research[75] 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).

4.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.[76] 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

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

[74] 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.

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

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

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.[77] 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.

4.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)[78] 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.[79]

4.B.1.10 Aggregate data

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

[77] 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.

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

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

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

Table 4.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)[80]
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)

4.B.2 Climate data

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

4.B.2.1 Historical climate data

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

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

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

average precipitation for all grid cells globally.

Berkeley Earth Surface Temperature The Berkeley Earth Surface Temperature (BEST) dataset provides temperatures from 1701-2018 over land from a combination of observational records (Rohde et al., 2013), with spatially disaggregated data available from 1753.[81] During the time periods used within this paper (varying between 1957-2014), as many as 37,000 station records, representing 14 separate databases of station data, are incorporated into the BEST data. Station data are incorporated using a kriging methodology that allows for the incorporation of more stations with shorter time series than other well-known global surface temperature interpolation data (like the UDEL temperature dataset). In particular, the spatial averaging method uses close neighbors of a station to identify discontinuities in a particular time series that may be due to instrumental change or re-positioning, and decreases the influence of these changes in the spatially averaged grid (Rohde et al., 2013). This does have the potential drawback of over-smoothing the spatial heterogeneity in temperatures (National Center for Atmospheric Research Staff (Eds), 2015). BEST data are provided at daily frequency on a $1°\times1°$ resolution grid, and we utilize the daily average 2m air temperature variable for each grid cell.

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

4.B.2.2 Climate projection data

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

To address both of these limitations, we use a high-resolution ($0.25°$ X $0.25°$) set of global, bias-corrected climate projections produced by NASA Earth Exchange (NEX): the Global Daily Downscaled Projections (GDDP) (Thrasher et al., 2012).[82] 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

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

[82] 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).

(Taylor, Stouffer, and Meehl, 2012). The statistical downscaling algorithm used to generate the NEX-GDDP dataset is the Bias-Correction Spatial Disaggregation (BCSD) method (Wood et al., 2004; Thrasher et al., 2012), which was developed 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.

4.B.2.3 SMME and model surrogates

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

The SMME method first divides the unit interval [0,1] into a set of bins. For this analysis, the bins are centered at the 1^{st}, 6^{th}, 11^{th}, 16^{th}, 33^{rd}, 50^{th}, 67^{th}, 82^{nd}, 89^{th}, 94^{th}, and

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

If the number of CMIP5 models in a bin is less than 2, surrogate models are generated to raise the total number of models to 2 in that bin. The surrogate models are produced by using the projected annual GMST of the SCM that is consistent with the bin's central quantile to scale the spatial pattern of a selected CMIP5 model, then adding the intercept and residual from the same model. There are two cases of selecting CMIP5 models for pattern and residual. When there is only one CMIP5 model in a bin, an additional model is selected that has a GMST projection close to GMST in the bin and a precipitation projection over the region of interest complementary to the model already in the bin (i.e., if the model in the bin is relatively dry, then a relatively wet pattern is selected, and vice versa.) When there is no CMIP5 model, two models are picked with GMST projections close to that of the bin, with one model being relatively wet and one being relatively dry. In the final probabilistic distribution, the total weight of the bin is equally divided among the CMIP5 models and surrogate models in the bin. For instance, if four models are in the bin centered at the 30th percentile, bounded by the 20th – 40th percentiles, each will be assigned a probability of $20\% \div 4 = 5\%$. The resulting distribution of GMST for all members of the SMME is shown in Figure 4.2B.

4.B.2.4 Aggregation of gridded climate data to administrative boundaries

We link gridded historical climate data to administrative mortality records by aggregating grid cell information to the same spatial and temporal level as the mortality records (see Table 4.1). 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 4.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 4.10 and 4.11. In this case, we begin with data on average temperatures for each day d at each grid cell z, generating observations T_{zd}. These grid-level values must then be aggregated to the level of an administrative unit i in year t. To do this, we first raise grid-level temperature to the power p, computing $(T_{zd})^p$ for $p \in \{1, 2, 3, 4\}$. We then take a spatial average of these values

over administrative unit i, weighting the average by grid-level population (and accounting for fractional grid cells that fall partially within administrative units). Population weights are time-invariant and calculated from the 2011 Landscan dataset (Bright et al., 2012). We then sum these daily polynomial terms T_{zd}^p over days in the year t. The vector of annual, administrative-level-by-year temperature variables we use for estimation is thus:

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

where w_{zi} is the share of i's population that falls into grid cell z, and where superscripts indicate polynomial powers. This nonlinear transformation performed prior to aggregation allows the aggregated measure of temperature to capture grid-by-day level exposure to very hot and very cold temperatures. In the econometric estimation of Equations 4.10 and 4.11, quadratic polynomials in precipitation are similarly calculated and weighted averages are taken over administrative units. In Appendix Figure 4.D.1, 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.

4.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.

4.B.3.1 Historical income data

Our main specification (Equation 4.11) 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**.[83] 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 4.B.1).

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

- **Eurostat (2013) subnational GDP**.[84] 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,[85] 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 4.11. 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:

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

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

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 4.11 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 4.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 4.2A).

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

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

[85] As noted in Appendix 4.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.

typically not annual, leading to an unbalanced panel. To construct annual values of income per capita using the Gennaioli et al. (2014) data, we linearly interpolate between years, before constructing the Bartlett kernel and taking averages across all years. For instances where we need to extrapolate backwards in time (i.e., when mortality data are available earlier than income data), we extrapolate backwards logarthmically. All subnational income data are in constant 2005 dollars PPP. A summary of the available years of data before interpolation is given in Table 4.B.2.

Country	ISO code	Years in mortality sample	Years in income sample[86]
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 4.B.2
Temporal coverage of mortality records and years of available subnational income data.

4.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

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

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.[87]

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 4.B.3.

TABLE 4.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%

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.[88] 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

[87] 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.

[88] 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.

levels across impact regions, each region grows in the future at the same rate as the national country projection from the SSPs. We detail these two steps below.

Estimation of the GDP-nightlights relationship While there exists a growing literature linking economic output to nightlights intensity, we take an unconventional regression approach to recovering this relationship because our goal is to apportion national income within a country, as opposed to predict the level of income at any given location. In particular, we are interested in the ratio $\frac{GDPpc_{rc}}{\sum_{r \in c} w_{rc} GCPpc_{rc}}$ 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 $\frac{GDPpc_{ict}}{\sum_{i \in c} w_{ict} GDPpc_{ict}}$.[89] 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:

$$\frac{GDPpc_{ict}}{\sum_{i \in c} w_{ict} GDPpc_{ict}} = \alpha + \beta ZNL_{ict} + \epsilon_{ict} \tag{4.B.1}$$

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

Allocation of national GDP to impact regions using relative nightlight intensity We use the estimated coefficients from Equation 4.B.1 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.

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

4.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).[90] 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:

$$\widehat{pop}_{rtv} = \begin{cases} pop_{ctv}^{SSP}\left(\frac{pop_{r,2011}^{LandScan}}{\sum_{r\in c} pop_{r,2011}^{LandScan}}\right), & \text{if } r \in C \\ pop_{r,2011}^{LandScan}, & \text{if } r \notin C \end{cases} \tag{4.B.2}$$

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.

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

4.C Spatial units for projection: "Impact regions"

We create a set of custom boundaries that define the spatial units for which location-specific projected damages of climate change are computed. To do so, we utilize politically defined regions, as opposed to a regular grid, as socioeconomic data are generally collected at this scale and because administrative regions are relevant to policy-makers. These regions, hereafter referred to as "impact regions", are constructed such that they are identical to existing administrative regions or are a union of a small number of administrative regions. We use version 2 of the Global Administrative Region dataset (GADM) (Global Administrative Areas, 2012), which contains 218,328 spatial units, to delineate boundaries. However, for computational feasibility and greater comparability across regions, we agglomerate these regions to create a set of 24,378 custom impact regions. To conduct this agglomeration, we establish a set of criteria that ensures these impact regions have approximately comparable populations and are internally consistent with respect to mean temperature, diurnal temperature range, and mean precipitation. A map of these regions is shown in Figure 4.C.1, and we detail this agglomeration algorithm below.

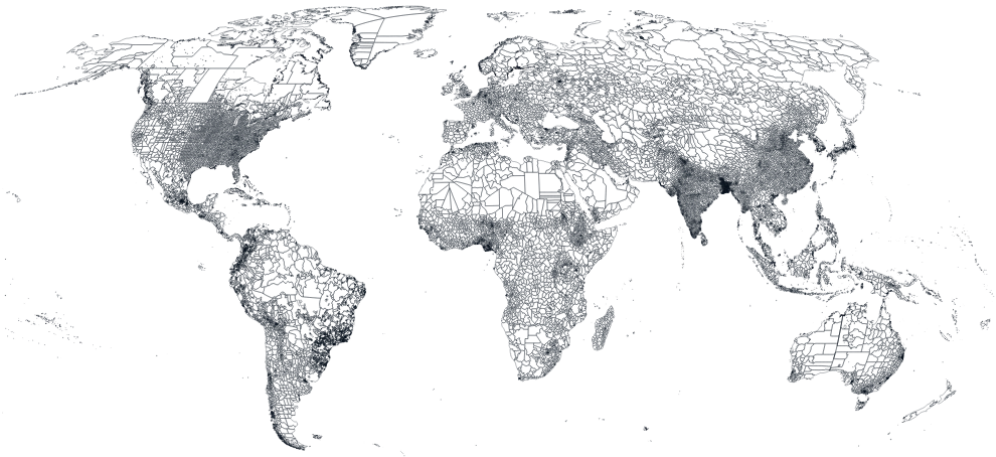

FIGURE 4.C.1
Map of the 24,378 "impact regions" for which location-specific projections are calculated. We use a clustering algorithm to form these regions from the full set of GADM administrative regions, such that they are roughly similar in total population, and so that they are approximately internally homogenous with respect to mean temperature, diurnal temperature range, and mean precipitation.

4.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 4.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 4.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 4.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 4.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 4.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 4.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

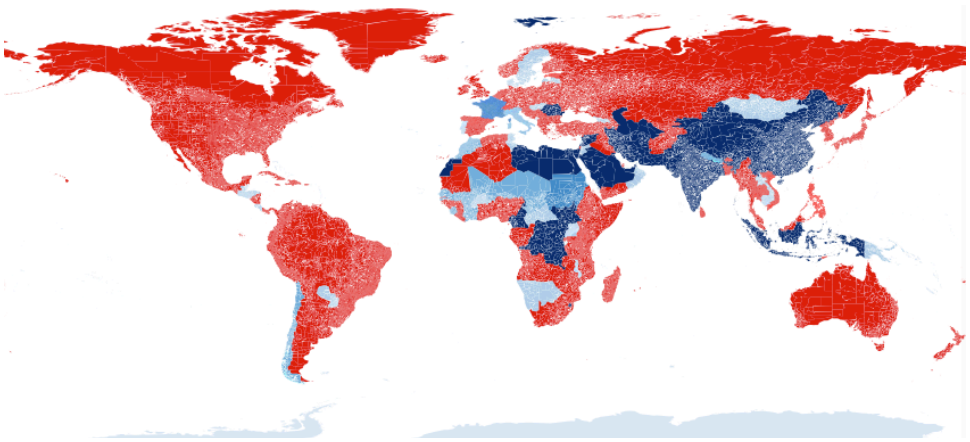

FIGURE 4.C.2

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.

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:

- The set of GADM regions within the agglomeration
- The set of neighboring agglomerated regions
- Population (pop),[91] 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.

[91] Population data are from Landscan (Bright et al., 2012), as in Appendix 4.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.

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

This appendix provides additional illustrations of and tabular results for the main econometric regressions used and discussed throughout the main text (Table 4.D.1, Figures 4.D.2 and 4.D.3, and Table 4.D.3), shows a set of robustness checks for those main results (Figures 4.D.1 and 4.D.4 and Table 4.D.4), and discusses an out-of-sample validation test designed to evaluate the accuracy with which our estimates predict mortality-temperature responses in locations that are not used for estimation (Figure 4.D.5).

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

Figure 4.D.1 displays the results of estimating a version of Equation 4.10 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 4.B.2 for details on these climate datasets). Here we show results for an all-age mortality response $g(\boldsymbol{T}_{it})$ in which an average treatment effect across all age categories is recovered (as in Table 4.2 in the main text). The four functional forms estimated are fourth-order polynomials, bins of daily average temperature, restricted cubic splines, and piecewise linear splines (details on these functional forms are in Section 4.4 of the main text). All regressions include $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects, and are population weighted. Robustness to alternative fixed effects specifications is shown in Table 4.2.

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

Table 4.D.1 displays the regression results from estimation of Equation 4.10 in the main text. These regression results reveal substantial heterogeneity across age groups within our multi-country sample (responses for column (2) are plotted in Figure 4.4). Selecting column (2) as our preferred specification, people over the age of 64 experience approximately 4.7 extra deaths per 100,000 for a day at 35°C compared to a day at 20°C, a substantially larger effect than that for younger cohorts, which exhibit little response.

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

The administrative regions in our sample display substantial heterogeneity in incomes, climates, and demographics, among many other characteristics. To begin to examine this heterogeneity before estimating the two-factor model of heterogeneous mortality-temperature responses in Equation 4.11, here we investigate variation in mortality responses to temper-

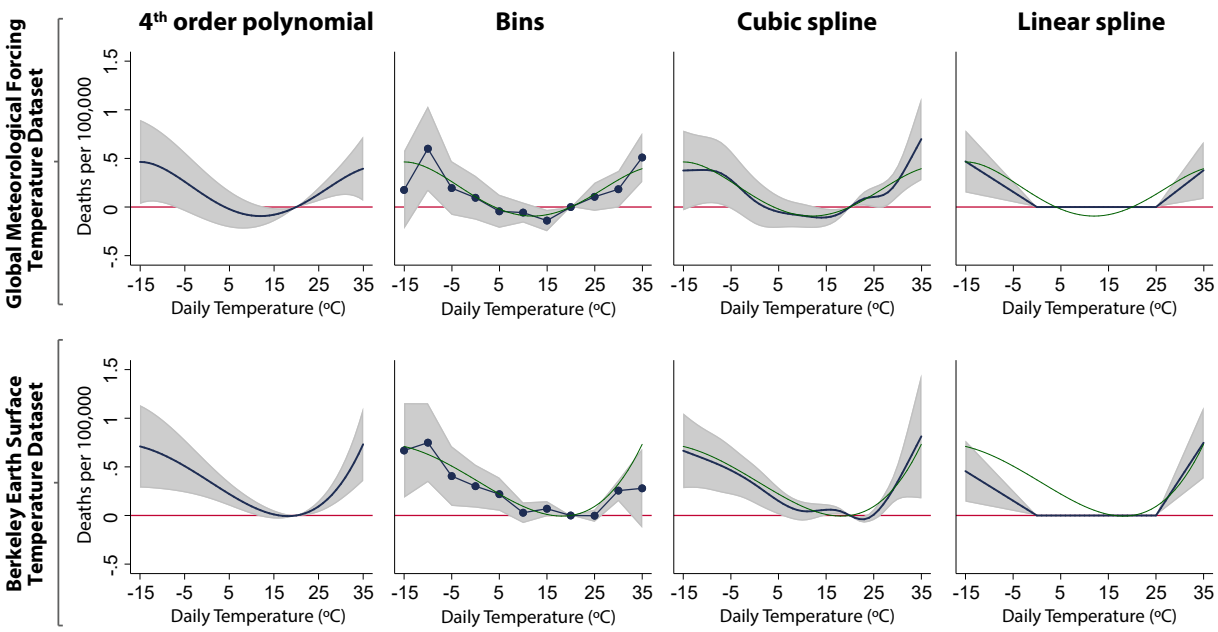

FIGURE 4.D.1

Robustness of the all-age mortality-temperature relationship to alternative functional forms and to different historical climate datasets. Row 1 shows the mortality-temperature response function as estimated using daily temperature and precipitation data from the Global Meteorological Forcing Dataset (GMFD). Row 2 shows the same response, using daily temperatures from Berkeley Earth Surface Temperature (BEST), and monthly precipitation from the University of Delaware. Each column displays a distinct functional form, with the fourth-order polynomial shown in column 1 overlaid in teal on each subsequent column. See Section 4.4 for details on each functional form.

ature at the country level across our sample.[92] Table 4.D.2 displays differential mortality-temperature response functions for each of the 9 countries or regions (i.e., the EU, which is composed of 33 countries) in our data.

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

TABLE 4.D.1

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

TABLE 4.D.2

Heterogeneity by country in the mortality-temperature response function. Regression estimates shown are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a sample that was winsorized at the 1% level. Point estimates indicate the effect of a single day at each daily average temperature value shown, relative to a day with an average temperature of 20°C. Country-specific coefficients are generated by interacting all climate variables and fixed effects with country dummies. Point estimates are only shown for daily average temperatures that are actually experienced in each country over our sample period.

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

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

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

The estimation of Equation 4.11 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 4.11 in practice, note that in Equation 4.10, 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 4.11, 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 4.D.3 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).[93] Although not all of the coefficients would be judged statistically significant by conventional criteria, it is noteworthy that higher incomes and warmer climates are associated with lower mortality consequences of hot days for all age categories. Income and climate are associated with cold day mortality differentially across age groups, with some evidence that higher income locations exhibit more extreme cold day sensitivity for the oldest age group. This relationship may arise due to age being positively correlated with income within the over 64 category, as older individuals are more susceptible to cold-related death risks (Deschênes and Moretti, 2009).

As these terms are difficult to interpret, we visualize this heterogeneity in the main text in Figure 4.5 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 4.D.3. This results in a set of predicted response functions that vary across the joint distribution of income and average temperature within our sample data, shown in Figure 4.5 for the >64 age category. Here we replicate this figure

[93] 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.

TABLE 4.D.3

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

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

Regression includes $age \times ADM2$ fixed effects and $age \times country \times year$ fixed effects. Adjusted $R^2 = 0.933$; N=820,237.

Standard errors clustered at the ADM1 level. *** p<0.01, ** p<0.05, * p<0.1

for the other two age groups in our analysis.

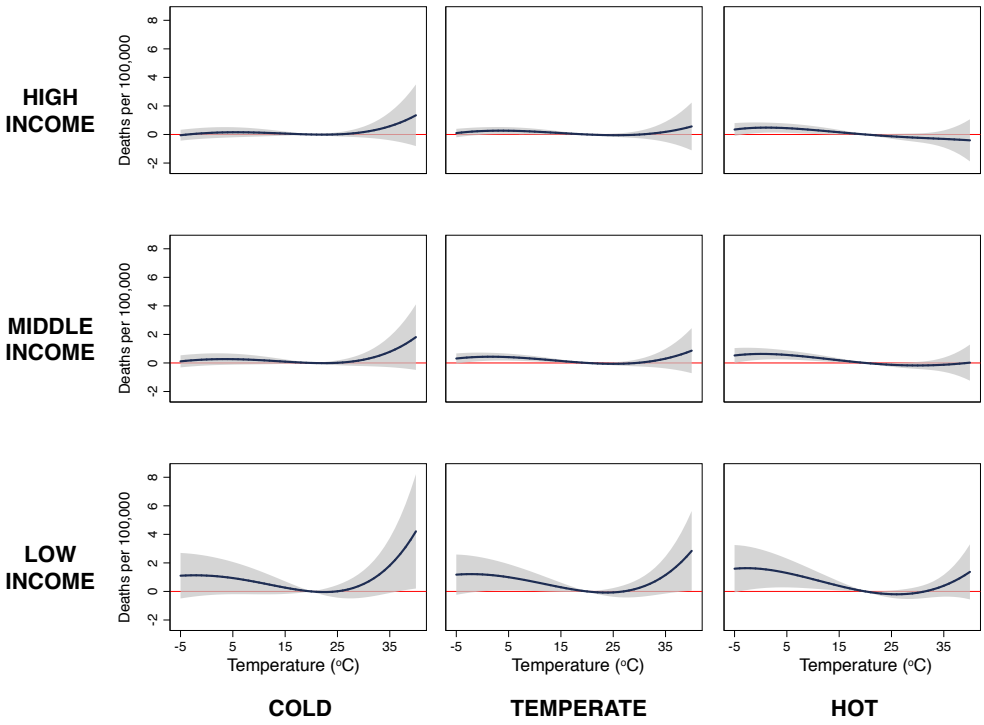

FIGURE 4.D.2

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.

4.D.5 Robustness of estimates of subnational heterogeneity in the mortality-temperature response function to an alternative characterization of long-run average climate

Our primary results rely on a parsimonious representation of the climate: to capture adaptation to long-run climate, we interact our nonlinear temperature variables ($\boldsymbol{T}$) with the long run average annual temperature ($TMEAN$), conditioning on income ($\log(GDPpc)$). In this specification, $TMEAN$ acts as a summary statistic of the long-run average climate, and we find that the mortality sensitivity to high temperatures declines as $TMEAN$ rises. To test the robustness of this finding, here we use a richer characterization of the climate, replacing our climate interaction term $TMEAN$ in Equation 4.11 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 4.11 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

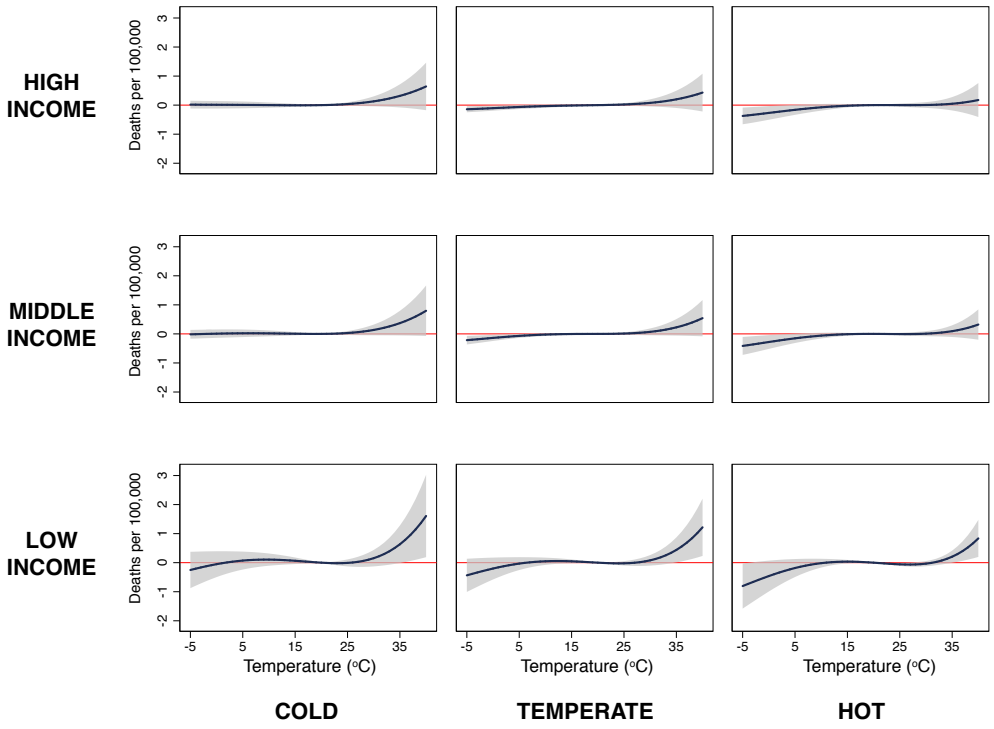

FIGURE 4.D.3

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.

is shown in Table 4.D.4. Consistent with our main results in Table 4.D.3, 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 4.D.4

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 4.D.4 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 4.D.4 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 4.D.3 and 4.D.4, 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.

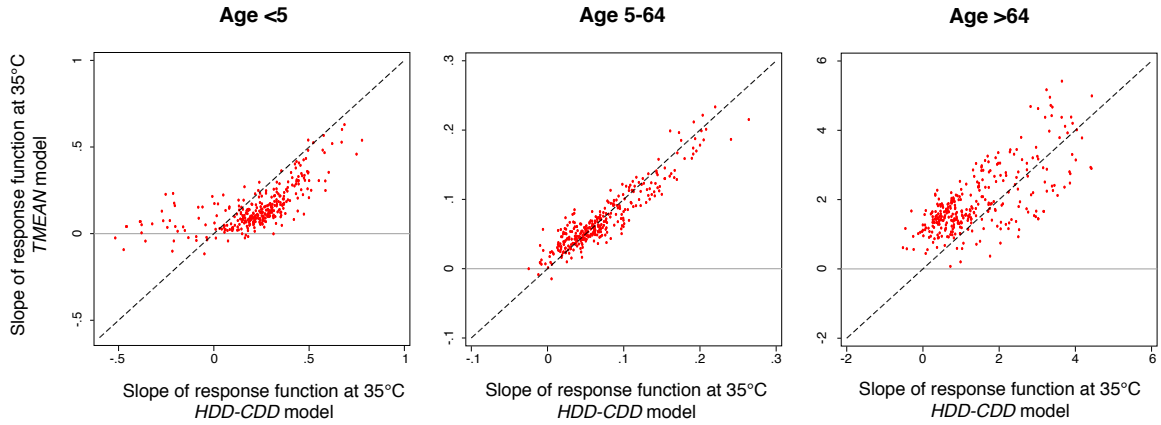

FIGURE 4.D.4

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 4.11 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 4.11 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.

4.D.6 Replication of Burgess et al. (2017) and out-of-sample model validation in India

Throughout our analysis, we use coefficients estimated from Equation 4.11 in the main text, in combination with local-level observations of $TMEAN$ and $\log(GDPpc)$, to generate predicted response functions in all regions of the world, including where mortality data are unavailable (see Section 4.4.3 for details). The accuracy of this extrapolation depends in part on the representativeness of the observed sample; as shown in Figure 4.3 in the main text, our observed sample lacks coverage for the poorest and hottest regions of the global income-climate distribution. To evaluate the performance of our interaction model in this region of the global distribution, we use mortality data from India to conduct an out-of-sample validation exercise. India represents the poorest and hottest country for which we have been able to obtain mortality records, and therefore provides an important check on our extrapolation performance.

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

After constructing an all-age mortality-temperature response function for India alone, we compare this result to the predicted response functions generated from our main analysis. These impact region-level response functions are generated using estimation of Equation 4.11, which relies on mortality, climate, and income data from the sample of countries shown in Table 4.1, but excludes India.

In Figure 4.D.5, we show the result of this replication and out-of-sample validation exercise for India. Figure 4.D.5 compares our predicted responses for all impact regions in India (in grey) to the mortality-temperature response estimated using India's data alone (in red). We generate the all-age average response function for each impact region from our age-specific interaction model by taking a population-weighted average of the responses predicted for each age category, using age-specific population values for the year 2015. Our model performs remarkably well, despite containing no information on Indian mortality rates: for the hotter end of the response function, where much of the low income world resides, our prediction is, if anything, conservative in extrapolating out-of-sample. Moreover, our results are very similar to the findings in Burgess et al. (2017), with an approximately linear increase in deaths for temperatures above 20°C.

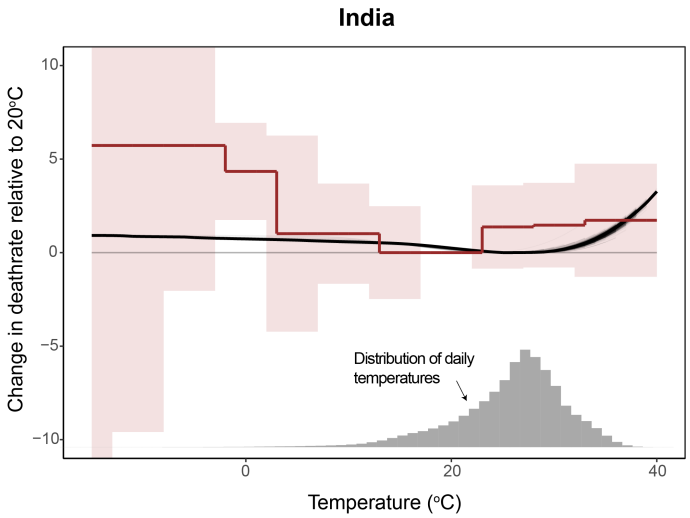

FIGURE 4.D.5

Out-of-sample validation of the mortality-temperature response function in India. Grey lines indicate predicted response functions for each impact region in India, predicted using coefficients from the interaction model in Equation 4.11 and using an estimation sample without Indian data. The solid black line is the unweighted average across all regions, while the red line is the estimated response function using only all-age mortality data from India. The red line is estimated using a nonparametric binned regression, as described in the text, to enable comparison with Burgess et al. (2017). The relative congruence between red and black lines shows that our interaction model generates reasonable predicted response functions in the poorest and hottest regions of the world, the subset of the covariate space for which we have the least representation.

4.E Implementation details for projection of future adaptation and benefits of income growth

In the main analysis, our estimates of the full mortality risk of climate change account for both the benefits and the costs of adaptation, as well as the benefits of income growth. In this appendix, we provide details on our implementation of adaptation and income benefits in future climate change projections. In Appendix 4.E.1 we detail the procedure we use to determine the temporal dynamics of income effects on the mortality-temperature relationship in future years, and in Appendix 4.E.2 we describe the assumptions we impose on the process of adaptation and income benefits over the course of the 21^{st} century.

4.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 4.11 using cross-sectional variation in climate and income in combination with year-to-year variation in daily average temperatures. In generating future projections of climate change impacts (i.e. results in Section 4.5.3), we apply the estimated coefficients from Equation 4.11 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 4.4.4 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:

$$M_{ait} - M_{ai,t-1} = \boldsymbol{\beta}_t \left[\boldsymbol{T}_{it} - \boldsymbol{T}_{i,t-1}\right] + q_a(\boldsymbol{R}_{it}) + \varepsilon_{it} \tag{4.E.1}$$

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 4.10). 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 4.11, 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.

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

As discussed in Section 4.4.3, 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 4.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 polynomi-

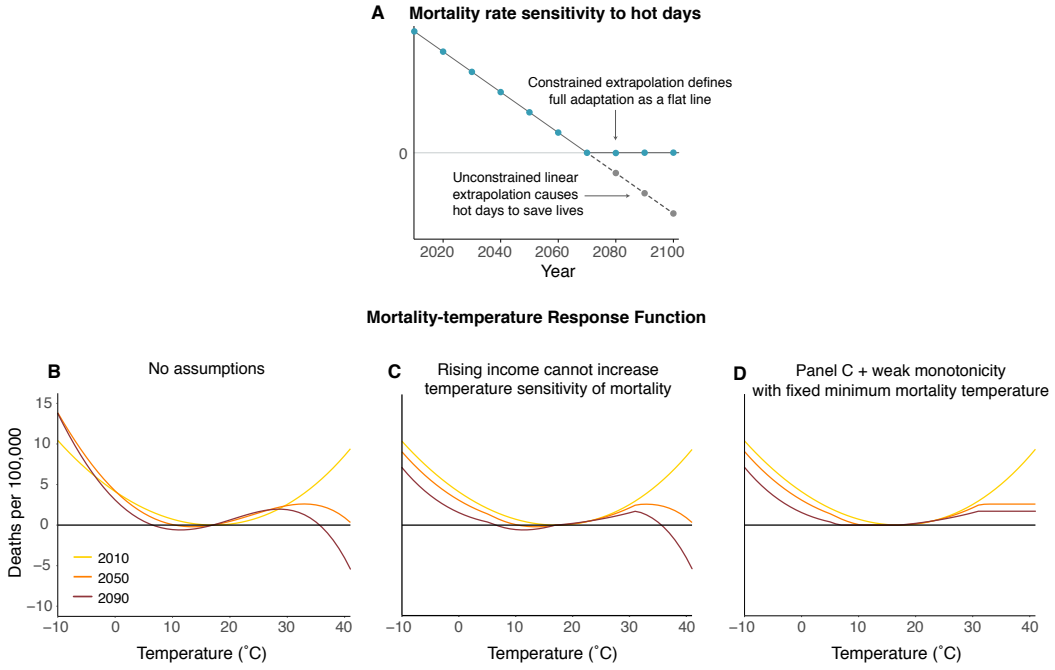

FIGURE 4.E.1

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).

als) and across diverse locations globally (e.g., Gasparrini et al., 2015; Burgess et al., 2017; Deschênes and Greenstone, 2011). As shown in Section 4.5, we also recover U-shaped relationships between mortality rates and daily temperatures across our multi-country sample. In our projections of future mortality responses to daily temperature, we ensure consistency with this literature and with our own estimates from historical data by imposing the constraint that the response function must remain weakly monotonic around an empirically estimated minimum mortality temperature. That is, we assume that temperatures farther from the minimum mortality temperature (either colder or hotter) must be at least as harmful as temperatures closer to the minimum mortality temperature.

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

the temperature at which mortality is minimized,[94] it is important to note that we allow these minimum mortality temperatures (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 4.11. 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 4.11 parameterizes the flattening of the U-shaped response function such that, with enough warming or sufficiently high income, the mortality-response function could become an inverted-U-shape. This is guaranteed to occur mechanically at some future date, as a result of extrapolating response functions out of the support of historically observed data. To avoid this unrealistic behavior, we impose weak monotonicity. An example of this assumption in practice is given in panel E of Figure 4.E.1.[95]

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 4.11 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.[96]

Assumption #2: Rising income cannot increase the temperature sensitivity of mortality. We assume that because increased income per capita strictly expands the choice set of individuals considering whether to make adaptive investments, future increases in income cannot raise the impacts of temperature on mortality rates. While we place no restrictions on the cross-sectional effect of income on the temperature sensitivity as estimated in Equation 4.11, we do not allow any income gains through time to raise the marginal effect of temperature on mortality. Note that this condition will only be binding if the marginal effect of income estimated in Equation 4.11 is positive for some nonempty set of temperatures. Further note that we impose this assumption first, before imposing weak monotonicity, as described under assumption #1. An example of this assumption in practice is given in panel C of Figure 4.E.1.

A visual example of the influence of these constraints can be seen for one example impact region (Chicago, Illinois) in Figure 4.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.

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

[95] See Appendix 4.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.

[96] 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.

4.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 4.5.3), as well as a robustness check and sensitivity analysis regarding the functional form of the mortality-temperature relationship, different assumptions about the behaviour of the relationship outside of the historical sample values, and assumptions regarding the rate of adaptation.

4.F.1 Additional climate change projection results

Alternative measures of climate change impacts In Figure 4.8 of the main text, we show a map of impact region-level mean estimates of the full mortality risk of climate change, accounting both for adaptation and income benefits as well as adaptation costs. However, in Section 4.5.3 we also define three other measures of expected climate change impacts: (i) mortality effects of climate change with neither adaptation nor income growth; (ii) mortality effects of climate change with benefits of income growth; (iii) mortality effects of climate change with benefits of income growth and adaptation. Panels A-C in Figure 4.F.1 below show projected impacts for each of these three alternative measures; for comparison, panel D repeats the full mortality risk of climate change map from the main text.

Climate change projections by age group In the main text, Figure 4.9 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 4.F.2, we show each of these age-specific projections for SSP3 and RCP8.5 (for reference, Table 4.1 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 4.6 and Appendix 4.G).

Climate change projections by socioeconomic scenario Throughout Section 4.5.3 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 mortality effects of climate change under two alternative scenarios (SSP2 and SSP4, alongside SSP3). In each column, we show results for two separate modeling groups that produce projections for each SSP (IIASA and OECD, as discussed in Appendices 4.B.3.2 and 4.B.3.3).

Gains from mitigation on climate change impacts spatially and in aggregate. Figure 4.8, displays climate change impacts spatially under the socioeconomic scenario SSP3 for the entire globe. Figure 4.F.4 shows a comparison between impacts under RCP8.5 and RCP4.5, showing the gains from mitigation. As expected, reducing emissions to the level of RCP4.5 would have substantial benefits in terms of reduced impacts. However, the spatial pattern of impacts remains, with clearly unequal distribution of impacts between places that

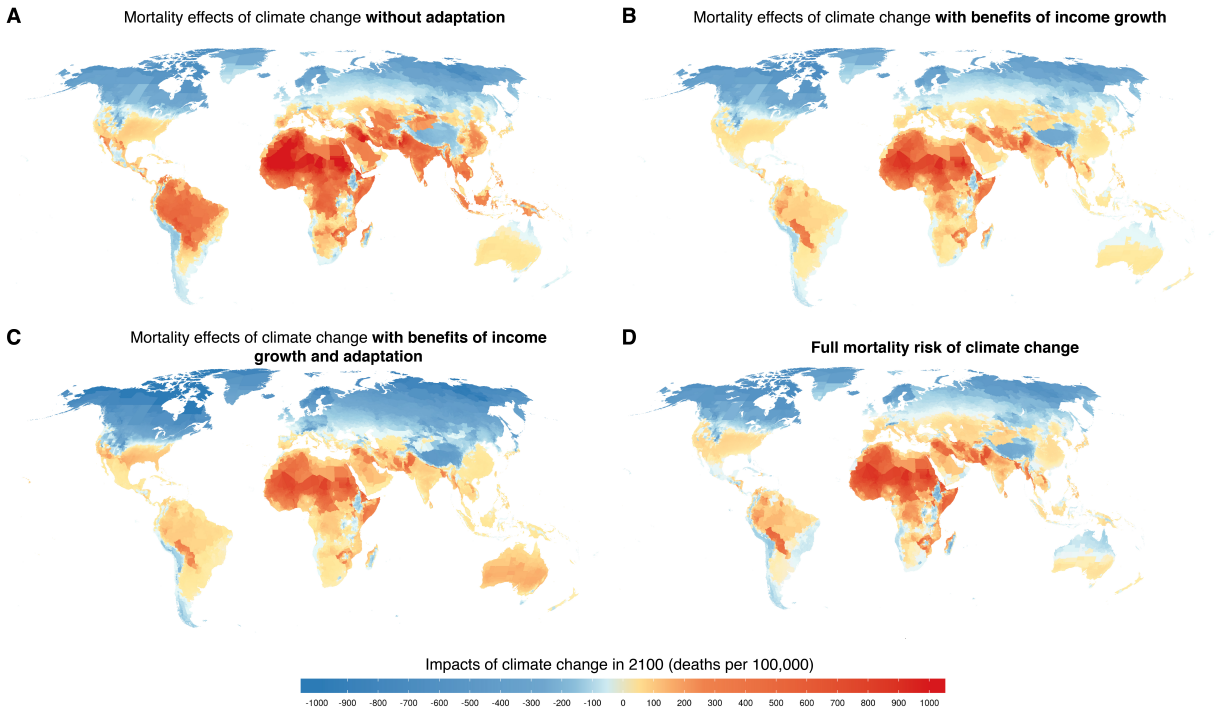

FIGURE 4.F.1

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

are relatively poorer today versus places that are relatively richer. Figure 4.F.5 replicates the aggregate timeseries (Figure 4.9 in the main text) for RCP4.5 under the SSP3 scenario. Clearly, the gains from reducing emissions are evident here, as is implied by Figure 4.F.4.

Correlation of uncertainty in climate change projections with present day income As shown in Figure 4.10 in the main text, there is a strong correlation between present day income and the composition of future damages between deaths and adaptation costs. In general, low income locations tend to suffer relatively large increases in mortality rates by end of century, while high income locations incur relatively large adaptation costs. In Figure 4.F.6 below, we examine whether low income locations also face higher *uncertainty* in projected impacts of climate change. We do so by demeaning each Monte Carlo simulation by an impact region-specific mean, and showing the spread in demeaned impacts for each decile of today's income distribution. As in Figure 4.10, panel A in Figure 4.F.6 shows projected deaths, panel B shows adaptation costs, and panel C shows the full mortality risk of climate change (the sum of deaths and adaptation costs measured in death equivalents); each panel shows impacts plotted against deciles of today's income distribution. There is some evidence that lower income regions face higher uncertainty in the magnitude of their

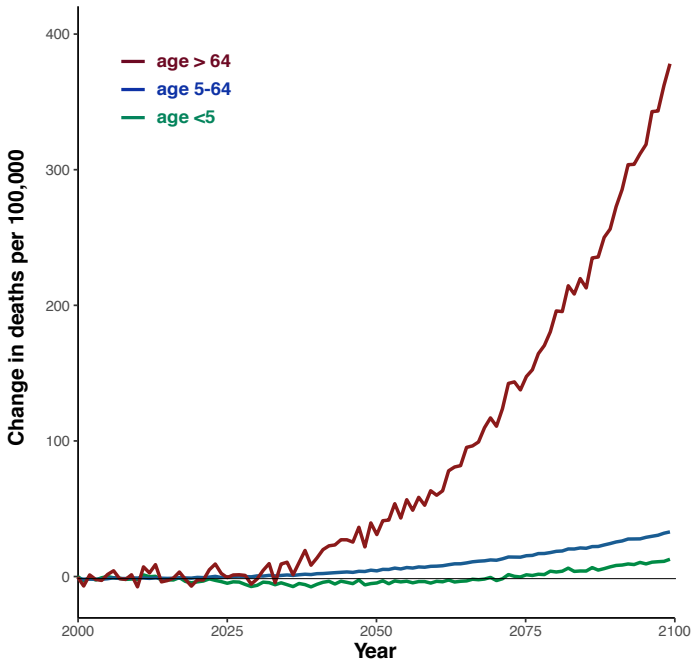

FIGURE 4.F.2

Heterogeneity in climate change impacts on mortality by age group. All lines show predicted mortality effects of climate change across all age categories and are represented by a mean estimate across a set of Monte Carlo simulations accounting for both climate model and statistical uncertainty. Each line represents one of the three age groups used in the analysis: <5, 5-64, and >64. Results are shown for the combination of SSP3 and RCP8.5 with a fourth-order polynomial functional form of temperature. Figure 4.9 in the main text represents the sum across these age-specific projections, where death rates are population weighted by age-specific population values.

projected full mortality risk of climate change, arising from higher uncertainty in impacts on death rates.

The impact of climate change in 2100 under RCP4.5 compared to contemporary leading causes of death. Figure 4.F.7 presents the same results as Figure 4.13 in the main text, but for RCP4.5. As can be seen, despite the overall decrease in the average impact under SSP3 and RCP4.5 when compared to RCP8.5, much of the inequality in both the impacts and the adaptation costs that was evident in Figure 4.13 remains.

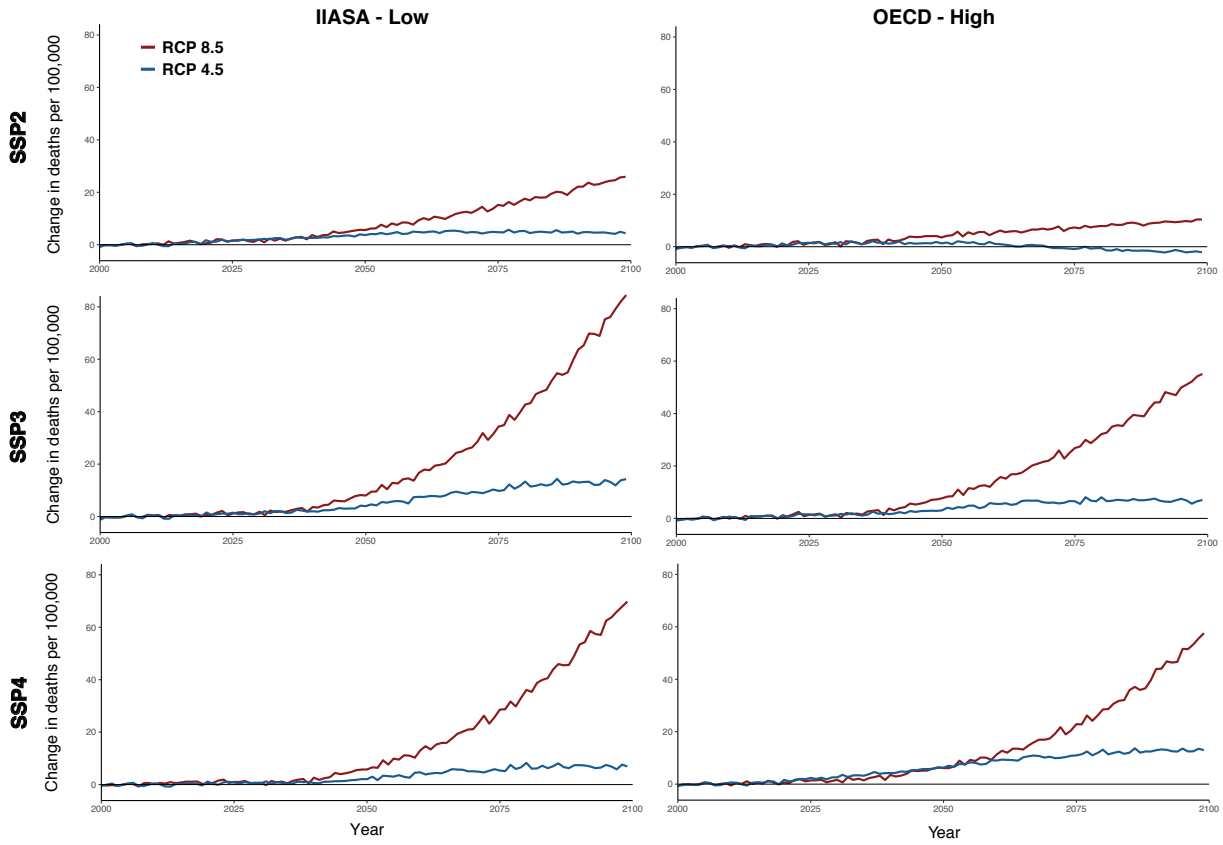

FIGURE 4.F.3

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

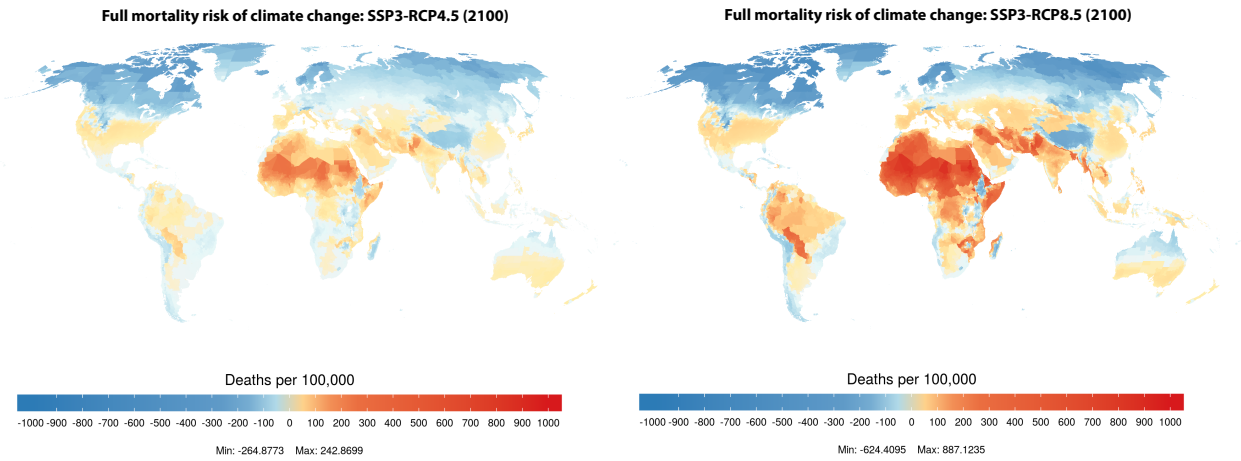

FIGURE 4.F.4

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

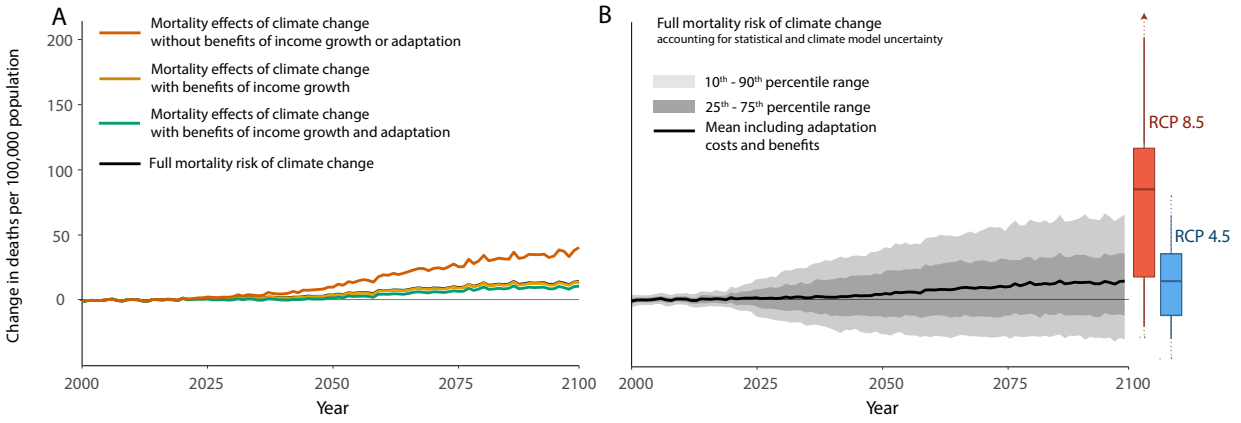

FIGURE 4.F.5

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

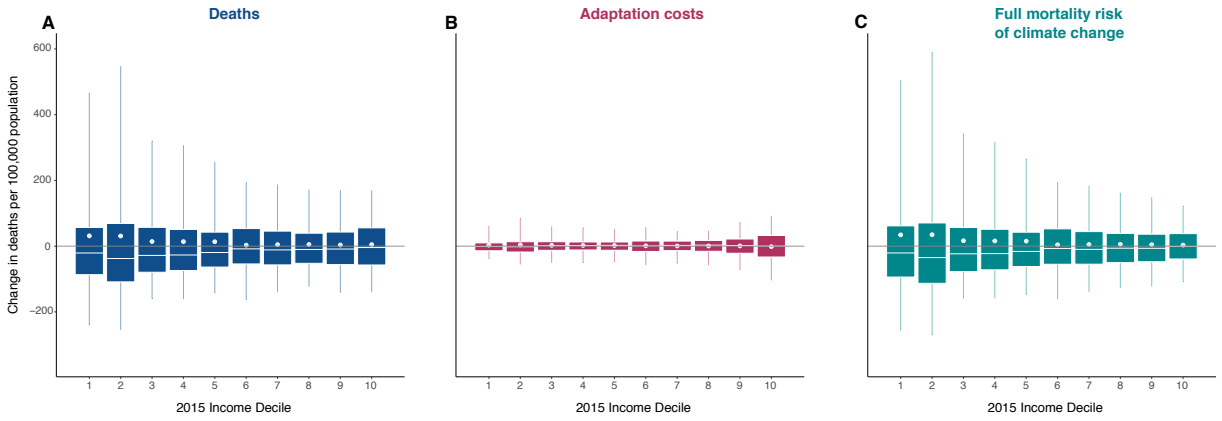

FIGURE 4.F.6

Uncertainty in climate change impacts and adaptation costs across present-day income groups. All three panels show impacts by present-day income decile for impact region estimates that have been demeaned across the full set of Monte Carlo simulations in order to demonstrate how uncertainty in future impacts may be correlated with contemporary incomes. This figure is analogous to Figure 4.10 in the main text, which shows distributions of impact region mean estimates. Panel A shows the range of annual mortality rates due to climate change in 2100, accounting for the benefits of adaptation, against deciles of 2015 income. Panel B shows the range of annual adaptation costs incurred due to climate change in 2100, measured in death equivalents. Panel C shows the range of full mortality risk due to climate change estimates, which are the sum of deaths and adaptation costs measured in death equivalents. Income deciles are calculated as in Figure 4.10. All box plots show moments of the distribution of impact region-specific demeaned impacts within an income decile. Solid vertical lines in each box plot extend to the 5^{th} and 95^{th} percentiles of this distribution, boxes indicate the interquartile range, white horizontal lines indicate the median, and white circles indicate the mean. All values shown refer to the RCP8.5 emissions scenario and the SSP3 socioeconomic scenario.

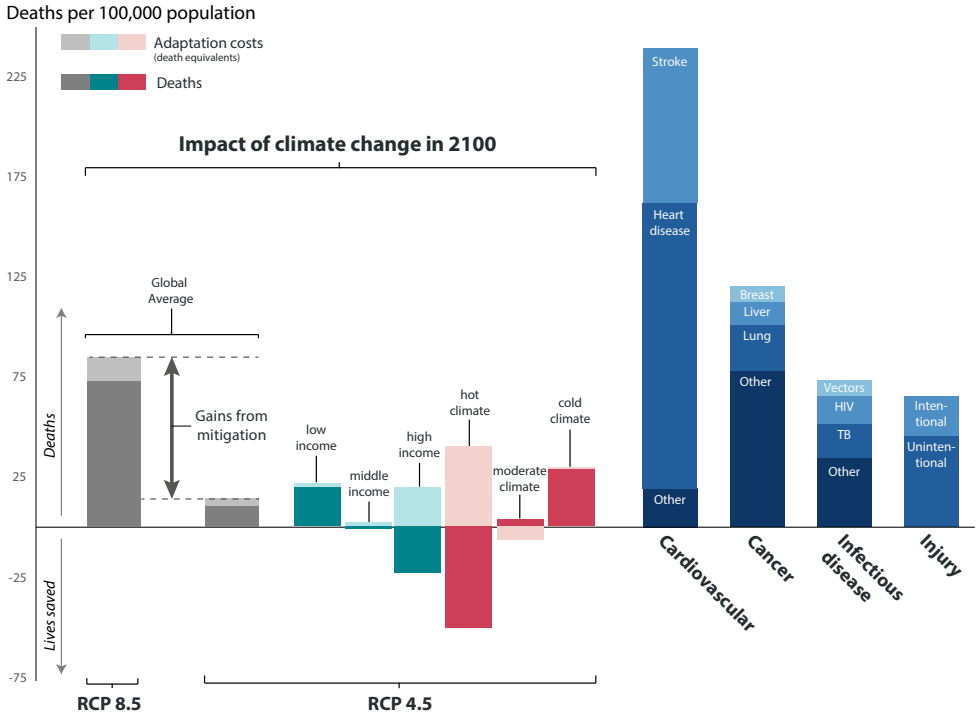

FIGURE 4.F.7

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).

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

As discussed in Section 4.4, we experiment with four distinct nonlinear transformations of daily temperature captured by $\boldsymbol{T}_{it}$ in Equations 4.10 and 4.11 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 4.11 (see Section 4.4.2). However, the binned specification, in which $\boldsymbol{T}_{it}$ contains binned daily temperatures with a fixed set of 5°C bins, is the most flexible functional form. In Figure 4.D.1, we show that the binned and fourth order polynomial functional forms recover remarkably similar mortality-temperature response functions across our pooled multi-country sample. Below in Figure 4.F.8, we show that this similarity carries through to generate very similar climate change impact projections across the binned and polynomial functional forms. Both projections are constructed using estimation of the interaction model in Equation 4.11 in combination with high-resolution covariates $TMEAN$ and $\log(GDPpc)$ to generate impact region-specific response functions (see Section 4.4.3 for details).

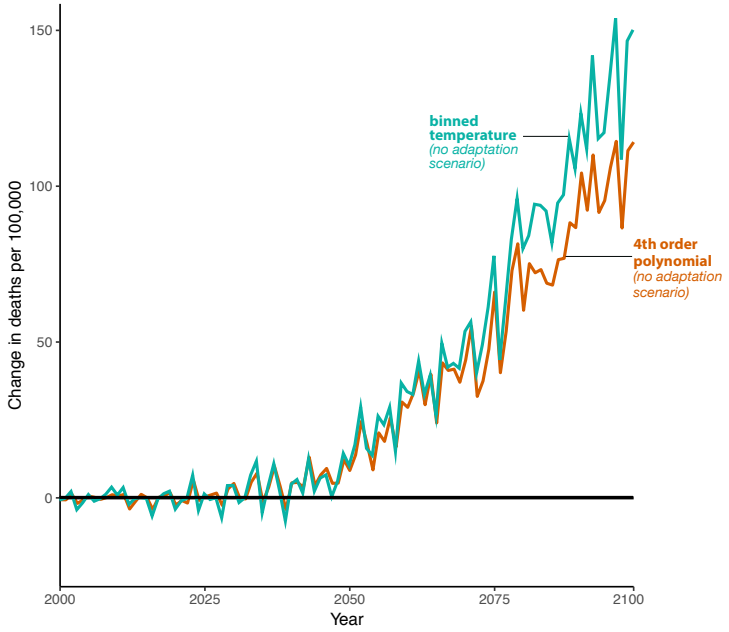

FIGURE 4.F.8

Robustness of impact projections to alternate functional forms of temperature. Each line represents the time series of changes to the mortality rate due to climate change under the socioeconomic scenario SSP3 and the emissions scenario RCP 8.5. Results shown are for a single climate model (CCSM4). Lines shown refer to estimates of mortality effects of climate change without adaptation or benefits of income growth, in which response functions do not evolve over time. In orange is the projected impact of climate change estimated using a fourth-order polynomial functional form of temperature in estimation of the regression model in Equation 4.11. 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.

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

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

Figure 4.F.9 probes the sensitivity of the projections of mortality risk changes up to 2100 to alternative assumptions about the relationship between mortality and temperature at temperatures that are not observed in available data sets. Specifically, for all temperatures above the maximum and below the minimum daily temperatures within our dataset, we alter the slope of the impact region-specific response functions in two ways. First for "constant out-of-sample extrapolation'', we set the marginal effect of temperature fluctuations to equal the value at the maximum if above the maximum temperature, and vice versa for temperatures below the minimum (Figure 4.F.9B). This implies that the response function is flat for all temperatures outside the observed range. For "linear out-of-sample" extrapolation, we set the marginal effect to be linearly increasing in the out-of-sample regions with a slope equal to the slope between the response function evaluated at the maximum (minimum) and the maximum minus 0.1C (plus 0.1C) (Figure 4.F.9C). It is apparent that neither of these alternatives have a meaningful effect on the overall projected impacts; looking at projections from a single GCM, the projected impact of climate change on mortality rates, including the benefits of income growth and adaptation, is 13.6 per 100,000 in 2100 under RCP8.5 in the paper's main specification (Panel A) and 12.6 per 100,000 and 13.4 per 100,000 in Panels B and C.

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

In our main results, we use the estimated coefficients from Equation 4.11 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 4.4, the estimation of Equation 4.11 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 4.E.1, we apply the estimated coefficients from Equation 4.11 over time when computing future climate change impacts; in doing so, we must make an assumption regarding the rate at which mortality sensitivity to temperature declines with changing covariates. As discussed previously, our main specification relies on a 13-year Bartlett kernel for $\log(GDPpc)$ and a 30-year Bartlett kernel for $TMEAN$.

Here, we conduct a sensitivity analysis where the speed at which the mortality-temperature response function changes with time-varying covariates is deterministically reduced by half. This exercise is used to understand how climate change impact projections change if the evolution of the response function towards zero (see Figure 4.E.1) occurs more slowly.

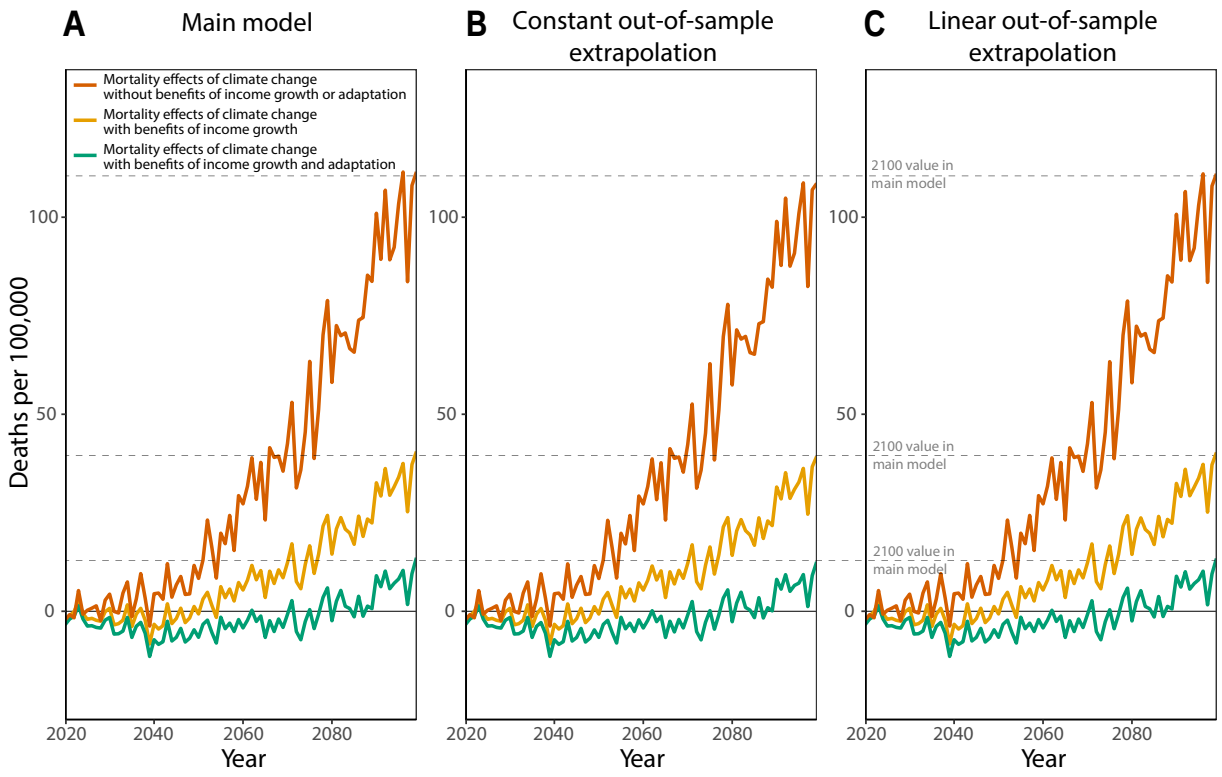

FIGURE 4.F.9

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

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

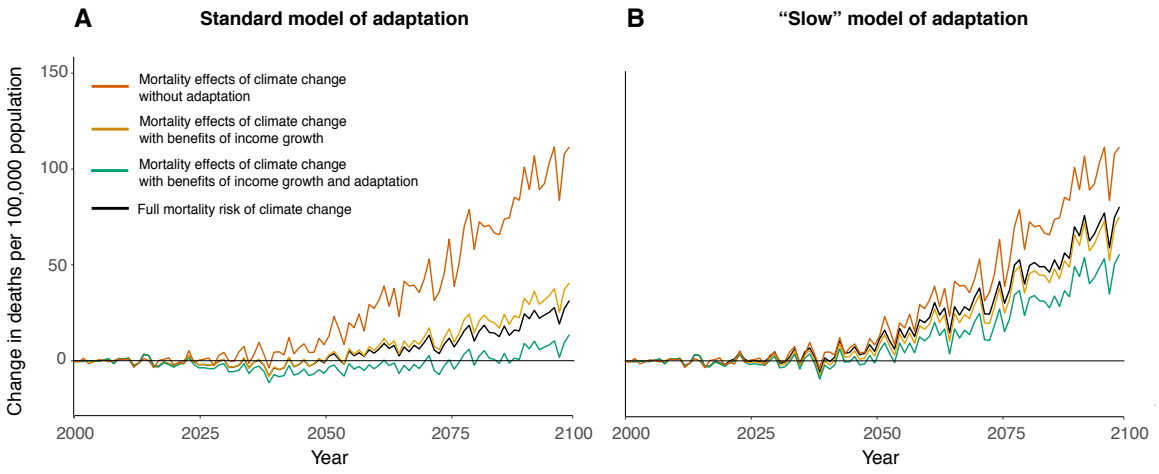

FIGURE 4.F.10

Impacts of climate change on mortality are qualitatively similar with a model of slower adaptation rates. Time series of projected mortality costs of climate change (black line), as compared to partial estimates from incomplete accounting of the costs and benefits of adaptation (other colors). All lines show predicted mortality impacts of climate change across all age categories under the RCP8.5 emissions scenario, for the socioeconomic scenario SSP3, and using a single climate model (CCSM4). Panel A shows results for our standard model of adaptation, as described in Section 4.4.3. Panel B shows results for an alternative model of adaptation in which the rate of adaptation to both income growth and to warming climate is cut in half.

4.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",[97] in combination with our empirically-derived damage functions, to construct mortality partial SCC estimates. We detail this implementation below.

4.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 4.3.2) 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 4.13 reveal that in the latter half of the 21^{st} century, full mortality damages are larger for a given level of warming if warming occurs later in time and damage functions become more convex with time at the end of the 21^{st} century. The finding that mortality costs rise over time is the net result of countervailing forces. On the one hand, later years are projected to have larger and older[98] 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

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

[98] 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.

values D_{irmt} from 2085-2100 and estimate a quadratic model similar to Equation 4.13, but interacting each term linearly with year t (we use 2085-2100 because the evolution of damages over time becomes roughly linear conditional on ΔGMST by this period). The temporal trend over the entire 21^{st} century is convex, implying that our linearization is, if anything, conservative. The specific interaction model we estimate is:

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

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.

4.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:

$$\frac{dT_i}{dt} = \frac{q_i F - T_i}{d_i}, i \in \{0, 1\}, \tag{4.G.1}$$

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:

$$\frac{dR_j}{dt} = a_j E - \frac{R_j}{\alpha_j \tau_j}, j \in \{0, 1, 2, 3\} \tag{4.G.2}$$

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.[99] As described below in Section 4.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 4.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 4.12; as described below and in Section 4.6, this difference is used to construct mortality partial SCC estimates.

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

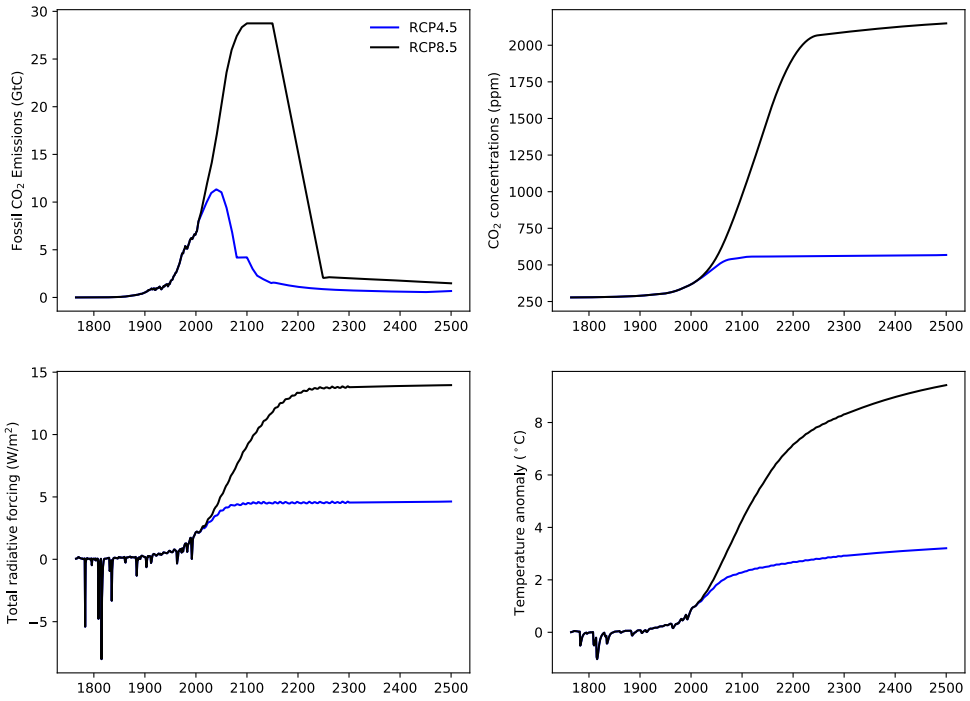

FIGURE 4.G.1

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.

4.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 4.G.1 and Figure 4.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 4.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 4.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 4.3 in the main text.

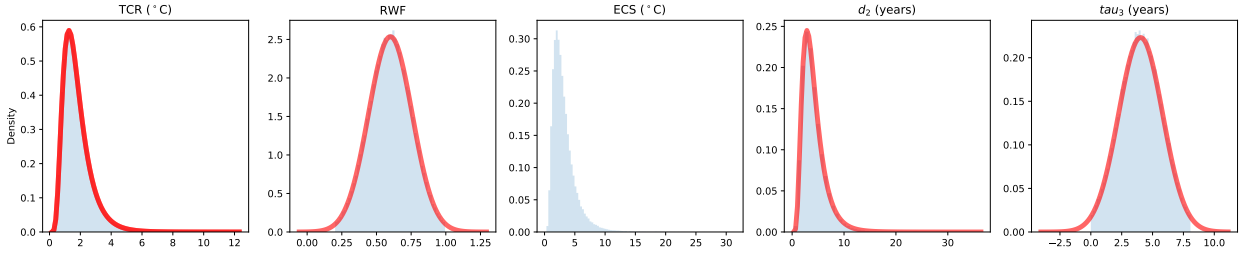

FIGURE 4.G.2

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→ 0; **D** short thermal adjustment time (d_2); **E** time scale of rapid carbon uptake by the ocean mixed layer (τ_3).

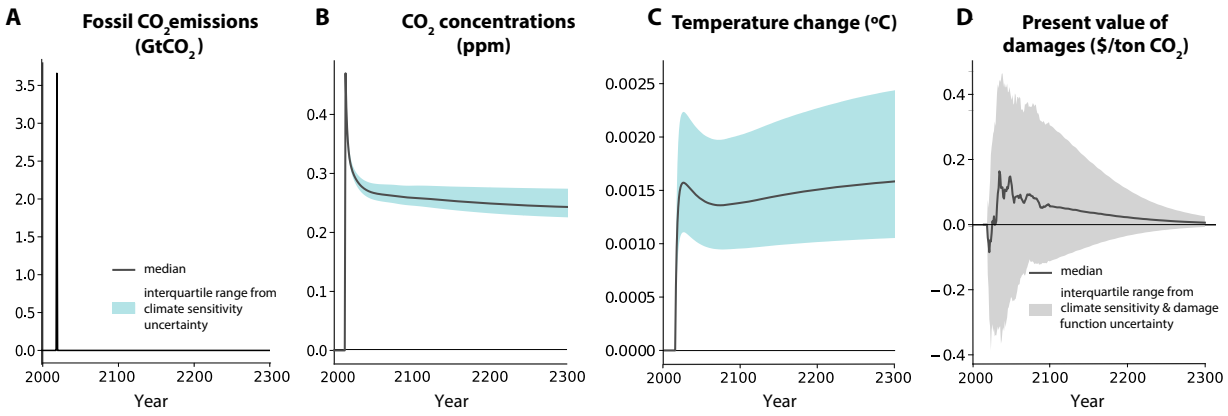

FIGURE 4.G.3

Change in emissions, concentrations, temperature, and damages due to a marginal emissions pulse in 2020. Panel A shows a 1GtC emissions pulse (equivalent to 3.66Gt CO_2) in 2020 for emissions scenario 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 EPA VSL. The shaded area indicates the inter-quartile range due to climate sensitivity and damage function uncertainty, while the solid line is the median estimate.

Graphical SCC calculation for RCP 4.5 Figure 4.G.3 replicates the SCC calculation graphically shown in Figure fig:pulse for RCP 4.5.

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 4.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).

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 4.G.4 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 would be expected based on the construction of the SMME).

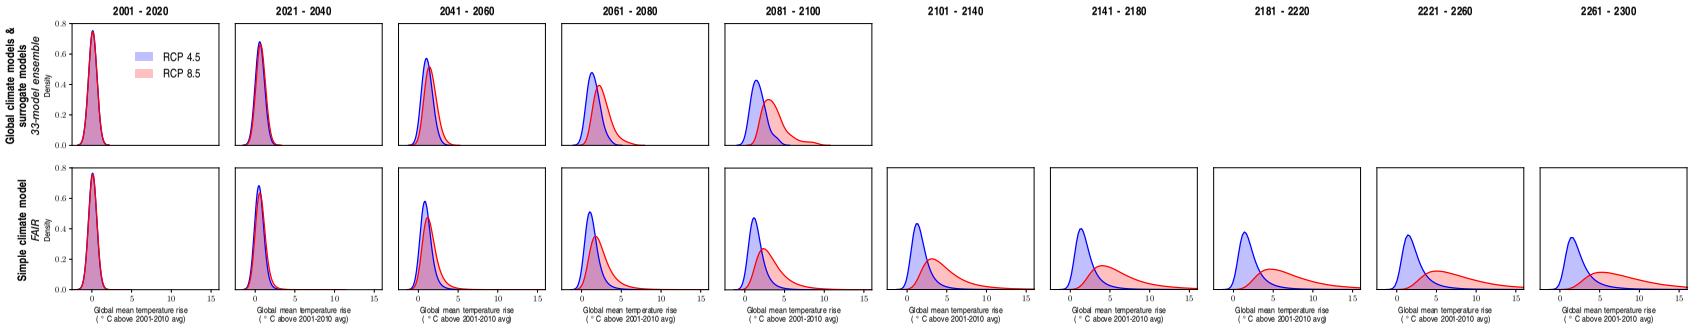

FIGURE 4.G.4

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.

4.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 4.6. These damage functions characterize valued mortality damages as a function of ΔGMST (changes in GMST relative to 2001-2010). Figure 4.G.5 shows these functions in 5-year time steps for each combination of valuation assumptions using the US EPA VSL (see Sections 4.5.3 and 4.6 for discussion of valuation of mortality-related costs of climate change). This figure contains the same information as Figure 4.11 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 4.G.5) .

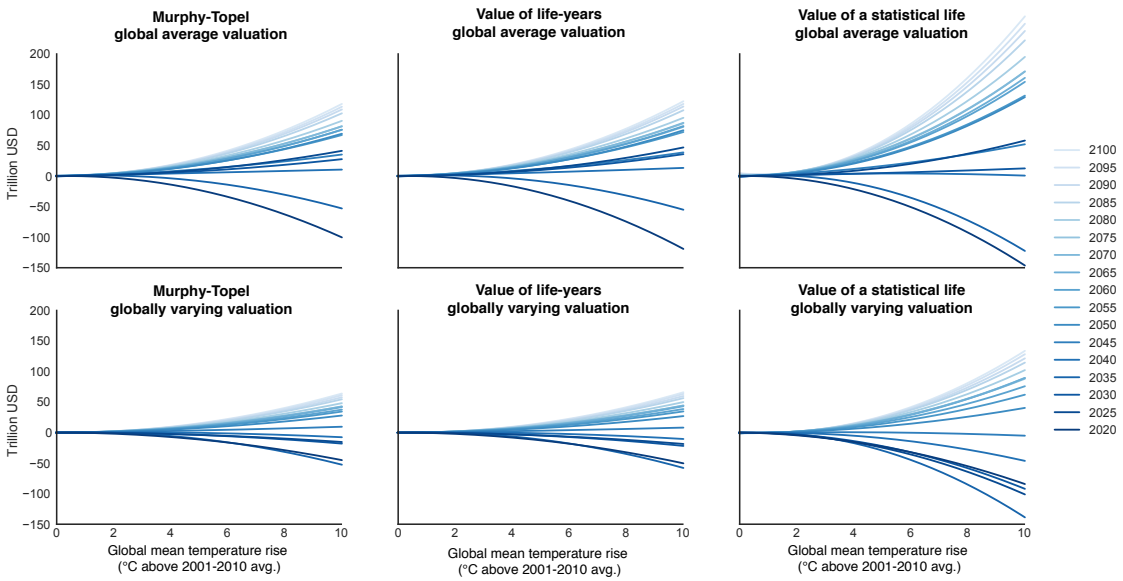

FIGURE 4.G.5

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 4.14 in Section 4.6. This final value is the net present value of the full mortality risks caused of a marginal emission of CO_2. A more robust estimate would make use of Ramsey-like discounting, accounting for the relationship between consumption growth and the discount rate, but we leave this for future study.

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

Mortality partial SCC estimates accounting for both climate sensitivity and damage function uncertainty: Using our Monte Carlo projections of damages, for each year from 2015 to 2100 we pool all Monte Carlo results for the associated 5-year window. We then run quantile regressions to fit quantile-specific damage functions for 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 4.6). 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.

4.H Sensitivity of the mortality partial social cost of carbon

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

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

As described in Section 4.6, panel A of Table 4.3 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 4.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[100] to convert this life-year valuation into a per life-year VSL for each impact region in each year. To calculate life-years lost for a given temperature-induced change in the mortality rate, we use the SSP projected population values, which are provided in 5-year age bins, to compute the implied conditional life expectancy for people in each age bin. We take the population-weighted average of remaining life expectancy across all the 5-year age bins in our broader age categories of <5, 5-64, and >64. This allows us to calculate total expected life-years lost, which we multiply by the impact-region specific VSL per life-year to calculate total damages. Note that this procedure assumes that our estimated climate change driven deaths occur with uniform probability for all people within an age category.

The above methodology values each life-year lost identically. In an alternative set of calculations (see results in Appendix 4.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 4.H.2 and 4.H.3.

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

TABLE 4.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

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

In the main text, mortality partial SCC values are shown using a combination of the US EPA VSL, an income elasticity of one, and valuation methods that value deaths using both an age-varying and an age-invariant value of a statistical life calculation (see Appendix 4.H.1). This appendix shows a range of mortality partial SCC estimates under alternative VSL values, alternative assumptions about the role of income in valuation, and with a life-years adjustment to the VSL that allows for age-specific values of remaining life, as derived by Murphy and Topel (2006).

Table 4.H.2 provides mortality partial SCC estimates across these distinct valuation approaches under the method shown in the main text Table 4.3, in which an income elasticity of one is used to adjust VSLs across the globe and over time.

Table 4.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 (recall Equation 4.4). However, the latter approach might be preferred by policymakers interested in valuing reductions in mortality risk equally for all people globally, regardless of how individuals value their own mortality risk.

Table 4.H.4 shows mortality partial SCC estimates under various socioeconomic projections (SSP3 is used throughout the main text; see Appendix 4.B.3.2 for a discussion of this choice). We note that under SSP4 and a moderate emissions scenario (RCP4.5), the cen-

tral 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 4.F.3), 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.

Finally, Table 4.H.5 shows mortality partial SCC estimates using a 1.5% discount rate, which more accurately reflects recent global capital markets than the discount rates shown in the main text (the average 10-year Treasury Inflation-Indexed Security value from 2003 to present is just 1.01% (Board of Governors of the US Federal Reserve System, 2020)).

TABLE 4.H.2

Globally varying valuation: Estimates of a mortality partial Social Cost of Carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions have heterogeneous valuation, based on local income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix 4.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 4.G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally varying valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	17.1	11.2	7.9	2.9	7.9	5.2	3.7	1.3
Climate sensitivity uncertainty	[8.3, 39.3]	[5.9, 24.1]	[4.4, 15.8]	[2.0, 4.3]	[3.9, 18.3]	[2.8, 11.2]	[2.1, 7.4]	[0.9, 2.0]
Damage function uncertainty	[-21.9, 50.8]	[-19.2, 32.1]	[-12.1, 26.6]	[-6.3, 12.0]	[-10.2, 23.7]	[-9.0, 15.0]	[-5.6, 12.4]	[-2.9, 5.6]
Full uncertainty	[-24.7, 53.6]	[-18.9, 36.0]	[-15.2, 26.3]	[-8.5, 11.5]	[-11.5, 25.0]	[-8.8, 16.8]	[-7.1, 12.2]	[-3.9, 5.3]
RCP 8.5	36.6	22.0	14.2	3.7	17.0	10.2	6.6	1.7
Climate sensitivity uncertainty	[18.8, 76.6]	[11.6, 45.2]	[7.7, 28.3]	[2.4, 6.2]	[8.7, 35.7]	[5.4, 21.0]	[3.6, 13.2]	[1.1, 2.9]
Damage function uncertainty	[-8.4, 74.2]	[-8.7, 48.2]	[-6.4, 35.6]	[-7.3, 14.1]	[-3.9, 34.6]	[-4.0, 22.4]	[-3.0, 16.6]	[-3.4, 6.6]
Full uncertainty	[-7.8, 73.0]	[-10.6, 46.8]	[-11.4, 32.9]	[-8.9, 13.0]	[-3.6, 34.0]	[-5.0, 21.8]	[-5.3, 15.3]	[-4.1, 6.1]
Panel B: Value of statistical life								
RCP 4.5	14.9	9.8	6.7	1.7	7.0	4.6	3.1	0.8
Climate sensitivity uncertainty	[2.4, 52.9]	[2.7, 30.4]	[2.5, 18.3]	[1.0, 2.1]	[1.1, 24.6]	[1.2, 14.2]	[1.2, 8.5]	[0.5, 1.0]
Damage function uncertainty	[-12.8, 44.1]	[-11.8, 33.1]	[-11.1, 25.6]	[-6.8, 12.6]	[-6.0, 20.5]	[-5.5, 15.4]	[-5.2, 11.9]	[-3.2, 5.9]
Full uncertainty	[-21.2, 63.5]	[-17.9, 43.5]	[-15.7, 32.1]	[-11.8, 14.7]	[-9.9, 29.6]	[-8.3, 20.3]	[-7.3, 15.0]	[-5.5, 6.9]
RCP 8.5	65.1	36.9	22.1	3.5	30.3	17.2	10.3	1.6
Climate sensitivity uncertainty	[30.0, 147.0]	[17.5, 82.3]	[10.8, 48.3]	[2.2, 5.6]	[14.0, 68.5]	[8.1, 38.3]	[5.0, 22.5]	[1.0, 2.6]
Damage function uncertainty	[18.4, 98.2]	[8.3, 63.1]	[2.3, 43.7]	[-7.0, 14.5]	[8.6, 45.7]	[3.9, 29.4]	[1.1, 20.3]	[-3.3, 6.7]
Full uncertainty	[3.0, 139.0]	[-2.4, 83.1]	[-5.6, 53.4]	[-9.3, 16.0]	[1.4, 64.7]	[-1.1, 38.7]	[-2.6, 24.9]	[-4.3, 7.5]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	17.5	11.6	8.3	3.1	8.1	5.4	3.9	1.5
Climate sensitivity uncertainty	[8.8, 39.6]	[6.3, 24.5]	[4.7, 16.3]	[2.1, 4.8]	[4.1, 18.4]	[2.9, 11.4]	[2.2, 7.6]	[1.0, 2.2]
Damage function uncertainty	[-16.4, 56.2]	[-16.6, 35.8]	[-12.2, 27.4]	[-6.0, 12.4]	[-7.7, 26.2]	[-7.7, 16.7]	[-5.7, 12.8]	[-2.8, 5.8]
Full uncertainty	[-25.3, 56.6]	[-19.3, 37.9]	[-15.6, 27.7]	[-8.6, 12.2]	[-11.8, 26.4]	[-9.0, 17.7]	[-7.3, 12.9]	[-4.0, 5.7]
RCP 8.5	36.3	22.0	14.3	4.0	16.9	10.3	6.7	1.9
Climate sensitivity uncertainty	[18.8, 75.5]	[11.7, 44.8]	[7.9, 28.3]	[2.6, 6.6]	[8.7, 35.2]	[5.5, 20.9]	[3.7, 13.2]	[1.2, 3.1]
Damage function uncertainty	[-8.2, 67.4]	[-7.5, 46.8]	[-8.3, 33.3]	[-5.5, 14.0]	[-3.8, 31.4]	[-3.5, 21.8]	[-3.9, 15.5]	[-2.6, 6.5]
Full uncertainty	[-8.0, 70.9]	[-11.0, 46.4]	[-11.6, 33.0]	[-8.8, 13.6]	[-3.7, 33.0]	[-5.1, 21.6]	[-5.4, 15.4]	[-4.1, 6.3]

TABLE 4.H.3

Globally uniform valuation: Estimates of a mortality partial Social Cost of Carbon (SCC) under different valuation assumptions. An income elasticity of one is used to scale either the U.S. EPA VSL, or the VSL estimate from (Ashenfelter and Greenstone, 2004). All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from results using the socioeconomic scenario SSP3. All regions are given the global median VSL, after scaling using income. Value of life years estimates (panel A) adjust death valuation by expected life-years lost. Value of statistical life estimates (panel B) use age-invariant death valuation. Murphy-Topel life years adjusted estimates (panel C) add an age-specific adjustment that allows the value of a life-year to vary with age, based on Murphy and Topel (2006) and described in Appendix 4.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 4.G for details).

Valuation	**EPA**				**A & G**			
Discount rate	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	Globally uniform valuation of mortality risk (2019 US Dollars)							
Panel A: Value of life years								
RCP 4.5	37.5	26.4	19.9	9.0	17.5	12.3	9.3	4.2
Climate sensitivity uncertainty	[19.4, 82.2]	[14.5, 54.4]	[11.4, 38.7]	[5.8, 15.1]	[9.0, 38.3]	[6.8, 25.3]	[5.3, 18.0]	[2.7, 7.1]
Damage function uncertainty	[-15.7, 87.9]	[-10.9, 63.1]	[-11.4, 44.2]	[-2.4, 23.1]	[-7.3, 40.9]	[-5.1, 29.4]	[-5.3, 20.6]	[-1.1, 10.7]
Full uncertainty	[-13.3, 101.7]	[-10.2, 68.7]	[-8.4, 50.0]	[-5.0, 21.7]	[-6.2, 47.4]	[-4.8, 32.0]	[-3.9, 23.3]	[-2.3, 10.1]
RCP 8.5	72.3	46.3	32.0	11.5	33.6	21.6	14.9	5.3
Climate sensitivity uncertainty	[37.8, 149.0]	[24.9, 93.4]	[17.7, 62.8]	[7.0, 20.1]	[17.6, 69.4]	[11.6, 43.5]	[8.2, 29.2]	[3.2, 9.4]
Damage function uncertainty	[7.2, 127.1]	[4.3, 86.3]	[-1.9, 59.0]	[-4.8, 24.8]	[3.4, 59.2]	[2.0, 40.2]	[-0.9, 27.5]	[-2.2, 11.5]
Full uncertainty	[4.6, 141.1]	[-0.5, 92.1]	[-2.9, 64.6]	[-4.6, 24.9]	[2.2, 65.7]	[-0.2, 42.9]	[-1.4, 30.1]	[-2.1, 11.6]
Panel B: Value of statistical life								
RCP 4.5	46.2	33.7	25.9	11.9	21.5	15.7	12.1	5.5
Climate sensitivity uncertainty	[15.3, 134.1]	[14.0, 86.6]	[12.3, 60.2]	[7.2, 21.4]	[7.1, 62.4]	[6.5, 40.3]	[5.7, 28.0]	[3.4, 10.0]
Damage function uncertainty	[14.3, 102.0]	[12.9, 75.4]	[3.5, 56.9]	[-0.7, 26.9]	[6.6, 47.5]	[6.0, 35.1]	[1.6, 26.5]	[-0.3, 12.5]
Full uncertainty	[2.8, 148.2]	[-1.8, 98.6]	[-4.1, 71.0]	[-4.2, 30.2]	[1.3, 69.0]	[-0.8, 45.9]	[-1.9, 33.1]	[-2.0, 14.1]
RCP 8.5	143.9	87.5	57.5	17.6	67.0	40.8	26.8	8.2
Climate sensitivity uncertainty	[68.8, 317.6]	[43.1, 189.7]	[29.2, 121.7]	[10.1, 32.9]	[32.0, 147.9]	[20.1, 88.3]	[13.6, 56.7]	[4.7, 15.3]
Damage function uncertainty	[59.5, 197.8]	[38.1, 130.2]	[23.8, 94.9]	[3.1, 33.2]	[27.7, 92.1]	[17.7, 60.6]	[11.1, 44.2]	[1.5, 15.4]
Full uncertainty	[39.0, 287.0]	[21.8, 176.9]	[11.9, 117.4]	[-2.0, 37.9]	[18.2, 133.7]	[10.1, 82.4]	[5.5, 54.7]	[-1.0, 17.6]
Panel C: Murphy-Topel life years adjusted								
RCP 4.5	35.8	25.3	19.0	8.6	16.7	11.8	8.8	4.0
Climate sensitivity uncertainty	[18.4, 79.1]	[13.8, 52.1]	[10.9, 37.0]	[5.5, 14.4]	[8.6, 36.8]	[6.4, 24.3]	[5.1, 17.2]	[2.6, 6.7]
Damage function uncertainty	[-15.1, 90.6]	[-8.0, 61.1]	[-4.9, 46.8]	[-4.1, 21.5]	[-7.0, 42.2]	[-3.7, 28.4]	[-2.3, 21.8]	[-1.9, 10.0]
Full uncertainty	[-14.2, 99.9]	[-10.8, 67.2]	[-9.0, 48.8]	[-5.7, 21.3]	[-6.6, 46.5]	[-5.0, 31.3]	[-4.2, 22.7]	[-2.7, 9.9]
RCP 8.5	70.1	44.6	30.7	10.9	32.6	20.8	14.3	5.1
Climate sensitivity uncertainty	[36.6, 144.7]	[24.0, 90.0]	[17.0, 60.2]	[6.6, 19.1]	[17.1, 67.4]	[11.2, 41.9]	[7.9, 28.0]	[3.1, 8.9]
Damage function uncertainty	[7.0, 123.2]	[0.0, 79.4]	[-0.8, 59.5]	[-4.4, 22.5]	[3.3, 57.4]	[0.0, 37.0]	[-0.4, 27.7]	[-2.0, 10.5]
Full uncertainty	[3.7, 134.5]	[-0.8, 87.9]	[-2.7, 61.7]	[-5.2, 24.0]	[1.7, 62.7]	[-0.4, 40.9]	[-1.3, 28.7]	[-2.4, 11.2]

TABLE 4.H.4

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

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel B: Age-adjusted globally VSL (2019 USD)				
RCP 4.5				
SSP2	25.7	15.8	10.4	2.9
SSP3	17.1	11.2	7.9	2.9
SSP4	-14.5	-10.0	-7.5	-3.7
RCP 8.5				
SSP2	33.3	18.7	11.0	1.2
SSP3	36.6	22.0	14.2	3.7
SSP4	22.5	13.0	7.9	1.2
Panel A: Globally varying VSL (2019 USD)				
RCP 4.5				
SSP2	2.0	0.3	-0.9	-3.3
SSP3	14.9	9.8	6.7	1.7
SSP4	-64.3	-46.6	-36.1	-18.5
RCP 8.5				
SSP2	43.9	22.0	10.7	-2.5
SSP3	65.1	36.9	22.1	3.5
SSP4	23.1	8.6	1.2	-6.4

TABLE 4.H.5

Estimates of a partial Social Cost of Carbon (SCC) for excess mortality risk incorporating the costs and benefits of adaptation, 1.5% discount rate. In both panels, an income elasticity of one is used to scale the U.S. EPA VSL value. All regions thus have heterogeneous valuation, based on local income. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix 4.H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Point estimates rely on the median values of the four key input parameters into the climate model FAIR and a conditional mean estimate of the damage function. The uncertainty ranges are interquartile ranges [IQRs] including both climate sensitivity uncertainty and damage function uncertainty (see Appendix 4.G for details).

	Annual discount rate $\delta = 1.5\%$
Panel A: Age-adjusted globally varying VSL (2019 US Dollars)	
Moderate emissions scenario (RCP 4.5)	28.5
Full uncertainty IQR	[-35.6, 88.5]
High emissions scenario (RCP 8.5)	66.4
Full uncertainty IQR	[-2.8, 126.5]
Panel B: Globally varying VSL (2019 US Dollars)	
Moderate emissions scenario (RCP 4.5)	24.6
Full uncertainty IQR	[-25.5, 102.9]
High emissions scenario (RCP 8.5)	123.9
Full uncertainty IQR	[13.7, 253.6]

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

As discussed in Section 4.6, 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 4.H.6 shows that this alternative approach to post-2100 damage estimation causes our central estimate of the SCC to fall by 21%.

TABLE 4.H.6

The influence of damage function extrapolation in years after 2100 on estimates of a mortality partial Social Cost of Carbon (SCC). In this table, an income elasticity of one is used to scale the U.S. EPA VSL value, and all SCC values are for the year 2020, 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 4.3 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 4.6. 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

4.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 4.6 for details). In Table 4.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 4.H.7

Estimates of a mortality partial Social Cost of Carbon (SCC) using a cubic polynomial damage function In this table, an income elasticity of one is used to scale the U.S. EPA VSL value. All SCC values are for the year 2020, measured in PPP-adjusted 2019 USD, and are calculated from damage functions estimated from projected results under the socioeconomic scenario SSP3. Damage functions are estimated as a cubic polynomial, instead of a quadratic (as in the main text). In panel A, SCC estimates use an age adjustment that values deaths by the expected number of life-years lost, using an equal value per life-year (see Appendix 4.H.1 for details). In panel B, SCC calculations use value of a statistical life estimates that do not vary with age. Estimates rely on the median values of the four key input parameters into the simple climate model FAIR and a conditional mean estimate of the damage function.

	Annual discount rate			
	$\delta = 2\%$	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
Panel A: Age-adjusted globally varying VSL (2019 USD)				
RCP 4.5	9.4	6.5	4.9	2.4
RCP 8.5	44.5	25.7	16.1	4.0
Panel B: Globally varying VSL (2019 USD)				
RCP 4.5	18.7	12.5	9.1	3.8
RCP 8.5	68.4	37.6	21.9	2.8

Appendix references

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

Auffhammer, Maximilian, 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.

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

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

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

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

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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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* .

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

Chapter 5: Estimating a Social Cost of Carbon for Global Energy Consumption

Chapter Summary

The global marginal damage caused by emitting a single ton of carbon dioxide (CO_2), or its equivalent, is key to climate policy,[1–3] but our current understanding of its value is based on spatially-coarse theoretical-numerical models[4–6] that are not tightly linked to data.[3,71] We develop the first architecture that integrates best-available data, econometrics, and climate science to estimate climate damages worldwide at the local level, as well as aggregated global marginal damages. Here we apply this architecture to construct the first global empirical estimates of the impact of climate change on total non-transport end-use energy consumption, one of the most uncertain impacts in current models.[8] At end-of-century, we project annual global electricity consumption to rise roughly 4 EJ (1100 TWh, 6% of current global consumption) for each 1°C increase in global mean temperature, reflecting increased cooling demand, while direct consumption of other fuels declines 10.1 EJ (6% of current global consumption) per 1°C, reflecting reduced heating. Together, these estimates indicate that emission of 1 ton of CO_2 today produces global net savings in future aggregate energy consumption of about $1 in net present value (3% discount rate). This finding is largely driven by a sharply nonlinear relationship between income and temperature-induced energy consumption, which indicates that for most of the 21^{st} century, much of the world is expected to remain too poor to increase energy consumption in response to warmer temperatures. By end-of-century, emerging economies in the tropics (e.g. India) are projected to increase electricity consumption dramatically, but these rising costs are offset by heating reductions in the wealthy economies of North America and Europe.

[1]This material first appeared as a working paper of the same title, with authors Ashwin Rode, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert Kopp, Kelly McCusker, Ishan Nath, James Rising, Justin Simcock, Jiacan Yuan. This project is an output of the Climate Impact Lab consortium that gratefully acknowledges funding from the Carnegie Corporation, Energy Policy Institute of Chicago (EPIC), International Growth Centre, National Science Foundation (#SES1463644), Sloan Foundation, and Tata Center for Development. Tamma Carleton acknowledges funding from the US Environmental Protection Agency Science To Achieve Results Fellowship (#FP91780401). James Rising acknowledges funding from the H2020-MSCA-RISE project GEMCLIME-2020 GA No. 681228. We thank Laura Alcocer, Thomas Bearpark, Trinetta Chong, Greg Dobbels, Radhika Goyal, Simon Greenhill, Dylan Hogan, Azhar Hussain, Theodor Kulczycki, Maya Norman, Sebastien Phan, Yuqi Song, Jingyuan Wang, and Jong-kai Yang for invaluable research assistance during all stages of this project, and we thank Jack Chang, Megan Landín, and Terin Mayer for excellent project management. We acknowledge the World Climate Research Programme's Working Group on Coupled Modeling, which is responsible for CMIP, and we thank the climate modeling groups for producing and making available their model output. For CMIP, the U.S. Department of Energy's Program for Climate Model Diagnosis and Intercomparison provides coordinating support and led development of software infrastructure in partnership with the Global Organization for Earth System Science Portals. We thank seminar participants at the UC Berkeley Energy Camp, University of Chicago EPIC Lunch Series, LSE Workshop in Environmental Economics, International Energy Workshop, International Workshop on Empirical Methods in Energy Economics, and University of Michigan Sustainability and Development Conference for helpful comments. The authors declare that data and code for replicating the findings of this study are available at GitLab: https://gitlab.com/ClimateImpactLab/Impacts/energy-code-release/.

5.1 Introduction

Quantifying the benefits of greenhouse gas mitigation is a topic of considerable importance to researchers and policy makers alike. The "social cost of carbon" (SCC), defined as the dollar value of climate change damages imposed globally by an additional (i.e. "marginal") ton of CO_2 emissions (or its equivalent), provides the means to determine the global social benefits of any mitigation policy.[3] To date, our understanding of the SCC has been informed by theoretical-numerical integrated assessment models[4,9,10] (IAMs). These pioneering models have produced numerous valuable insights and guided research and policy for decades.[2,11] Yet as research has progressed with advances in data and computing, new challenges and opportunities have emerged.[7,12]

Recent assessments[3,6,12,13] have raised the concerns that current IAMs are not tightly constrained by data, do not utilize best-available Earth System models, do not capture many known linkages between climate change and society, and only resolve damages at the geographic scale of large regions (e.g. continents). We address these concerns by designing a fully modular "bottom-up" architecture to develop "partial" SCC estimates for individual sub-sectors of the global economy (e.g. agriculture, health, labor), using representative data and detailed climate models.[14–22] Each global partial SCC is built up from econometrically derived, probabilistic, local damage estimates for thousands of geographic regions. In ongoing work, we are integrating these partial SCC estimates[23] to compute a total SCC, taking into account inter-sector and spatial dynamics. To our knowledge, no existing IAM transparently assembles a bottom-up SCC based on local econometric-based projections of damage.[3]

Here we develop the first empirically-derived estimates of the net cost of global energy consumption associated with an additional ton of CO_2 emissions, i.e. a *partial SCC for energy consumption*. IAM developers have themselves argued that uncertainty over this number is the most important uncertainty to resolve in the total SCC,[8] in part because some models predict that rising energy demand will be the single largest global cost from warming.[24] Prior econometric studies have measured the effect of local temperatures on local electricity consumption,[14,25–28] although they often omit non-electric energy consumption (e.g. natural gas used for heating) because the data are difficult to obtain. Moreover, these studies generally focus on residential end uses in regions that are wealthy and thus not globally representative (e.g. California). In contrast, an appealing feature of non-econometric studies using process-based models is that they simulate how climate change will affect all aspects of the production, conversion, delivery, and use of energy.[29–32] However, similar to IAMs, their drawback is that they are not generally constrained by plausibly causal econometric estimates of consumption behavior in response to warming. This analysis is the first to recover globally representative measurements of total energy consumption in response to rising temperatures, accounting for economic development and adaptive behavior, and to use these results to compute a partial SCC for energy consumption.

5.1.1 Implementation

Our modular approach to computing partial-SCC values has five steps, each of which can be implemented for each sector of the global economy. Here we apply these steps to compute the energy consumption component of the SCC.

First, we match globally representative, longitudinal data on energy consumption

with $0.25° \times 0.25°$ globally harmonized historical climate data.[33] This represents, to our knowledge, the most comprehensive global dataset compiled on energy consumption and temperature (*Methods*; *Appendix* 5.A). Energy consumption data are derived from International Energy Agency (IEA) data files[34] that describe electricity and direct fuel consumption across residential, commercial, industrial, and agricultural end-uses (excluding transportation) in 146 countries during 1971-2010. To make these data usable for global analysis, we harmonize data across diverse reporting systems and use econometric methods that minimize the influence of errors in record keeping (*Methods*; *Appendix* 5.A.1).

Second, we econometrically estimate the effect of historical temperature distributions on energy consumption using random year-to-year variation,[20] and measure how this *energy-temperature response* differs across energy types (*electricity* and *other fuels*), income levels, and climate zones.[23] This allows us to observe the effects of adaptive behaviors that populations undertake as they become richer[25,35] and/or are exposed to warmer climates[28] (e.g. AC adoption). Our approach accounts for all permanent differences between countries in energy consumption (e.g. due to geography or history) and all common trends in energy consumption (e.g. due to macroeconomic fluctuations, price changes, or technological innovations), thereby identifying a plausibly causal effect[14] of temperature distributions on energy consumption (*Methods*; *Appendix* 5.C).

Third, we project impacts of climate change in 24,378 geographic regions (e.g., US counties, See *Appendix* Figure 5.B.1) through 2099 (the final year high-resolution climate simulations are available) by combining the econometric results above with a probabilistic ensemble of downscaled climate projections (*Appendix* Figure 5.A.3)[36] based on CMIP5 models.[15] When projecting these impacts, we account for how the energy-temperature response will evolve as populations become richer and exposed to warmer climates (*Methods*). Standard socioeconomic scenarios[18] forecast that 95% of the end-of-century population will still remain within the range of historical temperatures and incomes that we currently observe around the world (See *Appendix* 5.C.3, Figure 5.C.3). In isolating the impact of future climate change on energy consumption, we hold constant the current energy supply mix, an assumption that should be relaxed in future work.

Fourth, we pool the empirically-derived damage estimates from the last step and fit *global energy damage functions* by aggregating impacts across locations and indexing them against the global mean surface temperature anomaly (ΔGMST) expressed in each climate model realization.[37] These functions describe the full distribution of global damage conditional on ΔGMST. We estimate damage functions that evolve over time to reflect expected changes in socioeconomics and adaptation[23] (*Methods*; *Appendix* 5.E).

Fifth, we adapt a probabilistic, simple climate-carbon cycle model[22] to project the distribution of annual ΔGMST up to 2300 that results from the release of 1 additional GtC of CO_2 (*Methods*; *Appendix* 5.F). Applying the distribution of impulse-responses of ΔGMST to damage functions from the last step generates a probability distribution for the stream of total global damages that result from the emission of a marginal ton of CO_2 today. This probability distribution accounts for uncertainty in our econometric estimates at all stages of the analysis as well as climatological uncertainty. Finally, the value of the flow of damage is discounted[3] to capture the partial SCC for global energy consumption.

5.2 Results

Empirically, we find that a population's average income per capita is a key determinant of how its end-use energy consumption responds to temperature. Electricity-temperature responses (Figure 5.1A) are "U"-shaped (i.e. increasing with hot and cold temperatures), but only in the seventh decile or higher of the global income distribution (annual per-capita income≥$11,258, 2019 USD PPP), while the other fuels-temperature response is "L-shaped" (i.e. increasing with cold temperatures) in the third decile or higher (annual per-capita income≥$2,849, 2019 USD PPP). Above these thresholds, rising incomes appear to amplify both responses, such that in the top decile, electricity consumption rises 0.017 GJ-per-capita (4.6 kWh-per-capita) on a 35°C day (relative to a 20°C day) and 0.0068 GJ-per-capita (1.9 kWh-per-capita) on a 0°C day, on average, while direct consumption of other fuels increases 0.034 GJ per-capita on a 0°C day. These differing responses likely reflect the use of electricity for cooling and heating, compared to the use of other fuels (e.g. natural gas, oil, and coal) for heating. Prior research has documented similarly "U"-shaped electricity-temperature responses in the top decile[14,26–28] (See *Appendix* 5.J for comparisons), but our data reveal that such responses do not generalize to other income levels nor do they capture the substantial other fuels-temperature response, which dominates on cold days (Figure 5.1B). To our knowledge, these findings represent the first empirical demonstration of how economic development shapes energy-temperature responses on a global, macroeconomic scale.

While income per capita is the dominant driver of the energy-temperature response, long-run climate also plays a smaller role in how populations adapt. For instance, higher AC adoption in hot locations may increase electricity use on hot days.[25,28] We empirically recover how income and long-run climate continuously and jointly shape the energy-temperature response (*Methods*, Equation 5.1; *Appendix* 5.C.3), thereby accounting simultaneously for effects of economic development and climate (Figure 5.1C). We find evidence that populations adapt to their long-run climate in ways that change their energy consumption during hot and cold periods, conditional on their income level. For instance, on a 35°C day, per capita electricity consumption is 0.0029 GJ greater in the hottest climate tercile relative to the coldest. Conversely, on a 0°C day, per capita consumption of electricity and other fuels are respectively 0.0024 GJ and 0.037 GJ greater in the coldest tercile relative to the hottest. These results are consistent with populations adopting more heating or cooling technologies when their climate is cooler or hotter, respectively.

Combining the measured relationships depicted in Figure 5.1C with projections for how incomes and climate will change over the next century, we project how the structure of all 24,378 local energy-temperature responses will evolve (See Figure 5.2A inset, *Methods*, and *Appendix* 5.C.3). This spatial granularity contrasts with existing IAMs used to develop SCC estimates, which partitioned the world into at most 16 units[10] (*Appendix* 5.B). Applying the ensemble of downscaled climate models and surrogates (*Methods*; *Appendix* 5.A.2.2) to our evolving projections of local energy-temperature responses, we isolate the *additional* energy consumption in each region caused by changes in the temperature distribution, over and above any changes to consumption that would occur without climate change, such as those increases associated with economic development (*Methods*, Equation 5.2).

In a high emissions scenario (RCP8.5), we project that by end-of-century, most of the

world is expected to increase net annual per-capita electricity consumption and decrease consumption of other fuels due to climate change (Figure 5.2A). The amplitude of these effects reflects differences in incomes and climates across locations. Hot and wealthy locations exhibit large net increases in electricity consumption, although very cold locations exhibit net declines where warming does not increase the number of hot days enough to offset the loss of cold days. Low income regions, such as much of Sub-Saharan Africa, do not increase electricity consumption as dramatically because they are projected to still have relatively low incomes at end-of-century (e.g. see Ethiopia in Figure 5.1A). Declines in consumption of other fuels are projected throughout the world, consistent with the use of these fuels for heating across a wider range of incomes.

To understand the scale of these impacts, they can be compared to current levels of energy consumption (Figure 5.2B). In many of today's rich countries, impacts at end-of-century are projected to be modest relative to current consumption, e.g. a +2% relative increase in annual US electricity consumption. This small magnitude is both due to high current consumption levels in conjunction with the fact that many rich countries are in temperate climates, where large projected increases and decreases in electricity consumption – from more hot and fewer cold days, respectively – offset one another. In contrast, in many of the poorest and/or most populous countries, the additional consumption imposed by climate change is projected to be substantial relative to current consumption, e.g. a +1600% relative increase in annual Nigerian electricity consumption. This is due both to uniformly hot temperatures and very low levels of current energy use.

Aggregating energy impacts globally, we project that in a high emissions scenario, annual electricity consumption will increase due to climate change by 0.97 GJ per capita (90% C.I. = $[0.40, 2.02]$, $p<0.001$) in 2099 (RCP8.5), while consumption of other fuels will decline 2.80 GJ per capita ($[-5.81, -0.99]$, $p<0.01$) (Figure 5.2C). Estimates in a moderate emissions scenario (RCP4.5) are 0.33 and 1.07 GJ per capita, respectively. (Electricity impacts do not include the primary energy lost in conversion to electricity.) It is notable that ignoring the effects of income growth and climate adaptation on the energy-temperature response would have resulted in dramatic underestimation of projected changes to global energy consumption due to warming (green lines, Figure 5.2C).

We monetize the climate-change induced changes in total energy consumption (electricity and other fuels combined) to develop a measure of the economic *damages* from climate change, i.e. all economic resources that would be available for other purposes in the absence of warming. In a baseline scenario of future real energy price growth of 1.4% per year (the historical growth rate of US energy prices), we project that end-of-century warming will cause net energy expenditure declines in much of the world, although there are net increases in many tropical and subtropical middle-income regions, such as portions of India, China, Indonesia and Mexico (Figure 5.3A). This pattern occurs because currently low income countries will likely be rich enough by end-of-century to consume other fuels on cold days but not rich enough to consume electricity on hot days, thus they experience savings from warming because it reduces other fuel costs (e.g. Ethiopia in Figure 5.1A). Hot middle income countries will be rich enough to spend on electricity for cooling in the future, so in some regions the additional spending on electricity during hot days outweighs the savings on cold days (e.g. India in Figure 5.1A). The largest overall savings are projected to occur among today's richest locations (*Appendix* 5.H).

Aggregating damages globally, we project modest net savings at end-of-century due to climate change, amounting to 0.2 % ($[-0.5\%, 0.1\%]$, $p<0.2$) and 0.1 % ($[-0.2\%, 0.0\%]$,

p<0.15) of 2099 world GDP in RCP 8.5 and RCP 4.5, respectively (Figure 5.3B). The magnitude of net global savings is similar across alternative pricing scenarios (*Appendix* 5.D, Figure 5.D.1). This result differs qualitatively from the increased energy spending reported in prior studies[14,24,26–28,37] that focused on electricity consumption in wealthy regions. Because low and middle-income populations spend little on electricity to cool and other fuels are consumed everywhere almost exclusively for heating, projections that use only electricity-temperature responses from high-income populations will overestimate new cooling expenditures and underestimate savings from reduced heating, leading to systematic overestimation of the total energy damages from climate change.

Damage functions describe the relationship between ΔGMST and global aggregate damages in a sector or the economy as a whole—they are at the heart of all IAMs used to develop SCC estimates,[1,4,9,10] informing mitigation policy implications by summarizing the costs of additional warming. We construct the first empirically-based global damage functions for energy consumption using the method in refs. [23,37], organizing global aggregate costs into functions of realized ΔGMST across 33,000 simulations (Figure 5.3C). These damage functions evolve over time, thereby capturing the influence of changing demographics, rising incomes, and warming local climates (See *Methods*). Damages are slightly quadratic in GMST anomaly, although essentially linear, with an additional +1° ΔGMST warming at end-of-century (relative to 2001-2010 average) increasing annual consumption of electricity 3.96 EJ ([3.92, 4.01], p<0.001) and decreasing consumption of other fuels 10.15 EJ ([10.03, 10.26], p<0.001), causing a net reduction in energy consumption costs by $181 billion ([177, 185], p<0.001). Earlier and later damage functions are less and more steep, respectively, primarily due to trends in income and population (*Appendix* 5.E, Figure 5.E.1).

Since CO_2 is long-lived in the atmosphere, the US National Academy of Sciences recommends computing SCC values that capture damages through to the year 2300.[3] To do this, we combine our empirically-derived damage functions with the Finite Amplitude Impulse Response (FAIR) climate model[22] to project through to 2300 the distribution of ΔGMST responses to the emission of a marginal 1 GtC of CO_2 (*Appendix* 5.F) (The high-resolution CMIP5 model runs used above to project spatially granular impacts end in 2100.) A CO_2 pulse emitted today perturbs the future trajectory of atmospheric CO_2 concentrations nonlinearly, affected by the half life of CO_2 in the atmosphere as it is stored and released in the oceans and biosphere (Figure 5.4A-B). This results in future ΔGMST that deviates from the baseline scenario, which in turn causes a stream of energy damages in future years (Figure 5.4C-D). The partial SCC from energy consumption is the net present value of these annual damages.

We find that one ton of CO_2 emitted today generates a total energy consumption burden valued at -$1.16 ([−4.76, −0.14], p<0.05) in net present value under the high-emissions RCP8.5 scenario and -$1.08 ([−4.29, −0.26], p<0.05) under a moderate-emissions RCP4.5 scenario (3% discount rate, 1.4% price growth scenario, Table 5.4E, Panel I). Our finding that the partial SCC from energy consumption is negative and small in magnitude is broadly robust across multiple pricing scenarios (Table 5.4E, Panels II-III; Table 5.G.1), although it is possible to obtain positive values (e.g. $1.20-9.96) by assuming that cooling and heating energy uses trend continuously at historical rates for the next several centuries, after accounting for the effects of income levels and climate (*Appendix* Table 5.I.1). This could occur, in principle, if the price of cooling technology falls indefinitely relative to other goods and services without a corresponding trend in

efficiency (*Appendix* 5.I.3), although we believe such a scenario is unlikely.

5.3 Discussion

Our estimates are the first, to our knowledge, to make use of globally comprehensive data and empirical relationships to compute a global partial SCC for energy consumption. This approach reveals the critical role of economic development in shaping how energy consumption patterns respond to climate change, as we find that much of the world will remain too poor in coming decades to spend substantially on energy-intensive cooling technologies. Our approach also demonstrates the importance of accounting for non-electricity energy consumption since global populations use other fuels to cope with cold temperatures even at low income levels. Together, these two factors explain why our analysis indicates that total global energy expenditures are not likely to increase dramatically in response to warming, and why marginal emissions today may in fact produce savings in global energy expenditures.

The modest magnitude of the aggregate global impacts, however, masks substantial important shifts in projected energy consumption. Most notably, projected impacts of warming in many of today's emerging economies may impose substantial costs and represent a large fraction of current consumption. For instance, we project that climate change will increase end-of-century electricity consumption in India by over 100% of its current consumption (Figure 5.2B).

Our approach is designed to isolate the effect of future climate change on energy consumption under given emissions and socioeconomic pathways. However, a natural area for future exploration is to account for feedback effects that climate change and climate policy may introduce via changes in the energy supply mix, the trajectory of CO_2 emissions, and technological innovation.[38,39] While future work should explore the potential for these factors to alter the partial SCC for energy consumption, we believe at least the CO_2 emissions feedbacks are likely immaterial, given the small size of our aggregate estimated impacts (*Appendix* 5.K).

Another extension to this analysis would be to account for additional types of future technological advancements that may affect the energy consumption response to climate change (e.g. changes in the relative cost of cooling technologies). Because we allow energy-temperature responses to evolve with rising incomes and temperatures in our projection, our estimates reflect historical trends in advancement and diffusion of technology that occur with changes in these two factors. In Appendix 5.I.3, we econometrically model more aggressive assumptions about technological change and find that substantially larger values of the partial SCC for energy consumption can arise. However, the particular assumptions required to generate this result (i.e. falling cost without efficiency gains) are unlikely to hold. The difficulty of predicting the direction and magnitude of unprecedented technological innovation under climate change underscores the critical need for further research in this area.[40]

The results of our analysis contrast with estimates derived from alternative approaches. For example, the numerical-theoretical FUND IAM[41] — the only modeling framework where direct comparison is possible — estimates a partial SCC for energy consumption[24] of \$8 per ton of CO_2, which constitutes 90% of its total SCC estimate (high emissions scenario, 3% discount rate). However, its developers acknowledge enormous uncertainty in this number.[8] Our findings are consistent with the concern that such prior estimates

rely on a limited empirical basis for calibration[42–45] and underscore the importance of developing a representative empirical approach.

We demonstrate the feasibility of combining global data, econometrics, detailed climate models, and modern computing to estimate a partial SCC for energy consumption. However, a total SCC, composed of many partial SCCs for different sectors of the economy, would be required to determine the full social cost of warming to global society. Our approach can be extended to the full range of outcomes potentially affected by climate (e.g. mortality,[23] agriculture, labor), thereby providing an empirically-based characterization of the total SCC.

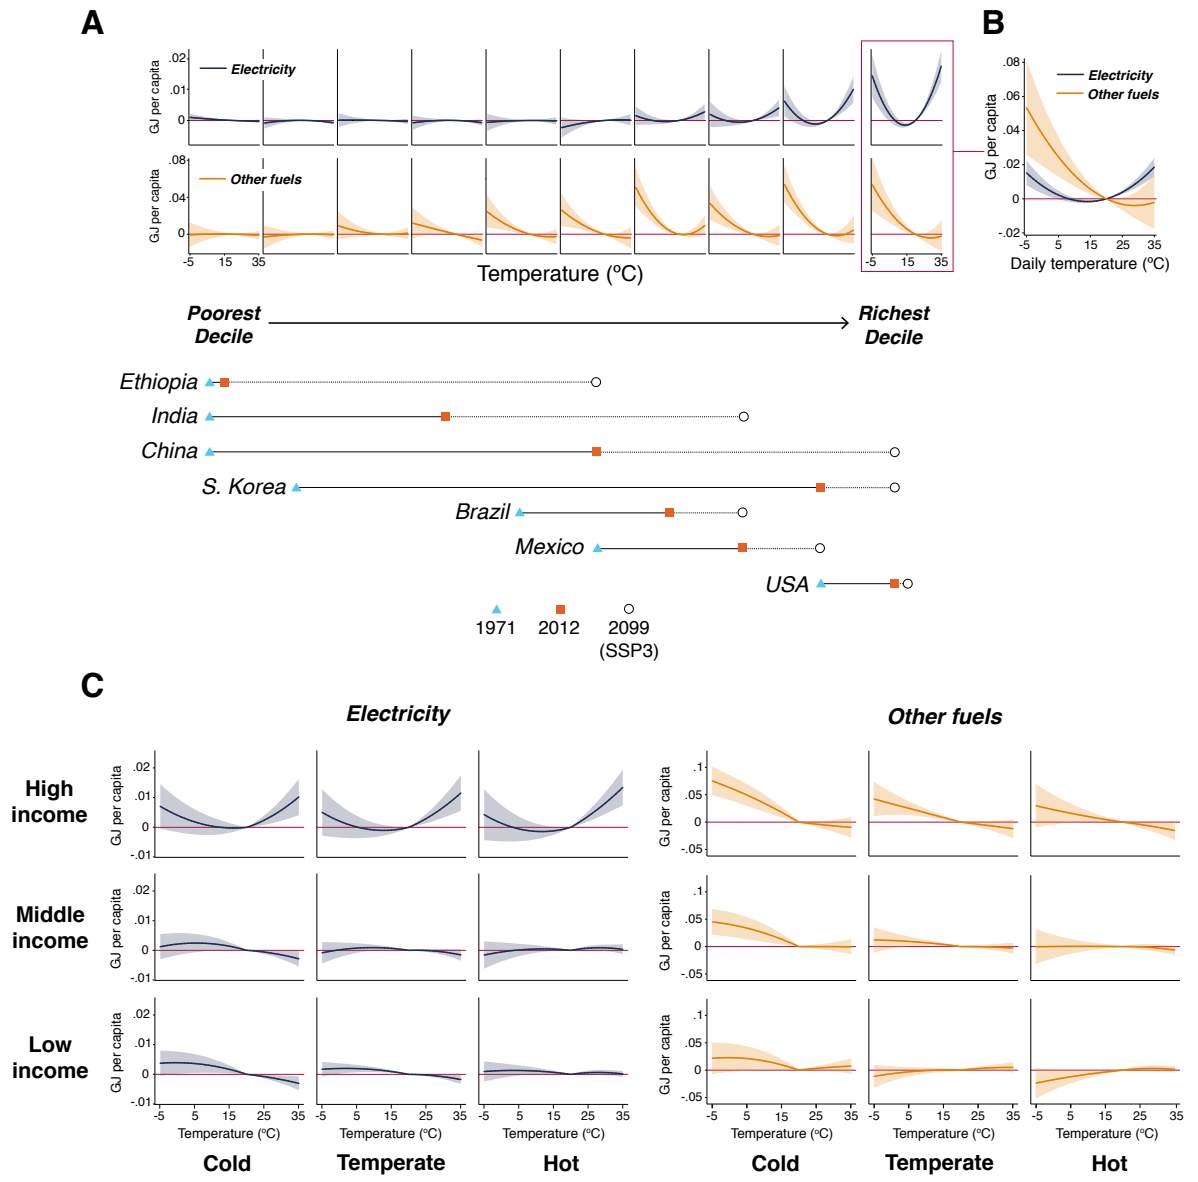

Figure 5.1: Energy consumption and temperature. Figure 5.1 displays various estimates of energy-temperature responses for electricity and end-use other fuels consumption that together account for all energy consumption. The height of each curve represents the additional per-capita daily consumption (in GJ) at a daily average temperature denoted on the horizontal axis, relative to a reference daily average temperature of 20°C. Shaded areas indicate 95% confidence intervals. **(A)** The energy-temperature response by income (*Appendix* Equation 5.C.3). Separate response curves are displayed for each decile of 15-year moving-average annual GDP per-capita in our sample, which consists of 146 countries over 1971-2010. The positions of selected country-years across these in-sample income deciles are indicated, with future income drawn from the Shared Socioeconomic Pathway 3 (SSP3) scenario.[18] **(B)** Reproduction of the electricity- and other fuels-temperature responses for the richest decile on a common scale. **(C)** The results from an econometric specification that models heterogeneity in the energy-temperature response due to both income and long-run climate (*Methods*, Equation 5.1. Each cell within a matrix illustrates a predicted response for a level of income and long-run climate. Cells are ordered vertically by income terciles (increasing income from bottom to top) and horizontally by terciles of annual cooling degree-days (increasingly warm climate from left to right, see *Methods* and *Appendix* 5.C.3).

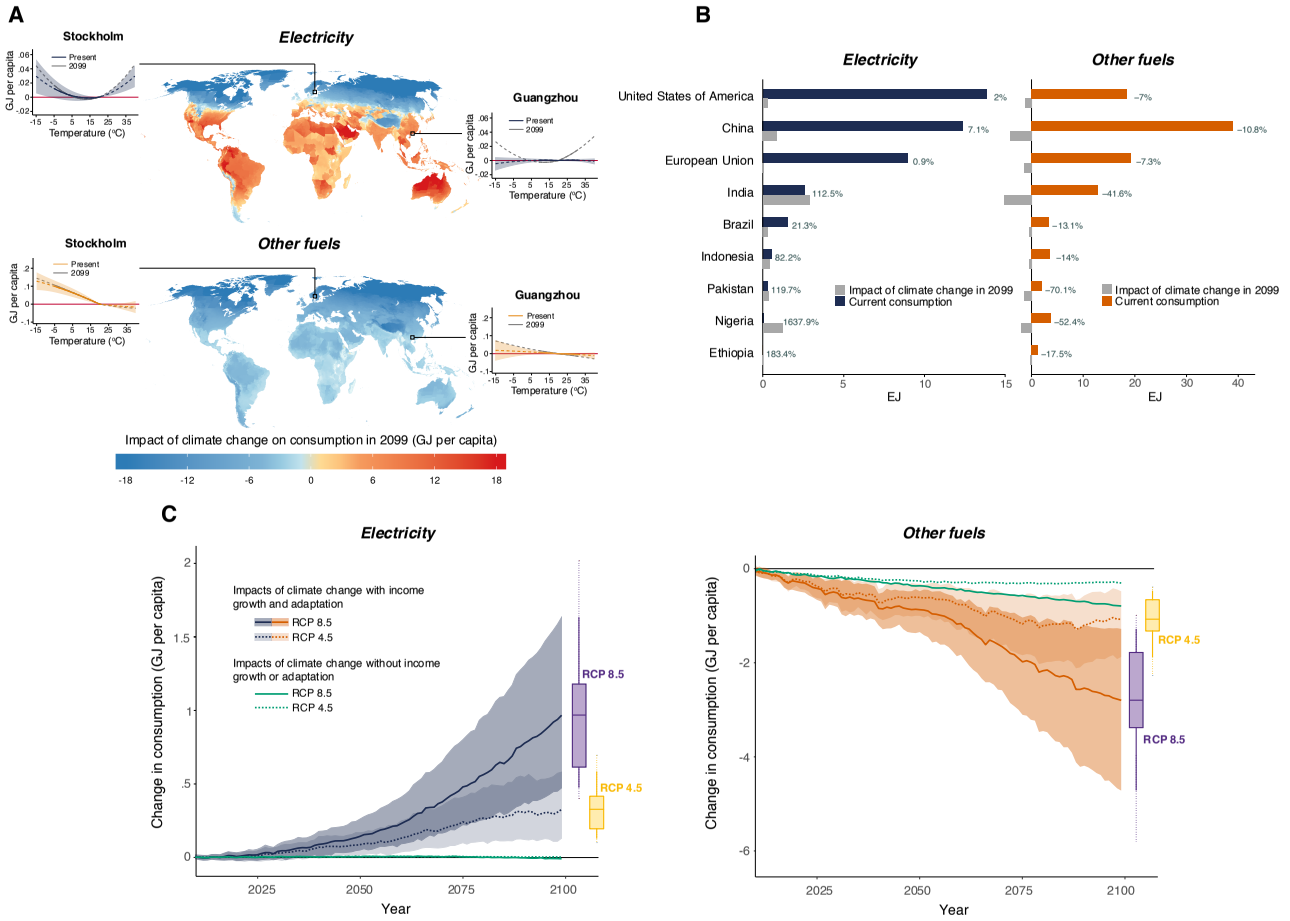

Figure 5.2: Projected impact of climate change on energy use in the twenty-first century. Projected impacts of climate change under high (RCP8.5) or moderate (RCP4.5) emissions scenarios and the SSP3 socioeconomic scenario on consumption of electricity and other fuels. Means are across 33 climate models and model surrogates. **(A)** The spatial distribution of consumption impacts under RCP8.5 in 2099 across 24,378 geographic regions. Energy-temperature responses reflect each region's projected income, climate, and adaptive adjustments in that year. Insets illustrate the change of energy-temperature responses between the present and 2099 for two example cities— Stockholm and Guangzhou. Dashed portions indicate regions of response functions that are outside of the range of temperatures historically experienced in each city. **(B)** Aggregating to the country level, gray bars show the total impact in 2099 (RCP8.5) for selected countries alongside their current consumption level of that type of energy (blue=electricity; orange=other fuels). Percentages are magnitude of gray bars relative to colored bars. **(C)** Aggregating globally, time series of total global impacts. Shaded areas indicate 10^{th}-90^{th} percentile of the distribution over projections accounting for climate model and econometric uncertainty (*Appendix* 5.C.5). Boxplots show full distribution of impacts in 2099 (boxes=inter-quartile range; solid whiskers= 10-90th percentiles; dashed whiskers=5-95th percentiles). Green lines illustrate projected impacts if present-day energy-temperature responses are held fixed and do not respond to rising incomes and changing temperatures.

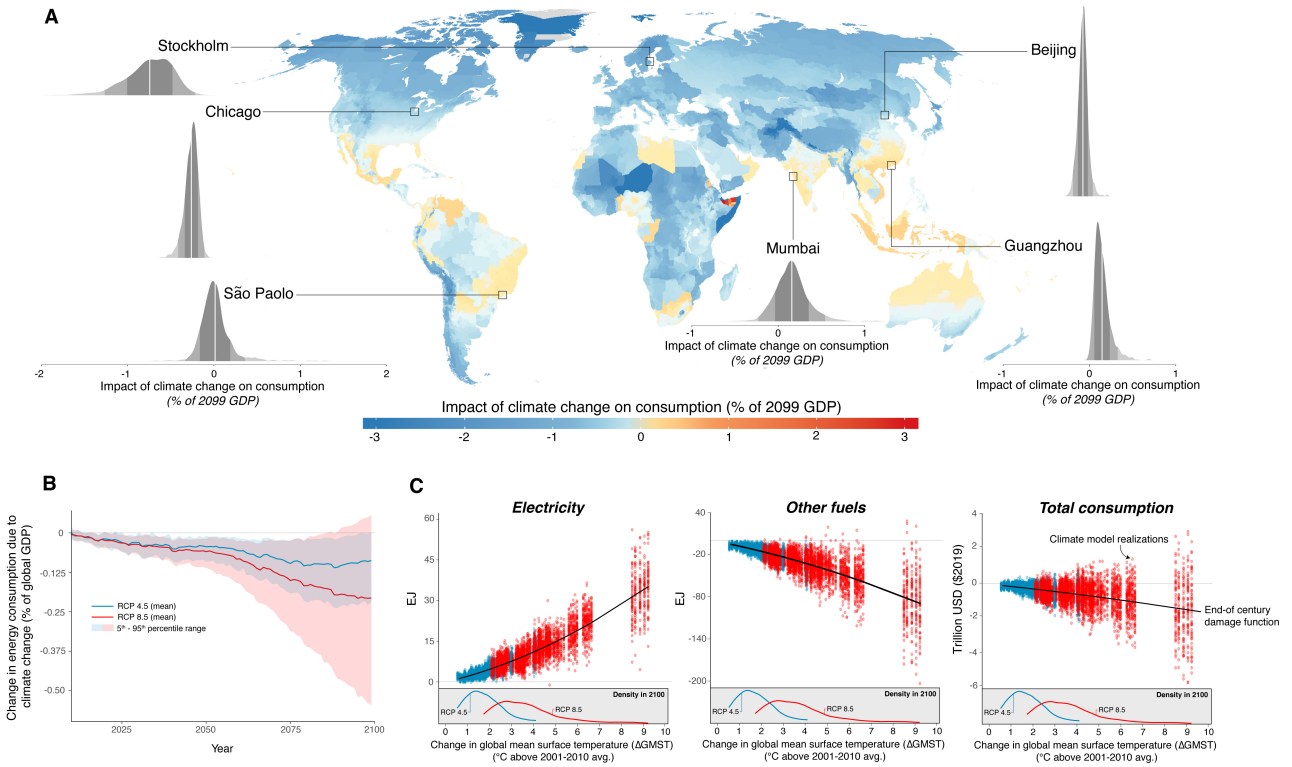

Figure 5.3: Economic damage from energy consumption impacts of climate change. (A) Map of annual climate-change induced changes in total energy expenditures (electricity and other fuels combined) in 2099 under a high emissions scenario (RCP8.5), the SSP3 socioeconomic scenario, and the 1.4% annual price growth scenario. Damage is expressed as a fraction of 2099 local GDP in each of 24,378 geographic regions. The map depicts mean estimates across 33 climate models and model surrogates, while probability density functions for selected cities plot the full distribution of projected impacts accounting for climate model and econometric uncertainty (*Appendix* 5.C.5). In each density plot, solid white lines indicate the mean estimate shown on the map, while shading indicates 1σ, 2σ, and 3σ deviations from the mean. **(B)** Time series of globally aggregated total energy consumption damages under both high (RCP8.5) and moderate (RCP4.5) emissions scenarios, with damages in each year expressed as a percent of global GDP in that year. Solid lines are means across models, shaded areas indicate the 5^{th}-95^{th} percentile of the distribution of impacts accounting for both climate model and econometric uncertainty. **(C)** Empirically-derived global energy damage functions at end-of-century. Total global electricity consumption impacts (left), other fuels consumption impacts (middle), and energy consumption damages (right) are indexed against the global mean surface temperature anomaly (ΔGMST) realized in each climate model simulation (blue dots=RCP 4.5; red dots=RCP 8.5; see *Methods*). Probability density functions below display the distribution of ΔGMST anomalies at end of century in each emissions scenarios across all 33 climate models and model surrogates.

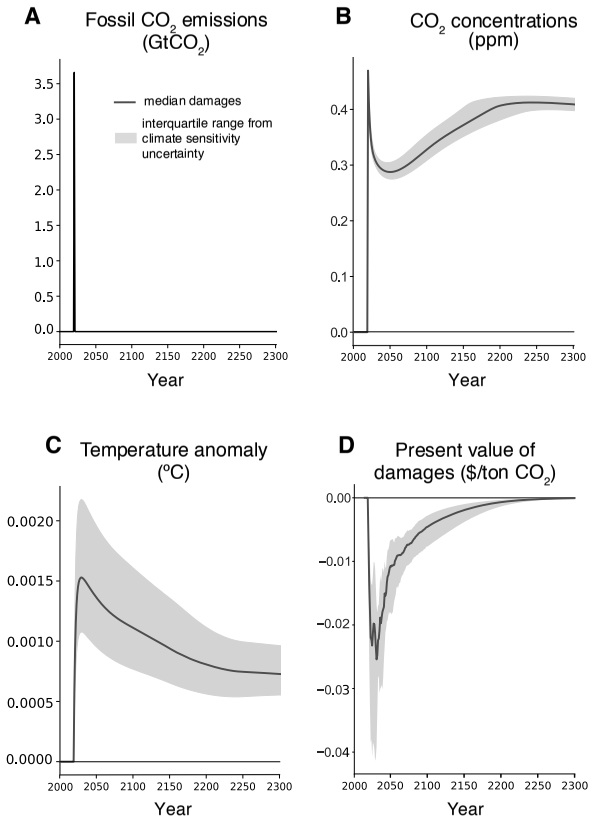

E. Partial Social Cost of Carbon Estimates

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	I: 1.4% price growth		
RCP 8.5	-1.51 [-6.59,0.06]	-1.16 [-4.76,-0.14]	-0.60 [-2.24,-0.19]
RCP 4.5	-1.37 [-6.00,-0.20]	-1.08 [-4.29,-0.26]	-0.58 [-1.98,-0.19]
	II: 0% price growth		
RCP 8.5	-0.72 [-2.63,-0.15]	-0.61 [-2.19,-0.17]	-0.39 [-1.39,-0.13]
RCP 4.5	-0.66 [-2.42,-0.21]	-0.57 [-2.00,-0.19]	-0.37 [-1.24,-0.13]
	III: MERGE-ETL 6.0 prices		
RCP 8.5	-1.12 [-3.88,-0.31]	-0.82 [-2.80,-0.24]	-0.39 [-1.38,-0.12]
RCP 4.5	-1.15 [-4.00,-0.40]	-0.83 [-2.83,-0.29]	-0.38 [-1.29,-0.12]

Figure 5.4: Social costs of carbon for global energy consumption. (A)-(D) The effects of a CO_2 pulse in the RCP8.5 emissions scenario, using the FAIR simple climate model.[22] Black line uses the default configuration, shaded area is the interquartile range of projections, sampling from a constrained joint distribution of climate parameters (see *Appendix* 5.F.2). **(A)** A 1GtC pulse released in 2020. **(B)** The effect on atmospheric CO_2 concentrations, relative to the baseline. **(C)** The impact on global mean surface temperature anomalies (ΔGMST). **(D)** The change in the discounted flow of damages from energy consumption estimated with the empirically-derived damage function in Figure 5.3C (3% annual discount rate). Integration of this flow is the partial social cost of carbon for energy consumption. **(E)** Estimates of the partial social cost of carbon for excess energy consumption costs under high (RCP8.5) and moderate (RCP4.5) emissions scenarios, assuming various annual discount rates recommended by refs. [2,3] (2.5%, 3%, %5) and future energy price scenarios (1.4% annual price growth, 0% price growth, and prices projected by the MERGE-ETL 6.0 IAM[46]). All estimates are under the socioeconomic scenario SSP3. Values in brackets are 5^{th}-95^{th} percentile ranges that account for uncertainty in the damage function (sampling from quantiles) and climate models (sampling from uncertain climate parameters; see Appendix 5.F.4). Additional partial SCC estimates exploring sensitivity to alternative pricing scenarios and socioeconomic scenarios are in Appendix 5.G.

5.4 Methods

Here we provide an overview of the data and methods used to complete each of the five stages of analysis that compose our modular approach to constructing a partial SCC. Details on each stage can be found in the Online Appendix.

5.4.1 Step 1: Data assembly

We compile a comprehensive dataset on historical energy consumption, climate, and income, as well as future projections of climate, income, populations, and energy prices (*Appendix* 5.A). Historical data are used to econometrically estimate the energy-temperature response, and how it differs by energy type (electricity and other fuels), income, and climate. Future projection data are used to generate high-resolution projected impacts of climate change, accounting for the effects of income growth and warming on the shape of the energy-temperature response.

Historical datasets Annual data on final consumption of electricity and other fuels for 146 countries from 1971 to 2010 were obtained from the International Energy Agency's (IEA) *World Energy Balances* dataset.[34] We take electricity consumption directly from the dataset, while other fuels consumption is constructed by aggregating over coal, peat, oil shale and oil sands, oil products, natural gas, solar, wind, goethermal, biofuels and waste, heat, and heat production from non-specified combustible fuels. For both electricity and other fuels, we aggregate over the industrial, commercial/public services, residential, agricultural, forestry, fishing, and non-specified sectors. Data inconsistencies and quality issues in the IEA's records are extensively documented.[34] We classify every such change in record-keeping methodology and employ specific data preparation and econometric techniques to address each individually (*Appendix* 5.A.1).

Historical data on daily average temperature and precipitation, as well as historical climatologies, are obtained from the Global Meteorological Forcing Dataset, v1 (GMFD),[33] a global gridded ($0.25° \times 0.25°$) daily climate record available from 1948 to 2010.[16] We link high-resolution daily climate data to country-level annual energy consumption data using a procedure detailed in Appendix 5.A.2.4 that preserves nonlinearity in the energy-temperature response (*Appendix* 5.A.2.4).

We obtain historical values of country-level annual income per capita from within the International Energy Agency's *World Energy Balances* dataset, which in turn sources these data from the World Bank.

Datasets of future projections We use a set of 21 high-resolution ($0.25° \times 0.25°$) bias-corrected global climate projections that provide daily temperature and precipitation through the year 2099 from the NASA Earth Exchange (NEX) Global Daily Downscaled Projections (GDDP) dataset.[47] We obtain climate projections based on two standardized emissions scenarios: Representative Concentration Pathways 4.5 (RCP4.5, an emissions stabilization scenario) and 8.5 (RCP8.5, a scenario with intensive growth in fossil fuel emissions).[48–50] Because this set of 21 climate models systematically underestimates tail risks of future climate change,[36,51] we assign probabilistic weights to climate projections and use 12 surrogate models that describe local climate outcomes in the tails of the climate sensitivity distribution.[36] The 21 models and 12 surrogate models are treated identically in our calculations and are referred to as the surrogate/model mixed ensemble (SMME). Full details on the SMME climate projections are in Appendix 5.A.2.3. Gridded output from these projections is aggregated to 24,378 globally comprehensive agglomer-

ated political units we call *impact regions* using the same method applied to historical climate data (*Appendix* 5.A.2.4). Impact regions are constructed to (i) respect national borders, (ii) be roughly equal in population across regions, and (iii) have approximately homogenous within-region climatic conditions (*Appendix* 5.B).

Projections of national populations and income per capita are derived from the Shared Socioeconomic Pathways (SSPs),[52] a set of scenarios of socioeconomic development over the 21^{st} century in the absence of climate impacts or policy. We utilize scenarios SSP2, SSP3, and SSP4.[53] National population and income per capita projections are respectively allocated to 24,378 impact regions based on current satellite-based within-country population distributions[54] (*Appendix* 5.A.3.3) and current nighttime light satellite imagery from the NOAA Defense Meteorological Satellite Program (DMSP) (*Appendix* 5.A.3.2).

The price trajectories we use to monetize estimated impacts of climate change are constructed based on either of two distinct data sources— present-day statistics from the IEA or price projections from integrated assessment models (IAMs). We obtain present-day average electricity generation costs by region of the world from the IEA's *World Energy Outlook 2017* (Figure 6.25); prices for other fuels are obtained from the IEA's *Energy Prices and Taxes Statistics* dataset. Price projections of electricity and other fuels' prices from 5 IAMs were obtained from IIASA's Scenario Explorer database.[46] Details on how prices are assigned across countries and over time can be found in Appendix 5.D.

5.4.2 Step 2: Econometric estimation of energy-temperature responses

Using historical data on energy consumption, climate, and income, we flexibly model annual electricity and other fuels consumption each as a function of daily average temperatures within a year, while accounting for heterogeneity in energy-temperature responses along the dimensions of both income and long-run climate. Let E_{jtc} denote consumption in GJ per capita in country j, year t, and fuel category c (electricity, other fuels). The temperature and precipitation vectors $\boldsymbol{T}_{jt}$ and $\boldsymbol{P}_{jt}$ contain country-by-year aggregations of nonlinear grid-cell-level transformations of daily temperature and precipitation, respectively. These vectors thus summarize the full distribution of daily average weather in country j, year t (*Appendix* 5.A.2.4).

To model heterogeneity by income, we use the 15-year moving average of a country's natural log of per-capita GDP ($\overline{LogGDPPC}_{jt}$), and to model heterogeneity by long-run climate we use a country's average annual cooling degree days and heating degree days over the sample period ($\overline{CDD}_j$ and $\overline{HDD}_j$). Annual cooling (heating) degree days are a common measure of exposure to warm (cold) temperatures and are defined as the cumulative deviations of daily average temperatures from a benchmark of 20° C, over all days in the year where the average temperature exceeded (fell below) 20° C. Because these measures do not change substantially over the historical record, we do not rely on long-run temporal variation within the timespan of the sample and instead use the average over the sample period (*Appendix* 5.C.3).

We estimate a specification of the following form:

$$E_{jtc} = f_c(\boldsymbol{T}_{jt}|\overline{LogGDPPC}_{jt}, \overline{CDD}_j, \overline{HDD}_j) + g_c(\boldsymbol{P}_{jt}) + \alpha_{jic} + \delta_{wtc} + \varepsilon_{jtc}. \tag{5.1}$$

In our main specifications, $\boldsymbol{T}_{jt}$ contains linear and quadratic terms of daily average tem-

peratures, each summed across the year, and $\boldsymbol{P}_{jt}$ contains linear and quadratic terms of daily cumulative precipitation, also each summed annually. Exploiting flexible interactions between the three income and long-run climate covariates ($\overline{LogGDPPC}_{jt}$, $\overline{CDD}_j$, and $\overline{HDD}_j$) and all terms in the temperature vector ($\boldsymbol{T}_{jt}$), we estimate an energy-temperature response f_c for fuel category c that is conditional on income and long-run climate. Details of this procedure can be found in Appendix 5.C.3.

All our econometric specifications include a full set of country-by-reporting regime intercepts, referred to here as "fixed effects" (α_{jic}), where reporting regimes (i) are time spans within a country where observations for a given fuel category are documented by the IEA to be comparable (*Appendix* 5.A.1). These fixed effects flexibly account for all permanent differences in energy consumption across country-regimes. In addition, we include world region-by-year fixed effects (δ_{wtc}) for each fuel category, where w indexes world regions based on UN classifications (Oceania, N. America, N. Europe, S. Europe, W. Europe, E. Europe, E. Asia, S.E. Asia, Central America/Caribbean, South America, sub-Saharan Africa, N. Africa/W. Asia, S. Asia). These fixed effects flexibly account for all world region-level trends and shocks in energy consumption. The use of fixed effects is more reliable than trying to individually control explicitly for determinants of energy consumption because it accounts for time-invariant and time-trending factors non-parametrically. We thus exploit random within-country, year-to-year variation in realized daily temperatures to identify a plausibly causal effect of historical temperature distributions on energy consumption. Finally, ε_{jtc} denotes the stochastic error term.

Due to evidence of unit root behavior in the dependent variable, we estimate Equation 5.1 in first-differences (*Appendix* 5.A.1). Furthermore we employ inverse variance weighting to address differences in data quality across reporting regimes (*Appendix* 5.C.1). Standard errors are clustered by country-fuel category-reporting regime.

While Equation 5.1 is designed to causally identify the effect of daily temperatures on energy consumption, it does not identify overall levels of energy consumption as these are absorbed in the spatial and temporal fixed effects. We therefore express estimated electricity- or other fuels-temperature responses as predicted consumption relative to a "mild" day with an average temperature of 20°C. The matrices of electricity- and other fuels-temperature responses in Figure 5.1C summarize the results from estimating Equation 5.1, while Figures 5.1A and 5.1B display responses that are estimated for each decile of the in-sample income distribution, but do not differ by long-run climate (*Appendix* 5.C.2).

Estimating Equation 5.1 enables us to use observable characteristics – $\overline{LogGDPPC}$, $\overline{CDD}$, and $\overline{HDD}$ – to predict energy-temperature responses at different points in time for each of 24,378 impact regions. For each impact region r and year t, we use the estimated function $\widehat{f}_c(\cdot)$ from Equation 5.1, along with 15-year moving averages of the covariates ($\overline{LogGDPPC}_{rt}$,$\overline{CDD}_{rt}$,$\overline{HDD}_{rt}$), to predict energy-temperature responses for each fuel category. Responses evolve over time as 15-year moving averages of the covariates change for a given impact region r, thereby reflecting the effects of adaptive behaviors that populations undertake as they become richer and/or are exposed to warmer climates. Figure 5.2A (inset) plots examples of such evolving responses for impact regions corresponding to the cities of Stockholm, Sweden (exhibiting small changes) and Guangzhou, China (exhibiting large changes).

5.4.3 Step 3: Projecting the impacts of climate change

To estimate future energy consumption impacts for each fuel category c and impact region r for each year from 2015 to 2099, we apply a set of probabilistic climate change projections to the spatially and temporally heterogenous energy-temperature responses. The distribution of future daily average temperatures under a given emissions scenario (RCP8.5 or RCP4.5) is obtained from the 33 projections in the SMME (*Appendix* 5.A.2.3).

Let $\boldsymbol{T}_{rt}$ represent a vector containing impact region-by-year aggregations of nonlinear grid-cell-level transformations of daily temperature in a future year t, under a warmer climate. In contrast, let $\boldsymbol{T}_{r2015}$ represent the counterfactual temperature vector for the same impact region under a climate that is the same as that of 2015. These vectors are constructed in exactly the same way as is done for the temperature vectors used in estimating Equation 5.1 (*Appendix* 5.A.2.4). The impact of climate change on fuel category c is expressed as the estimated change in consumption relative to a no-climate change counterfactual in which the future climate is the same as in 2015:

$$\begin{aligned} ImpactOfClimateChange_{crt} = & \underbrace{\widehat{f}_c(\boldsymbol{T}_{rt} \mid \overline{LogGDPPC}_{rt}, \overline{CDD}_{rt}, \overline{HDD}_{rt})}_{\substack{\text{Temperature-induced energy consumption under climate change} \\ \text{(with income growth and climate-driven adaptation)}}} \\ & - \underbrace{\widehat{f}_c(\boldsymbol{T}_{r2015} \mid \overline{LogGDPPC}_{rt}, \overline{CDD}_{r2015}, \overline{HDD}_{r2015})}_{\substack{\text{Temperature-induced energy consumption without climate change} \\ \text{(with income growth)}}}. \quad (5.2) \end{aligned}$$

The object $ImpactOfClimateChange_{crt}$ represents the change in annual per-capita electricity or other fuels consumption due to a shift in the temperature distribution under climate change, accounting for the evolution of energy-temperature responses as locations warm and incomes rise. It isolates the *additional* impact of climate change net of other factors (e.g. income) that will change in the future. The two projections are identical in every way, except for the climate. Thus, we evaluate the second term using future levels of income but use $\overline{CDD}$ and $\overline{HDD}$ values drawn from the 2015 climate distribution. All fixed effects and other controls cancel out and are therefore omitted.

We construct estimates of Equation 5.2 for all impact regions up to 2099 under emissions scenarios RCP8.5 and RCP4.5, using each of the 33 climate projections in the SMME. Figure 5.2A maps mean impact estimates across these 33 climate projections at year 2099 under RCP8.5, while 5.2B and 5.2C display impacts aggregated to the country and global levels respectively. Confidence intervals around the means are constructed to reflect both climatological and econometric sources of uncertainty. The distribution of impacts across 33 climate projections captures uncertainties in the climate system through to 2099, and we additionally capture uncertainty arising from econometric estimation of Equation 5.1 using the delta method.[55] Appendix 5.C.5 details the method used to combine both these independent sources of uncertainty.

To highlight the critical importance of income growth and climate-driven adaptation in shaping future energy-temperature responses, we also consider a "no-adaptation" impact

projection that ignores these factors (green lines in Figure 5.2C). To do this, we project

$$ImpactOfClimateChange_{crt}^{NoAdaptation} = \underbrace{\widehat{f}_c(\boldsymbol{T}_{rt}|\overline{LogGDPPC}_{r2015}, \overline{CDD}_{r2015}, \overline{HDD}_{r2015})}_{\substack{\text{Temperature-induced energy consumption under climate change} \\ \text{(no adaptation)}}} - \underbrace{\widehat{f}_c(\boldsymbol{T}_{r2015}|\overline{LogGDPPC}_{r2015}, \overline{CDD}_{r2015}, \overline{HDD}_{r2015})}_{\substack{\text{Temperature-induced energy consumption without climate change} \\ \text{(no adaptation)}}} \tag{5.3}$$

which captures the change in consumption responses due to future temperature, holding each impact region's income and climate fixed at 2015 values for all years in the projection.

5.4.4 Step 4: Estimating global energy damage functions

Our fourth step is to pool empirical estimates of climate change impacts constructed using Equation 5.2 to fit global energy damage functions, which express global energy consumption costs of climate change as a function of the change in global mean surface temperature relative to the 2001-2010 average level (ΔGMST).[1] These damage functions summarize the economic costs of all impacts measured in the detailed empirical analysis, demonstrating how they vary with the change global mean surface temperature.

Damage functions through 2099 are directly built from estimates of global costs (D_{tlps}, denominated in either EJ or dollars) in each year (t) using 33 climate models (l), two emissions scenarios (p), and a resampling of estimates (s) that captures uncertainty in the estimation of Equation 5.1. We interpret each of the resulting 33,000 simulation outputs D_{tlps} as a potential realization of damages that result from the spatial distribution of warming in model l, given the overall ΔGMST that is exhibited by that model under the emissions scenario p. Multiple simulations lead to an empirically-derived distribution of potential outcomes that are conditional on the ΔGMST value for the year, climate model, and emissions scenario used to generate that projection. To construct damage functions, we use these outcomes to estimate a conditional distribution of damages[23,37] using ordinary least squares, to obtain expected values, and quantile regressions, to capture uncertainty in damages conditional on ΔGMST.

In our projections of the future, the underlying population distribution and level of per capita income are evolving over time, thereby shaping the sensitivity of energy consumption to warming and through it, global damages. These changes over time require the construction of year-specific damage functions. Thus, we separately estimate a quadratic damage function in each year:

$$D(\Delta GMST, t)_{tlps} = \psi_0^t + \psi_1^t \Delta GMST_{tlp} + \psi_2^t \Delta GMST_{tlp}^2 + \varepsilon_{tlps}, \tag{5.4}$$

using all simulations within a 5-year window of t, thereby allowing the shape of the function $D(\Delta GMST, t)_{tlps}$ to evolve flexibly and smoothly over the century. Figure 5.3C displays examples of damage functions (expected values) at end-of-century, with each point in the scatterplot representing an individual realization of D_{tlps}. The leftmost and middle panels demonstrate examples of separate damage functions for electricity and other fuels respectively, where the realizations are denominated in EJ. The rightmost panel of Figure

5.3C displays a damage function for total energy consumption, denominated in dollars. To monetize the projected impacts of climate change on energy consumption, we apply country-specific real prices for electricity and other fuels to the projected quantity impacts, thus reflecting differential costs across geographies and fuels. Price trajectories up to 2099 are constructed in either of two ways: i) by extrapolating present-day prices under various price growth scenario assumptions or ii) by utilizing price projections developed in existing IAMs (*Appendix* 5.D).

In addition to estimating expected damages, we estimate 19 quantile regressions (for every fifth percentile from the 5^{th} to 95^{th} percentiles) to capture the full distribution of damages conditional on ΔGMST (*Appendix* 5.F.4). Quantile regressions also use a quadratic functional form (Equation 5.4), but with different coefficients and residuals. The resulting conditional distribution reflects econometric uncertainty in the impact estimates from which damages are constructed, as well as differences in the spatial patterns of warming exhibited across different climate models within the SMME.

As described in the next step, we use the estimated dollar-denominated damage functions to compute the net cost of global energy consumption associated with an additional ton of CO_2. Because CO_2 is long-lived in the atmosphere, the US National Academy of Sciences recommends computing SCC values that capture damages through to the year 2300.[3] Because CMIP5 models are not run beyond 2099, the SMME sample ends in 2099. Therefore, it is necessary to develop a separate approach to extend these damage functions beyond 2099. Details of this approach can be found in Appendix 5.E.

5.4.5 Step 5: Calculating the partial social cost of carbon

In the final step, we combine a probabilistic, simple climate-carbon cycle model with the set of damage functions described above to compute the partial SCC. The partial SCC at time t_0 is defined as the marginal social cost from elevated energy consumption imposed by the emission of a marginal ton of CO_2 at t_0 holding all other factors fixed (including the forecast trajectory of baseline greenhouse gas emissions). For a discount rate δ, this is expressed as:

$$Partial\ SCC_{t_0} = \sum_{t_0}^{2300} e^{-\delta t} \frac{d\hat{D}(\Delta GMST, t)}{d\Delta GMST_t} \frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}}, \tag{5.5}$$

where $\frac{d\widehat{\Delta GMST}_t}{dCO2_{t_0}}$ is the estimated increase in ΔGMST that occurs at each moment in time along the baseline climate trajectory (e.g. RCP8.5) as a result of a marginal unit of emissions at time t_0, which we approximate with an infinitesimally small pulse of CO2 emissions occurring at time t_0. The values $\frac{d\hat{D}(\Delta GMST,t)}{d\Delta GMST_t}$ are the marginal damages at each moment in time that occur as a result of this small change in future global temperatures; they are computed using the damage functions described in Equation 5.4. Following ref. [2] we hold δ fixed throughout a projection, but use the range of values $\delta \in \{0.025, 0.03, 0.05\}$ recommended by ref. [3] to explore the influence of the discount rate.

To calculate the change in ΔGMST_t due to a marginal pulse of CO_2 in 2020, we adapt a version of the Finite Amplitude Impulse Response (FAIR) simple climate model that has been developed especially for this type of calculation (*Appendix* 5.F).[22] Specifically, we use FAIR to calculate ΔGMST_t trajectories for emissions scenarios RCP4.5 and RCP8.5,

both with and without an exogenous impulse of 1Gt C (equivalent to 3.66Gt CO2) in the year 2020, an approximation of an infinitesimal emission for which the model numerics are stable. In FAIR, this emissions impulse perturbs the trajectory of atmospheric CO_2 concentrations and ΔGMST_t for 2020-2300, with dynamics that are influenced by the baseline RCP scenario. In each scenario, the trajectory of ΔGMST_t in the "RCP + pulse" simulation is differenced from the baseline RCP simulation to compute $\frac{d\widehat{\Delta GMST_t}}{dCO2_{t_0}}$, and the resulting damages are converted into USD per 1t CO_2 and discounted to the present.

To capture uncertainty in the climate physic represented in FAIR, we generate a distribution of future temperature trajectories by resampling the equilibrium climate sensitivity, the transient climate response, the short thermal adjustment time, and the time scale of rapid carbon uptake by the ocean mixed layer from a joint distribution that we constrain using findings from the literature (*Appendix* 5.F.2). The solid lines in Figures 5.4 indicate trajectories arising from the median values in FAIR's configuration parameters and the shaded areas indicate interquartile ranges calculated through this resampling. The final range of final uncertainty in projected damages combines this uncertainty in climate sensitivity with uncertainty in damages, conditional on the climate sensitivity, by also resampling from quantiles of the damage function.

Panels I, II, and III of Table 5.4E present partial SCC estimates under RCP8.5 and RCP4.5, assuming various discount rates and future energy price scenarios. Interquartile ranges (in brackets) account for econometric and climatological uncertainty (*Appendix* 5.F.4). Additional partial SCC estimates demonstrating sensitivity to alternative pricing scenarios, approaches to estimating post-2100 damages, and socioeconomic scenarios can be found in Appendix 5.G. Appendix 5.I.3 presents partial SCC estimates under a scenario of persistent technological progress where the cost of energy services continues to decline indefinitely based on historically measured rates without a corresponding trend in efficiency.

Chapter Review and Looking Ahead

In first chapter, we presented the first architecture that integrates best-available data, econometrics, and climate science to estimate climate damages worldwide at the local level, as well as aggregated global marginal damages. We applied this architecture to construct the first global empirical estimates of the impact of climate change on total non-transport end-use energy consumption, one of the most uncertain impacts in current integrated assessment models.[8] At end-of-century, we project annual global electricity consumption to rise roughly 4 EJ (1100 TWh, 6% of current global consumption) for each 1°C increase in global mean temperature, reflecting increased cooling demand, while direct consumption of other fuels declines 10.1 EJ (6% of current global consumption) per 1°C, reflecting reduced heating. Together, these estimates indicate that emission of 1 ton of CO_2 today produces global net savings in future aggregate energy consumption of about $1 in net present value (3% discount rate). By end-of-century, emerging economies in the tropics (e.g. India) are projected to increase electricity consumption dramatically, but these rising costs are offset by heating reductions in the wealthy economies of North America and Europe.

The next, and final, chapter closes with some concluding remarks.

References

[1] Nordhaus, W. D. An optimal transition path for controlling greenhouse gases. *Science* **258**, 1315–1319 (1992).

[2] Interagency Working Group on Socal Cost of Carbon. Social cost of carbon for regulatory impact analysis - under executive order 12866. Tech. Rep., United States Government (2010).

[3] National Academies of Sciences, Engineering, and Medicine. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide* (The National Academies Press, Washington, DC, 2017).

[4] Stern, N. Stern review report on the economics of climate change (2006).

[5] Tol, R. S. The economic effects of climate change. *Journal of economic perspectives* **23**, 29–51 (2009).

[6] Pindyck, R. S. Climate change policy: What do the models tell us? *Journal of Economic Literature* **51**, 860–872 (2013).

[7] Burke, M. *et al.* Opportunities for advances in climate change economics. *Science* **352**, 292–293 (2016).

[8] Anthoff, D. & Tol, R. S. The uncertainty about the social cost of carbon: A decomposition analysis using FUND. *Climatic Change* **117**, 515–530 (2013).

[9] Nordhaus, W. D. Estimates of the social cost of carbon: background and results from the rice-2011 model. Tech. Rep., National Bureau of Economic Research (2011).

[10] Tol, R. S. On the optimal control of carbon dioxide emissions: an application of fund. *Environmental Modeling & Assessment* **2**, 151–163 (1997). URL http://dx.doi.org/10.1023/A:1019017529030.

[11] Greenstone, M., Kopits, E. & Wolverton, A. Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy* **7**, 23–46 (2013).

[12] Pizer, W. *et al.* Using and improving the social cost of carbon. *Science* **346**, 1189–1190 (2014).

[13] Revesz, R. L. *et al.* Global warming: Improve economic models of climate change. *Nature News* **508**, 173 (2014).

[14] Deschênes, O. & Greenstone, M. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics* **3**, 152–185 (2011).

[15] Taylor, K. E., Stouffer, R. J. & Meehl, G. A. An overview of cmip5 and the experiment design. *Bulletin of the American Meteorological Society* **93**, 485 (2012).

[16] Auffhammer, M., Hsiang, S. M., Schlenker, W. & Sobel, A. Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy* **7**, 181–198 (2013).

[17] Kopp, R., Hsiang, S. & Oppenheimer, M. Empirically calibrating damage functions and considering stochasticity when integrated assessment models are used as decision tools. *Impacts World, Potsdam, Germany, May 27-30* (2013).

[18] O'Neill, B. C. *et al.* A new scenario framework for climate change research: the concept of shared socioeconomic pathways. *Climatic Change* **122**, 387–400 (2014).

[19] Rasmussen, D. J. & Kopp, R. E. Appendix A: Physical climate projections. Economic risks of climate change: An American prospectus. *Economic Risks of Climate Change: An American Prospectus* (2015). URL https://cup.columbia.edu/book/economic-risks-of-climate-change/9780231174565. 1510.00313.

[20] Hsiang, S. Climate econometrics. *Annual Review of Resource Economics* **8**, 43–75 (2016).

[21] Diaz, D. & Moore, F. Quantifying the economic risks of climate change. *Nature Climate Change* **7**, 774 (2017).

[22] Millar, R. J., Nicholls, Z. R., Friedlingstein, P. & Allen, M. R. 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**, 7213–7228 (2017).

[23] Carleton, T. *et al.* Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2018-51* (2019). URL https://ssrn.com/abstract=3224365.

[24] Diaz, D. B. Evaluating the Key Drivers of the US Government's Social Cost of Carbon: A Model Diagnostic and Inter-Comparison Study of Climate Impacts in DICE, FUND, and PAGE (2014).

[25] Davis, L. W. & Gertler, P. J. Contribution of air conditioning adoption to future energy use under global warming. *Proceedings of the National Academy of Sciences* 201423558 (2015).

[26] Auffhammer, M., Baylis, P. & Hausman, C. H. Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the united states. *Proceedings of the National Academy of Sciences* 201613193 (2017).

[27] Wenz, L., Levermann, A. & Auffhammer, M. North–south polarization of european electricity consumption under future warming. *Proceedings of the National Academy of Sciences* **114**, E7910–E7918 (2017).

[28] Auffhammer, M. Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data. Tech. Rep., National Bureau of Economic Research (2018).

[29] Hadley, S. W., Erickson, D. J., Hernandez, J. L., Broniak, C. T. & Blasing, T. Responses of energy use to climate change: A climate modeling study. *Geophysical research letters* **33** (2006).

[30] Zhou, Y., Eom, J. & Clarke, L. The effect of global climate change, population distribution, and climate mitigation on building energy use in the US and China. *Climatic Change* **119**, 979–992 (2013).

[31] Isaac, M. & Van Vuuren, D. P. Modeling global residential sector energy demand for heating and air conditioning in the context of climate change. *Energy policy* **37**, 507–521 (2009).

[32] Clarke, L. *et al.* Effects of long-term climate change on global building energy expenditures. *Energy Economics* **72**, 667–677 (2018).

[33] Sheffield, J., Goteti, G. & Wood, E. F. Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling. *Journal of Climate* **19**, 3088–3111 (2006).

[34] International Energy Agency. World energy balances (edition 2017) (2018). URL https://www.oecd-ilibrary.org/content/data/9ddec1c1-en.

[35] van Ruijven, B. J., De Cian, E. & Wing, I. S. Amplification of future energy demand growth due to climate change. *Nature Communications* **10**, 2762 (2019).

[36] Rasmussen, D. J., Meinshausen, M. & Kopp, R. E. Probability-weighted ensembles of US county-level climate projections for climate risk analysis. *J. Appl. Meteor. Climatol.* **55**, 2301–2322 (2016). URL http://journals.ametsoc.org/doi/abs/10.1175/JAMC-D-15-0302.1.

[37] Hsiang, S. *et al.* Estimating economic damage from climate change in the United States. *Science* **356**, 1362–1369 (2017).

[38] Tong, D. *et al.* Committed emissions from existing energy infrastructure jeopardize 1.5 C climate target. *Nature* **572**, 373–377 (2019).

[39] Woodard, D. L., Davis, S. J. & Randerson, J. T. Economic carbon cycle feedbacks may offset additional warming from natural feedbacks. *Proceedings of the National Academy of Sciences* **116**, 759–764 (2019).

[40] Acemoglu, D., Hemous, D., Barrage, L., Aghion, P. *et al.* Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. In *2019 Meeting Papers*, 1302 (Society for Economic Dynamics, 2019).

[41] Anthoff, D. & Tol, R. Fund-climate framework for uncertainty, negotiation and distribution. *Available at www. fund-model. org. Accessed October* **26**, 2015 (2014).

[42] Downing, T. E., Greener, R. A. & Eyre, N. The economic impacts of climate change: Assessment of fossil fuel cycles for the ExternE project. Environmental Change Unit, University of Oxford, and Eyre Energy Environment (1995).

[43] Downing, T. E., Eyre, N., Greener, R. & Blackwell, D. Full fuel cycle study: Evaluation of the global warming externality for fossil fuel cycles with and without CO2 abatament and for two reference scenarios. Report to the International Energy Agency greenhouse gas R & D programme, Environmental Change Unit (1996).

[44] Miller, K. & Hodgson, D. Modelling UK energy demand. In *Global Warming and Energy Demand*, 150–161 (Routledge, 2005).

[45] Anthoff, D. & Tol, R. S. The climate framework for uncertainty, negotiation and distribution (FUND): Technical description, version 3.8. *Technical Document* (2014).

[46] Huppmann, D. *et al.* IAMC 1.5 °C Scenario Explorer and Data hosted by IIASA. Integrated Assessment Modeling Consortium & International Institute for Applied Systems Analysis (2018).

[47] Thrasher, B., Maurer, E. P., McKellar, C. & Duffy, P. Technical note: Bias correcting climate model simulated daily temperature extremes with quantile mapping. *Hydrology and Earth System Sciences* **16**, 3309–3314 (2012).

[48] Riahi, K. *et al.* RCP 8.5—A scenario of comparatively high greenhouse gas emissions. *Climatic Change* **109**, 33–57 (2011).

[49] Thomson, A. M. *et al.* RCP 4.5: A pathway for stabilization of radiative forcing by 2100. *Climatic Change* **109**, 77 (2011). URL https://doi.org/10.1007/s10584-011-0151-4.

[50] Van Vuuren, D. P. *et al.* The representative concentration pathways: An overview. *Climatic Change* **109**, 5 (2011).

[51] Tebaldi, C. & Knutti, R. 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**, 2053–2075 (2007). URL http://rsta.royalsocietypublishing.org/content/365/1857/2053.

[52] IIASA Energy Program. Ssp database, version 1.1 [data set]. Tech. Rep., National Bureau of Economic Research (2016). URL https://tntcat.iiasa.ac.at/SspDb. Accessed25December,2016.

[53] Riahi, K. *et al.* The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: an overview. *Global Environmental Change* **42**, 153–168 (2017).

[54] Bright, E. A., Coleman, P. R., Rose, A. N. & Urban, M. L. LandScan 2011 (2012). Digital dataset: web.ornl.gov/sci/landscan/index.shtml.

[55] Wooldridge, J. M. *Econometric analysis of cross section and panel data* (MIT Press, Cambridge and London, 2002).

[56] Wood, A. W., Leung, L. R., Sridhar, V. & Lettenmaier, D. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic change* **62**, 189–216 (2004).

[57] Meinshausen, M., Raper, S. C. B. & Wigley, T. M. L. Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – part 1: Model description and calibration. *Atmos. Chem. Phys.* **11**, 1417–1456 (2011). URL http://www.atmos-chem-phys.net/11/1417/2011/.

[58] Mitchell, T. D. Pattern scaling: An examination of the accuracy of the technique for describing future climates. *Climatic Change* **60**, 217–242 (2003). URL http://link.springer.com/article/10.1023/A%3A1026035305597.

[59] Dellink, R., Chateau, J., Lanzi, E. & Magné, B. Long-term economic growth projections in the shared socioeconomic pathways. *Global Environmental Change* (2015).

[60] Samir, K. & Lutz, W. 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* (2014).

[61] Gennaioli, N., Porta, R. L., Silanes, F. L. D. & Shleifer, A. Growth in regions. *Journal of Economic Growth* **19**, 259–309 (2014). URL https://ideas.repec.org/a/kap/jecgro/v19y2014i3p259-309.html.

[62] Global Administrative Areas. GADM database of global administrative areas, version 2.0. Tech. Rep. (2012). URL www.gadm.org.Accessed25December,2016.

[63] Rahman, M., Rahman, R. & Pearson, L. Quantiles for finite mixtures of normal distributions. *International Journal of Mathematical Education in Science and Technology* **37**, 352–358 (2006).

[64] Kopp, R. E. & Mignone, B. K. The US government's social cost of carbon estimates after their first two years: Pathways for improvement. *Working paper* (2012).

[65] Collins, M., Knutti, R. *et al. Chapter 12: Long-term Climate Change: Projections, Commitments and Irreversibility* (Intergovernmental Panel on Climate Change, 2013). URL http://www.ipcc.ch/report/ar5/wg1/.

[66] Joos, F. *et al.* Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: A multi-model analysis. *Atmospheric Chemistry and Physics* **13**, 2793–2825 (2013).

[67] Geoffroy, O. *et al.* 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**, 1841–1857 (2013).

[68] Ricke, K. L. & Caldeira, K. Maximum warming occurs about one decade after a carbon dioxide emission. *Environmental Research Letters* **9**, 124002 (2014).

[69] Millar, R. J. *et al.* Model structure in observational constraints on transient climate response. *Climatic Change* **131**, 199–211 (2015).

[70] Mastrandrea, M. D. *et al.* Guidance note for lead authors of the IPCC Fifth Assessment Report on consistent treatment of uncertainties. Tech. Rep., Intergovernmental Panel on Climate Change (2010).

[71] Barreca, A., Clay, K., Deschenes, O., Greenstone, M. & Shapiro, J. S. Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy* **124**, 105–159 (2016).

[72] McGrath, G. Natural gas-fired electricity conversion efficiency grows as coal remains stable (2017). Accessed from: https://www.eia.gov/todayinenergy/detail.php?id=32572.

[73] Pachauri, R. K. *et al. Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change* (IPCC, 2014).

[74] U.S. Environmental Protection Agency. Emission Factors for Greenhouse Gas Inventories (2018). Accessed from: https://www.epa.gov/sites/production/files/2018-03/documents/emission-factors_mar_2018_0.pdf.

Appendix

Several of the raw datasets, numerical methods, and intermediate results described in this appendix are also described in the appendix of ref. [23]. We provide this information here for completeness.

5.A Data

5.A.1 Energy consumption data

Data assembly As described in *Methods*, we obtain data on final consumption of electricity and other fuels from the International Energy Agency's (IEA) *World Energy Balances* dataset. Electricity consumption is taken from the ELECTR variable code, and consumption of other fuels is obtained by aggregating over the following variable codes: COAL (coal and coal products); PEAT (peat and peat products); OILSHALE (oil shale and oil sands); TOTPRODS (oil products); NATGAS (natural gas); SOLWIND (solar/wind/other); GEOTHERM (geothermal); COMRENEW (biofuels and waste); HEAT (heat), and HEATNS (heat production from non-specified combustible fuels). For both electricity and other fuels, we aggregate over the following sectoral codes: TOTIND, which encompasses consumption in the industrial sector, and TOTOTHER, which encompasses consumption in the commercial/public services, residential, agricultural, forestry, fishing, and non-specified sectors. The non-specified sector includes consumption in the other sectors within TOTOTHER if disaggregated figures are not provided for those sectors.

Harmonization of energy consumption data across diverse reporting regimes The IEA extensively documents a range of data quality issues related to its energy consumption data. Data quality issues can arise for a variety of reasons, including lack of data, revisions to data, imputed data, and reporting inconsistencies. The documentation identifies such issues for individual countries over specific years, often explicitly noting "breaks in time series".[2] Our data preparation approach is based on upon reading and categorizing all documentation pertaining to each year of each country in the data.

When carrying out our analyses, we address documented data quality issues in one of two ways– imposing fixed effects that account for mean differences in energy consumption across time spans of observations documented to be incomparable, or dropping of these incomparable observations. Although there are 146 countries in the data, we use the IEA documentation to identify 275 distinct reporting "regimes" for electricity consumption (i.e. time spans within a country where reporting practices are documented to be comparable) and 294 regimes for other fuels consumption. All our regressions contain fixed effects at the country-regime level (i.e. an indicator variable for each country-by-regime), which account for mean differences in energy consumption across reporting regimes. While country-regime fixed effects are a powerful way to deal with known data quality issues, they are not always a sufficient remedy if observations fail to meet even basic standards of comparability. In some cases it is necessary to drop observations altogether. We drop an observation whenever its definition of a fuel or sector category is documented to be at odds with the standard definitions. For instance, we drop ob-

[2] For instance, in the documentation for Denmark, it is noted that "major revisions were made by the Danish administration for the 1990 to 2001 data, which may cause breaks in time series...".

servations from Sweden prior to 1993, as certain road transport fuel consumption was included under the commercial sector during those years, but not during other years for Sweden or any years for other countries. A total of 412 observations are dropped for electricity consumption and 1,117 observations are dropped for other fuels consumption. In our released code base, we provide cleaning scripts depicting precisely the fixed effects imposed and the observations dropped.

Unit root behavior in energy consumption data Despite data preparation measures to guard against quality issues, there continue to exist what appear to be persistent shocks in energy consumption, even within a country-regime. For example, Figure 5.A.1 displays the time series for other fuels consumption in Italy, with each color indicating a distinct regime. Persistent shocks to consumption do not always appear to be tied to a change in regime. Motivated by such patterns, we formally test for a unit root in electricity and other fuels consumption for each of the country-regime time series, using both the augmented Dickey-Fuller and Phillips-Perron tests under various lag lengths (including a time trend in each case). The null hypothesis is that there exists a unit root. Figure 5.A.2 plots the histogram of p-values from the country-regime time series tests. The left panel displays tests on electricity consumption time series while the right panel displays tests on other fuels consumption time series. Each histogram within a panel represents a variant of the unit root test (augmented Dickey-Fuller or Phillips-Perron, each with different lag length). In all variants of the test for both electricity and other fuels consumption, the mass of very high p-values (i.e. close to $p = 1$ and far above the conventional rejection threshold of $p = 0.05$) suggests that it is very difficult to reject a unit root in a large number of time series.[3]

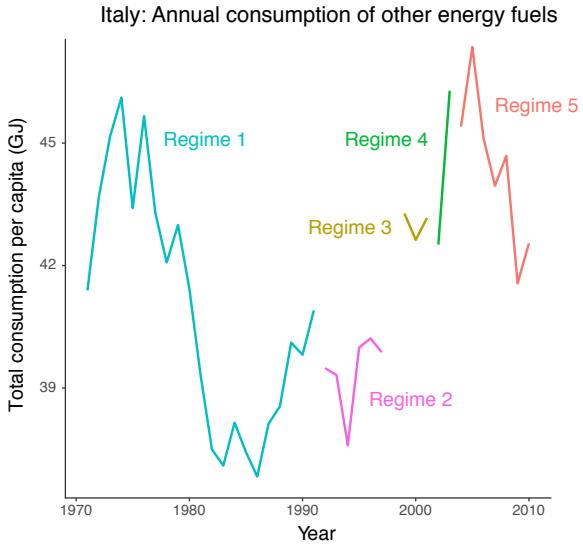

Figure 5.A.1: Per-capita consumption of other fuels in Italy (1971-2012). Each color represents a time span of years where reporting practices are documented by the IEA to be comparable (i.e. a "regime"). Persistent shocks to consumption do not always correspond to known changes in the regime (e.g. in 1989).

With strong suggestive evidence of unit root behavior in our energy consumption data, we estimate all regressions (detailed in Appendix 5.C) in first differences, as first

[3]It should be noted that unit root tests tend to be underpowered. Hence we consider the full distribution of p-values, rather than just how many p-values fall below a conventional rejection threshold such as $p = 0.05$.

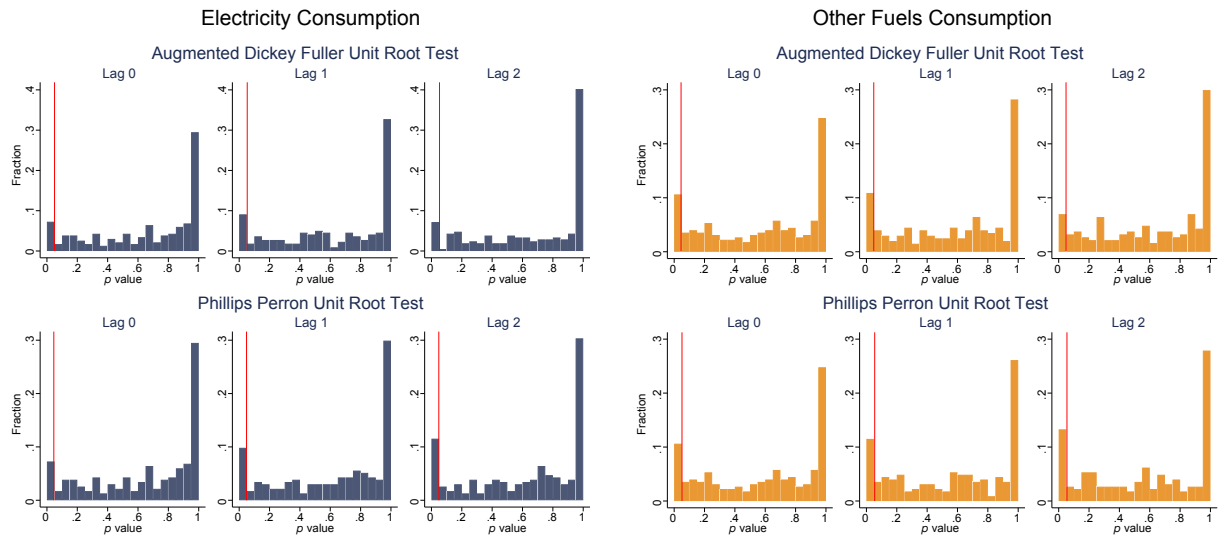

Figure 5.A.2: Unit root tests by country-regime time series. The left panel depicts the histogram of p-values from unit root tests on every country-regime time series for electricity consumption; the right panel does the same for other fuels consumption. The tests are repeated using different testing procedures (Augmented Dickey-Fuller, Phillips-Perron) and lag length (0,1,2) combinations, in each case including a trend. The existence of a unit root is always the null hypothesis, which fails to be rejected for a substantial number of time series (often with very high p-values), regardless of testing procedure and lag length. Vertical red lines mark p-values of 0.05.

differences remove confounding sources of spurious correlation that can enter due to unit root behavior.[55]

5.A.2 Climate data

This appendix describes the climate data that we use in this analysis as well as some of the methods employed to make these data consistent with the scale and resolution of the energy consumption data. Broadly speaking, we use two classes of climate data, the first being historical data to estimate energy-temperature responses, and the other being future climate data which are used to project the damages of climate change into the future under various emissions scenarios. We begin by describing the historical data, followed by the future projection data, and finally we detail the method we use to spatially and temporally aggregate these outputs to match the lower resolution energy consumption data.

5.A.2.1 Historical climate data

Data on historical climate exposure is used to estimate the energy-temperature response as well as the heterogeneity in the response by average climatology. For this estimation, we use the Global Meteorological Forcing Dataset, v1 (GMFD).[33] 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.[33] Data are available on a $0.25° \times 0.25°$ resolution grid from 1948-2010. The

temporal frequency is up to 3-hourly, but the daily data are used for this analysis. We obtain daily average temperatures and monthly average precipitation for all grid cells globally. A primary reason for using GMFD in our regression analysis is that GMFD is used to bias-correct the climate model projections (described below).

5.A.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 renders 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).[474] The NEX-GDDP dataset comprises 21 climate projections, which are downscaled from the output of GCM runs in the Coupled Model Intercomparison Project Phase 5 (CMIP5) archive.[15] The statistical downscaling algorithm used to generate the NEX-GDDP dataset is the Bias-Correction Spatial Disaggregation (BCSD) method,[47,56] 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.[33] A daily, quantile-specific relation between GCM historical period outputs and historical observations is derived from this comparison. This relation 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.[5]

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;[48,50] 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.

[4]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).

[5]Details are available in Appendix A of the NEX-GDDP documentation: https://gdo-dcp.ucllnl.org/downscaled_cmip_projections/techmemo/downscaled_climate.pdf

5.A.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.[36,51] 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[36] 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)[57] and (2) a form of linear pattern scaling[58] that preserves high-frequency variability to construct model surrogates to fill the tails of probability distribution that are not captured by the GCM ensembles. This method provides us with an additional 12 surrogate models.

The SMME method first divides the unit interval [0,1] into a set of bins. For this analysis, the bins are centered at the 1^{st}, 6^{th}, 11^{th}, 16^{th}, 33^{rd}, 50^{th}, 67^{th}, 82^{nd}, 89^{th}, 94^{th}, and 99^{th} percentiles. Bins are narrower in the tails to ensure samples are created for portions of the GMST probability distribution function that are not captured by CMIP5 models. The bounds and center of each bin are assigned corresponding quantiles of GMST anomalies for 2080-2099 from simple climate model (SCM) output; in the application here and that of ref.,[36] this output came from the MAGICC6 model,[57] 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 5.A.3, and the weights assigned to each GCM and surrogate under each emissions scenario are shown in Table 5.A.1.

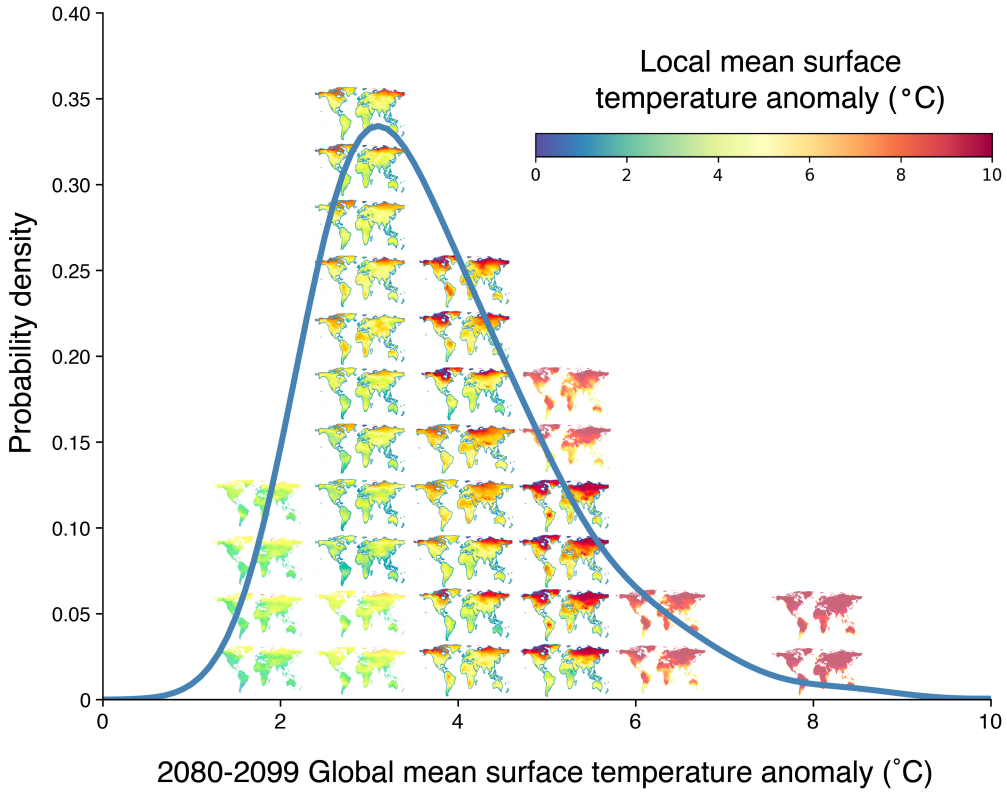

Figure 5.A.3: Future climate projections used in generating empirically-based climate change impact projections. This figure shows the 21 climate models (outlined maps) and 12 model surrogates (dimmed maps) that are weighted in climate change projections so that the weighted distribution of the 2080 to 2099 global mean surface temperature anomaly (ΔGMST) exhibited by the 33 total models matches the probability distribution of estimated ΔGMST responses (blue-gray line) under RCP8.5. For this construction, the anomaly is relative to values in 1986-2005.

5.A.2.4 Aggregation of gridded climate data to administrative boundaries

We link gridded daily historical climate data to country-year energy consumption data by aggregating grid cell information across space and time. Similarly, to generate future climate change impact projections at each of our 24,378 custom impact regions (*Appendix* 5.B), we aggregate grid cell climate projections to impact region scale. In both cases, nonlinear transformations of temperature and precipitation are computed at the grid-cell-by-day level before averaging values across space using population weights and finally summing over days within a year.

To see how this calculation is operationalized, consider the second-order polynomial specification for temperature used in our main set of results for estimation of Equation 5.1 in *Methods* (equivalent to 5.C.4 in Appendix 5.C). In this case, we begin with data on average temperatures for each day day d at each grid cell z, generating observations T_{zd}. These grid-level values must then be aggregated to the level of county j in year t. To do this, we first raise grid-level temperature to the power k, computing $(T_{zd})^k$ for $k \in \{1, 2\}$. We then take a spatial average of these values over country j, weighting the average by grid-level population (and accounting for fractional grid cells that fall partially within administrative units). Population weights are time-invariant and calculated from the 2010 Gridded Population of the World dataset.[6] We then sum these daily polynomial terms T_{zd}^p over days in the year t. The vector of annual, administrative-level-by-year

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

GCM or surrogate model name	Weight in RCP4.5	Weight in RCP8.5
ACCESS1-0	3.12%	3.00%
bcc-csm1-1	2.38%	2.29%
BNU-ESM	3.12%	3.00%
CanESM2	3.12%	3.00%
CCSM4	2.38%	2.29%
CESM1-BGC	2.38%	2.29%
CNRM-CM5	2.38%	2.29%
CSIRO-Mk3-6-0	3.12%	3.00%
GFDL-CM3	4.17%	4.00%
GFDL-ESM2G	4.17%	4.00%
GFDL-ESM2M	4.17%	4.00%
inmcm4	4.17%	4.00%
IPSL-CM5A-LR	3.12%	3.00%
IPSL-CM5A-MR	3.12%	3.00%
MIROC5	2.38%	2.29%
MIROC-ESM	4.17%	4.00%
MIROC-ESM-CHEM	4.17%	4.00%
MPI-ESM-LR	2.38%	2.29%
MPI-ESM-MR	2.38%	2.29%
MRI-CGCM3	9.37%	9.00%
NorESM1-M	9.37%	9.00%
surrogate_CanESM2_89	1.04%	1.00%
surrogate_CanESM2_94	4.17%	4.00%
surrogate_CanESM2_99	1.04%	1.00%
surrogate_GFDL-CM3_89	1.04%	1.00%
surrogate_GFDL-CM3_94	4.17%	4.00%
surrogate_GFDL-CM3_99	1.04%	1.00%
surrogate_GFDL-ESM2G_01	1.04%	1.00%
surrogate_GFDL-ESM2G_06	-	4.00%
surrogate_GFDL-ESM2G_11	1.04%	1.00%
surrogate_MRI-CGCM3_01	1.04%	1.00%
surrogate_MRI-CGCM3_06	4.17%	4.00%
surrogate_MRI-CGCM3_11	1.04%	1.00%

Table 5.A.1: Model weights in the surrogate model mixed ensemble (SMME). This table lists all 33 models in the SMME, and their corresponding model weights, which are used to generate probabilistic climate change impact projections. Details on the SMME method can be found in ref.[36] Details on how we use the SMME to capture climate model uncertainty when generating climate change impact projections can be found in Appendix 5.C.5.

temperature variables we use for estimation is thus:

$$\boldsymbol{T}_{jt} = \left[\sum_{d \in t} \sum_{z \in j} w_{zj} (T_{zd})^1, \sum_{d \in t} \sum_{z \in j} w_{zj} (T_{zd})^2 \right]$$

where w_{zj} is the share of j's population that falls into grid cell z, and where superscripts indicate polynomial powers. This procedure recovers grid-by-day-level nonlinearities in the energy-temperature (and energy-precipitation) response, because energy consumption is additive across time and space.[20] 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.

5.A.3 Socioeconomic data and downscaling methodologies

This section provides details of the socioeconomic data used throughout our analysis, which includes historical national incomes, future projections of incomes, and future projections of population counts. Additionally, because we require these variables at

high spatial resolution for future projections, we detail the downscaling procedures we use to disaggregate available socioeconomic projections, which are generally provided at relatively low resolution.

5.A.3.1 Historical income data

Our main specification (Equation 5.1 in *Methods* and Equation 5.C.4 in Appendix 5.C) estimates heterogeneity in energy-temperature responses as a function of income and long-run average climate in each location. We obtain country-level annual income data from within the International Energy Agency's World Energy Balances dataset; these income data originally are sourced by the IEA from the World Bank. For each country-year, the historical income variable is calculated as the 15-year moving average of the natural log of per-capita GDP.

5.A.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[59] and the International Institute for Applied Systems Analysis (IIASA) GDP model,[60] 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.[7]

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 (*Appendix* 5.B) 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 ref.[61] 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.[8] 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

[7] 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.

[8] 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.

approach models heterogeneity in income levels across impact regions, each region grows in the future at the same rate as the national country projection from the SSPs. We detail these two steps below.

Estimation of the GDP-nightlights relationship While there exists a growing literature linking economic output to nightlights intensity, we take an unconventional regression approach to recovering this relationship because our goal is to apportion national income within a country, as opposed to predicting the level of income at any given location. In particular, we are interested in the ratio $\frac{GDPpc_{rc}}{\sum_{r \in c} GCPpc_{rc}}$ for impact region r in country c, which will allow us to predict income at the impact region level, given projections of national GDP per capita from the SSPs, $\sum_{r \in c} GDPpc_{rc} = GDPpc_c^{SSP}$. Thus, we estimate a regression relating *relative* GDP per capita to *relative* nightlights intensity, where each administrative region's values are calculated as relative to the country mean. The dependent variable for administrative region i in country c and year t is thus $\frac{GDPpc_{ict}}{\sum_{i \in c} GDPpc_{ict}}$.[9] 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 the stable nightlights data product from 1992-2012 is used to construct time-varying measures of average nightlights intensity across an administrative region, NL_{ict}.

The regression we estimate is as follows:

$$\frac{GDPpc_{ict}}{\sum_{i \in c} GDPpc_{ict}} = \alpha + \beta ZNL_{ict} + \epsilon_{ict} \tag{5.A.1}$$

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

Allocation of national GDP to impact regions using relative nightlight intensity We use the estimated coefficients from Equation 5.A.1 to compute income at the 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 $GDPpc_{ct}^{SSP}$ 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.

5.A.3.3 Population projections and downscaling methodology

Future projections of national populations are derived from the International Institute for Applied Systems Analysis (IIASA)[60] population projections as part of the Shared Socioe-

[9] The income data available from ref.[61] are at the first administrative level (i.e. ADM1).

conomic Pathways (SSPs).[10] 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 ref.[18]

To assemble population projections for each of our 24,378 impact regions (*Appendix* 5.B), we downscale the country-level projections from the SSPs using 2011 high-resolution LandScan estimates of populations.[54] 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:

$$\widehat{pop}_{rtv} = \begin{cases} pop_{ctv}^{SSP}\left(\frac{pop_{r,2011}^{LandScan}}{\sum_{r\in c} pop_{r,2011}^{LandScan}}\right), & \text{if } r \in C \\ pop_{r,2011}^{LandScan}, & \text{if } r \notin C \end{cases} \tag{5.A.2}$$

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.

[10] The population data are accessed from the SSP database.[52]

5.B Spatial units for projection: "impact regions"

We create a set of custom boundaries that define the spatial units for which location-specific projected damages of climate change are computed. To do so, we utilize politically defined regions, as opposed to a regular grid, as socioeconomic data are generally collected at this scale and because administrative regions are relevant to policy-makers. These regions, hereafter referred to as "impact regions", are constructed such that they are identical to existing administrative regions or are a union of a small number of administrative regions. We use version 2 of the Global Administrative Region dataset (GADM),[62] 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 5.B.1, and we detail this agglomeration algorithm below.

Figure 5.B.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.

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:

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

where pop_c is population of country c in 2011 from Landscan (see Appendix 5.A.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 approximately 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 5.B.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 5.B.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 5.B.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 5.B.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 5.B.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:

- The set of GADM regions within the agglomeration

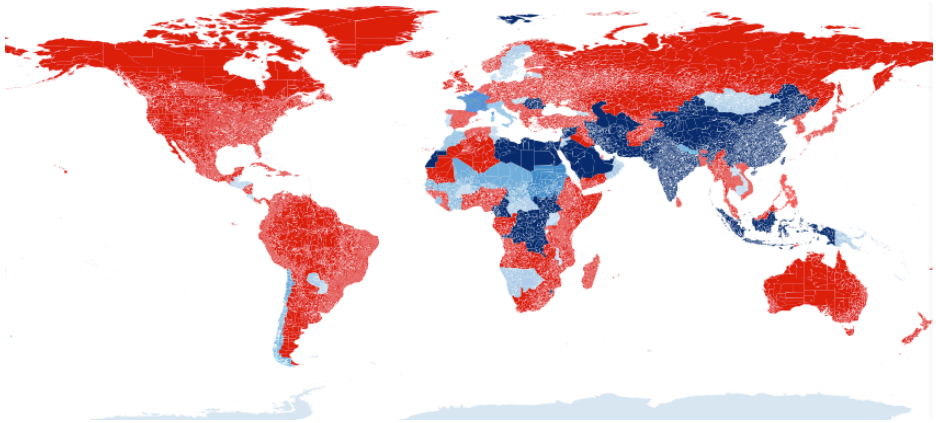

Figure 5.B.2: Categorization of countries based on the method used to construct impact regions out of GADM administrative regions. A country's target number of impact regions is $target_c$, as computed in the text. Countries in shades of blue have target values that can be approximated by one of the 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 neighboring agglomerated regions
- Population (pop),[11] 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.
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:

[11] Population data are from Landscan,[54] as in Appendix 5.A.3.3.

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.

5.C Estimating energy-temperature responses and projecting impacts of climate change

This appendix provides details on the estimation of energy-temperature responses and projection of climate change impacts. Section 5.C.1 demonstrates estimation of global average energy-temperature responses, while Section 5.C.2 introduces heterogeneity by income. Section 5.C.3 explains how we estimate and project energy-temperature responses as a function of both income and long-run climate. Finally, Section 5.C.5 details the procedure for characterizing econometric and climatological uncertainty in our estimates of the energy consumption impacts of climate change.

5.C.1 Global average energy-temperature response

Our energy-temperature responses flexibly model annual electricity and other fuels consumption each as a function of daily average temperatures within a year. Let E_{jtc} denote consumption in GJ per capita in country j, year t, and fuel category c (electricity, other fuels). The temperature and precipitation vectors $\boldsymbol{T}_{jt}$ and $\boldsymbol{P}_{jt}$ contain country-by-year aggregations of nonlinear grid-cell-level transformations of daily temperature and precipitation, respectively. These vectors thus summarize the full distribution of daily average weather in country j, year t (*Appendix*, 5.A.2.4). The basic econometric specification for a global average energy-temperature response is:

$$E_{jtc} = f_c(\boldsymbol{T}_{jt}) + g_c(\boldsymbol{P}_{jt}) + \alpha_{jic} + \delta_{wtc} + \varepsilon_{jtc}. \tag{5.C.1}$$

Our primary object of interest is the effect of temperature on energy consumption, represented by the response function $f_c(\cdot)$, which differs for electricity and other fuels. In our estimation of Equation 5.C.1, the vector $\boldsymbol{T}_{jt}$ contains polynomials of daily average temperatures (up to fourth order), each summed across the year, and the vector $\boldsymbol{P}_{jt}$ contains polynomials of daily cumulative precipitation (up to second order), also each summed annually. We estimate $f_c(\cdot)$ and $g_c(\cdot)$ as *linear* functions of the *nonlinear* elements of $\boldsymbol{T}_{jt}$ and $\boldsymbol{P}_{jt}$, respectively. This construction allows us to estimate a linear regression model while preserving the nonlinear relationship between energy consumption and temperature that takes place at the grid-cell-by-day level.[20]

The econometric specification includes a full set of country-by-reporting regime fixed effects (α_{jic}), where reporting regimes (i) are time spans within a country where observations for a given category are documented by the IEA to be comparable (*Appendix* 5.A.1). In addition, we include world region-by-year fixed effects (δ_{wtc}) for each category, where w indexes world regions based on UN classifications (Oceania, N. America, N. Europe, S. Europe, W. Europe, E. Europe, E. Asia, S.E. Asia, Central America/Caribbean, South America, sub-Saharan Africa, N. Africa/W. Asia, S. Asia). Finally ε_{jtc} denotes the stochastic error term; standard errors are clustered by country-category-reporting regime.

Due to evidence of unit root behavior in the dependent variable (*Appendix* 5.A.1), we estimate Equation 5.C.1 and all other regressions in first-differences. Additionally, to address differential data quality across reporting regimes, we employ an inverse variance weighting procedure in all regressions. In particular, we utilize Feasible Generalized Least Squares (FGLS) weights to downweight low credibility country-reporting regimes based

on their residual variance.[12] To implement this weighting, we first estimate the regression in first differences. Using the residuals from this regression, we calculate a country-category-reporting regime level weight equal to the inverse of the average value of the squared residuals, where the average is taken across all year observations that fall within a given country-category-reporting regime. We then apply these weights to the regression in a second stage. Observations from all years within a given country-category-reporting regime are given the same weight in the second stage; country-category-reporting regimes with higher average residual variance thus receive lower weight.

Formally, for Equation 5.C.1, each observation in country j, reporting regime i, category c receives a weight $w_{jic} = \frac{1}{Var(\Delta\varepsilon_{jtc \in jic})}$, and the weighted, first-differenced regression specification is:

$$w_{jic}\Big[\Delta E_{jtc}\Big] = w_{jic}\Big[\Delta f_c(\boldsymbol{T}_{jt}) + \Delta g_c(\boldsymbol{P}_{jt}) + \Delta\delta_{wtc} + \Delta\varepsilon_{jtc}\Big], \tag{5.C.2}$$

where Δ here denotes the first-difference operator. To operationalize Equation 5.C.2, we estimate w_{jic} with $\widehat{w_{jic}} = \frac{1}{\widehat{Var}(\widehat{\Delta\varepsilon_{jtc \in jic}})}$, where $\widehat{\Delta\varepsilon_{jtc}}$ denotes the residual for the country j, year t, category c observation from the unweighted, first-differenced regression, and $\widehat{Var}$ refers to the sample variance. Note that the term α_{jic} from Equation 5.C.1 drops out of Equation 5.C.2 due to first-differencing.

As discussed in *Methods*, the estimated electricity- or other fuels-temperature response is expressed as the difference between predicted consumption on a day where the average temperature is T and predicted consumption on a "mild" day with an average temperature of 20°C. The blue curve in Figure 5.C.1 displays the global, population-weighted electricity-temperature response estimated from Equation 5.C.2, while the orange curve displays the other-fuels temperature response.[13] Over the time period represented by the data (1971-2010), electricity consumption exhibits little sensitivity to temperature on average, globally speaking (blue curve, Figure 5.C.1), while other fuels consumption is seen to increase at cold temperatures (orange curve, Figure 5.C.1). The lack of an electricity-temperature response, particularly on hot days, reflects the fact that during the sample period, most of the global population is too poor to access electricity-intensive protective technologies such as air-conditioning. In contrast, consumption of other fuels is seen to increase on cold days, reflecting the use of these fuels for heating across a broader range of the global population.

5.C.2 Energy-temperature response heterogeneity by income

The global average energy responses to temperature shown in the previous section are likely to mask substantial heterogeneity by income, as consumption of electricity and other fuels is strongly correlated with wealth. To demonstrate how energy-temperature responses vary by income, we estimate responses for each decile of GDP per-capita as follows:

$$E_{jtc} = f_{cq}(\boldsymbol{T}_{jt}) + g_c(\boldsymbol{P}_{jt}) + \phi_{cq} + \alpha_{jic} + \delta_{wtc} + \varepsilon_{jtc}, \tag{5.C.3}$$

[12] We assume constant residual variance within each country-category-reporting regime.

[13] These are obtained by estimating Equation 5.C.2 with population weights in addition to inverse-variance weights. Population weighting assigns each country's observations a weight proportional to its average population over the years of the sample. Given that Equation 5.C.2 does not model heterogeneity in the energy-temperature response by income and long-run climate, population weights give us the response for the average global person.

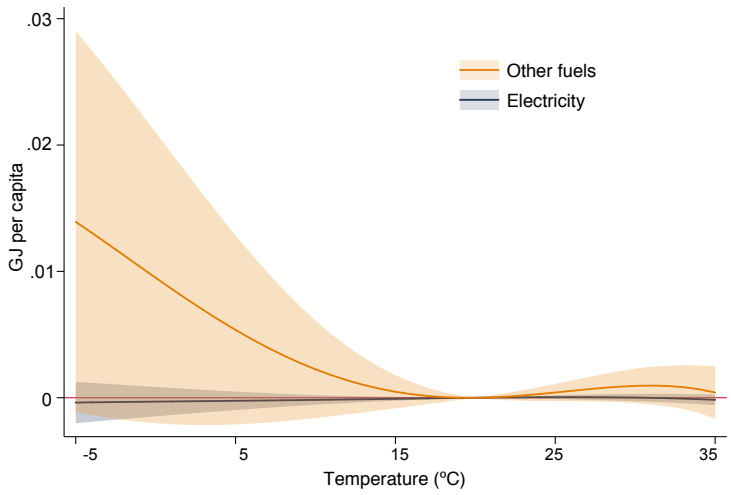

Figure 5.C.1: Population-weighted global average electricity-temperature response (blue) and other fuels-temperature response (orange). Each point on the curve denotes the additional per-capita daily consumption at a daily average temperature denoted on the horizontal axis, relative to a daily average temperature of 20°C. Shaded areas indicate 95% confidence intervals.

where q indexes the in-sample decile of the 15-year moving average GDP per-capita that country j in year t falls into, and ϕ_{cq} is a fuel category-specific decile fixed effect. As with Equation 5.C.1, we take into account nonlinearities in the temperature response, f_{cq}, but do so via a second-order polynomial due to the large number of additional parameters. Thus, the vector $\boldsymbol{T}_{jt}$ contains linear and quadratic terms of daily average temperatures, each summed across the year (*Appendix*, 5.A.2.4). The electricity- and other fuels-temperature responses in main text Figure 5.1A and 5.1B are obtained from estimating Equation 5.C.3 in first-differences using inverse variance weighting (as shown in Equation 5.C.2).

5.C.3 Energy-temperature response heterogeneity by income and long-run climate

All climate impact projections computed throughout our analysis incorporate heterogeneity in the energy-temperature response not only by income, but also by long-run average climate. To represent heterogeneity in the energy-temperature response along the dimensions of both income and long-run climate, we estimate Equation 5.1 (*Methods*) including flexible interactions between temperature and the following covariates:

1. **Income** measured as the 15-year moving average of a country's natural log of per-capita GDP ($\overline{LogGDPPC}$).
2. **Long-run climate** is measured by a country's average annual values of the following two variables over the time period of the sample:[14]
    - *Heating degree days* ($\overline{HDD}$) are defined as the cumulative deviations of daily average temperatures from a benchmark of 20° C, over all days where the

[14] Because long-run climate is slow moving, it is difficult to rely on temporal variation within the timespan of the sample. We therefore utilize only cross-sectional variation to characterize heterogeneity in the energy-temperature response due to long-run climate.

average temperature fell below 20° C. Formally, heating degree days in year t are defined as $\sum_{d \in t} |T_d - 20| * \boldsymbol{I}_{T_d<20}$, where T_d is the average temperature on day d and $\boldsymbol{I}_{T_d<20}$ is an indicator variable equal to one if $T_d < 20$.

- *Cooling degree days* ($\overline{CDD}$) are defined similarly over days where the daily average temperature exceeds 20° C, i.e. cooling degree days in year t are $\sum_{d \in t} |T_d - 20| * \boldsymbol{I}_{T_d>20}$.

These variables were carefully chosen based on the intersection of prior evidence from the literature, economic theory, and inclusion in standard projections of the global economy developed for integration with physical climate models.[18] Prior literature has emphasized the adaptive significance of average climate[28] and income per-capita.[25] For example, higher per-capita GDP entails greater capability to invest in protective measures such as air-conditioning, which can amplify the energy-temperature response, as demonstrated in main text Figure 5.1A. Furthermore, the energy-temperature response may also depend on long-run exposure to extreme temperatures as places with greater previous exposure may differ in their adaptive behaviors.[23]

We flexibly model the heterogeneity in the temperature response due to each of the covariates. The vector $\boldsymbol{T}_{jt}$ contains linear and quadratic terms of daily average temperatures, each summed across the year (*Appendix*, 5.A.2.4). To capture the role of income-driven adaptation, we interact a linear spline of $\overline{LogGDPPC}$ with each element of $\boldsymbol{T}_{jt}$. A knot point of the income spline is set separately for electricity and other fuels consumption, based on the point in the in-sample income distribution at which consumption starts to become responsive to temperature. Figure 5.1A (main text) suggests that this point is the beginning of the seventh decile for electricity ($\overline{GDPPC} = \$11,258$, 2019 USD PPP) and the beginning of the third decile for other fuels ($\overline{GDPPC} = \$2,849$, 2019 USD PPP). In addition to interactions of income with temperature exposure, we also flexibly control for the direct effect of income on energy consumption through a piecewise linear function of $\overline{LogGDPPC}$ that also allows for step-wise jumps at these points.

To capture the role of climate-driven adaptation, the long-run average climate measures $\overline{CDD}$ and $\overline{HDD}$ are interacted with the daily temperature vector $\boldsymbol{T}_{jt}$, but only over specific daily temperature ranges holding adaptive significance. Specifically, we allow the cooling degree day measure $\overline{CDD}$ to modulate the energy-temperature response on warm days, where the average temperature is at least 20° C. Conversely, the heating degree day measure $\overline{HDD}$ modulates the energy-temperature response on cool days, where the average temperature is below 20° C. Figure 5.C.2 provides graphical intuition for this split interaction on either side of 20° C. In order to model this type of interaction effect when energy consumption data are annual rather than daily, it is necessary to exploit information on the number of days in a year where the average temperature is $\geq$ 20° C and $<$ 20° C. For this purpose, we construct the variables $\mathbf{I}_{T_{jd} \geq 20}$ and $\mathbf{I}_{T_{jd}<20}$, which denote the share of country j's population experiencing an average temperature on day d that is $\geq$ 20° C and $<$ 20° C, respectively.[15] These variables are also included in the $\overline{CDD}$ and $\overline{HDD}$ interaction terms.

The form of Equation 5.1 (*Methods*) that we estimate thus specifies energy consumption in country j, year t, and category c as a function of temperature, and income and climate covariates, as follows:

[15] As with our other temperature exposure measures, we generate these variables at the country level by averaging across grid cells within a country using population weights (*Appendix* 5.A.2.4).

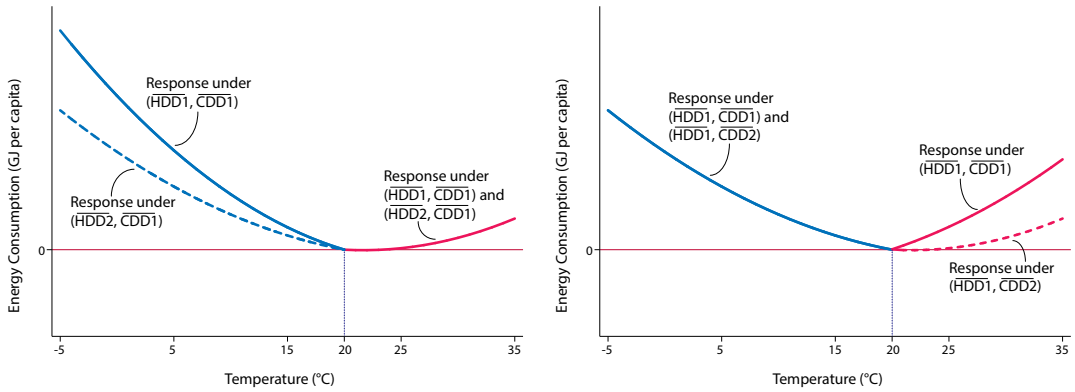

Figure 5.C.2: Modeling adaptation to long-run climate. The schematics below illustrate how we model adaptation to long-run climate separately for two sides of the energy-temperature response (below and above 20° C), through long-run average heating degree day (HDD) and cooling degree day (CDD) measures. The left panel illustrates hypothetical responses under two different long-run average HDD values ($\overline{HDD1}$ and $\overline{HDD2}$), holding fixed the long-run average CDD value at $\overline{CDD1}$. Changing the long-run average HDD value from $\overline{HDD1}$ to $\overline{HDD2}$ only alters the response to temperatures less than 20° C. Similarly, the right panel illustrates hypothetical responses under two different long-run average CDD values ($\overline{CDD1}$ and $\overline{CDD2}$), holding fixed the long-run average HDD value at $\overline{HDD1}$. Changing the long-run average CDD value from $\overline{CDD1}$ to $\overline{CDD2}$ only alters the response to temperatures $\geq$ 20° C.

$$
\begin{aligned}
E_{jtc} &= f_c(\boldsymbol{T}_{jt} \mid \overline{lGDPpc}_{jt}, \overline{CDD}_j, \overline{HDD}_j) + g_c(\boldsymbol{P}_{jt}) + \alpha_{jic} + \delta_{wtc} + \varepsilon_{jtc} \\
&= \underbrace{\boldsymbol{\beta}_c \cdot \boldsymbol{T}_{jt}}_{\text{Effect of temperature}} \\
&+ \underbrace{[\boldsymbol{\eta}_{1c} \cdot \boldsymbol{T}_{jt}](\bar{I}_c - \overline{lGDPpc}_{jt})\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c} + [\boldsymbol{\eta}_{2c} \cdot \boldsymbol{T}_{jt}](\overline{lGDPpc}_{jt} - \bar{I}_c)\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}}_{\text{Effect of income growth on energy-temperature response}} \\
&+ \underbrace{\sum_{k=1}^{2} \gamma_{kc}\overline{CDD}_j \sum_{d\in t}(T_{jd}^k - 20^k)\mathbf{I}_{T_{jd}\geq 20}}_{\substack{\text{Effect of climate adaptation on energy-temperature response} \\ \text{for days} \geq 20^\circ \text{ C}}} + \underbrace{\sum_{k=1}^{2} \lambda_{kc}\overline{HDD}_j \sum_{d\in t}(20^k - T_{jd}^k)\mathbf{I}_{T_{jd}<20}}_{\substack{\text{Effect of climate adaptation on energy-temperature response} \\ \text{for days} < 20^\circ \text{ C}}} \\
&+ \underbrace{\left[\kappa_{1c}\overline{lGDPpc}_{jt} + \phi_1\right]\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c} + \left[\kappa_{2c}\overline{lGDPpc}_{jt} + \phi_2\right]\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}}_{\text{Direct effect of income on energy consumption}} \\
&+ \underbrace{\boldsymbol{\theta}_c \cdot \boldsymbol{P}_{jt}}_{\text{Precipitation controls}} + \underbrace{\alpha_{jic} + \delta_{wtc}}_{\text{Fixed effects}} + \underbrace{\varepsilon_{jtc}}_{\text{Error Term}}, \qquad \text{(5.C.4)}
\end{aligned}
$$

where $\bar{I}_c$ denotes the income knot point for category c, and $\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c}$ and $\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}$ are indicator variables for whether $\overline{LogGDPPC}$ in country j, year t, is $<$ or $\geq$ the category c income knot point. We estimate this model in first-differences using inverse variance weighting (*Appendix* 5.A.1 and 5.C.1). The matrices in main text Figure 5.1C summarize the results, with each cell in a matrix displaying a predicted energy-temperature response evaluated at a particular point in the income $\times$ long-run climate space within the estimation sample. The cells are ordered vertically by $\overline{LogGDPPC}$ terciles (increasingly rich from bottom to top) and horizontally by $\overline{CDD}$ terciles (increasingly warm climate

from left to right). The predicted energy- temperature response function in each cell is evaluated at the mean values of $\overline{LogGDPPC}$, $\overline{CDD}$, and $\overline{HDD}$ within their respective terciles.

Extrapolating energy-temperature responses spatially and temporally

We use the estimated parameters from Equation 5.C.4 to extrapolate energy-temperature responses across locations over time based on projected future incomes and climate (this procedure is detailed in *Methods*). Figure 5.C.3 demonstrates the overlap in the joint income-climate distributions at 2010 and 2090. Although the future distribution is shifted towards higher incomes, greater cooling degree days, and fewer heating degree days, the substantial overlap in the two distributions allows for credible extrapolation of energy-temperature responses into the future.

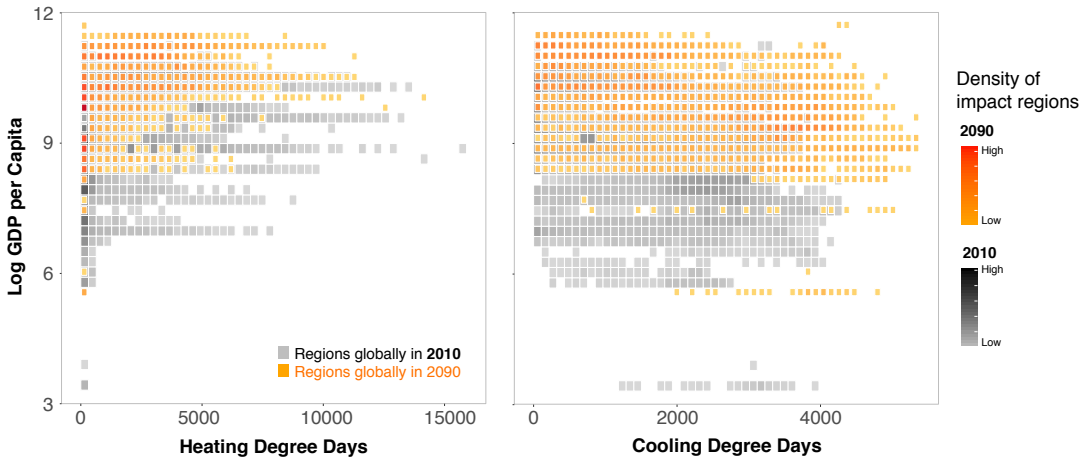

Figure 5.C.3: Sample overlap between present and future. The density plots below depict joint distributions of income and long-run climate as measured by heating degree days and cooling degree days. Distributions are for 24,378 impact regions, in 2010 (grey-black) and 2090 under the RCP8.5 emissions scenario and SSP3 socioeconomic scenario (red-orange).

5.C.4 Projecting energy consumption impacts of climate change

To estimate future energy consumption impacts for each fuel category c and impact region r for each year from 2015 to 2099, we apply a set of probabilistic climate change projections to the spatially and temporally heterogenous energy-temperature responses described above (and detailed in *Methods*). The distribution of future daily average temperatures under a given emissions scenario (RCP8.5 or RCP4.5) is obtained from the 33 projections in the SMME (*Appendix* 5.A.2.3).

Let $\boldsymbol{T}_{rt}$ represent a vector containing impact region-by-year aggregations of nonlinear grid-cell-level transformations of daily temperature in a future year t, under a warmer climate. In contrast, let $\boldsymbol{T}_{r2015}$ represent the counterfactual temperature vector for the same impact region under a climate that is the same as that of 2015. These vectors are constructed in exactly the same way as is done for the temperature vectors used in estimating Equation 5.C.4 (*Appendix* 5.A.2.4).

The impact of climate change on consumption in fuel category c is expressed as the estimated change in consumption relative to a no-climate change counterfactual in which

the future climate is the same as in 2015. The precise form of Equation 5.2 (*Methods*) with which we project impacts is thus:

$$
\begin{aligned}
ImpactOfClimateChange_{crt} = \; & \underbrace{\widehat{f}_c(\boldsymbol{T}_{rt} \mid \overline{lGDPpcpc)}_{rt}, \overline{CDD}_{rt}, \overline{HDD}_{rt})}_{\substack{\text{Temperature-induced energy consumption under climate change} \\ \text{(with income growth and climate-driven adaptation)}}} \\
& - \underbrace{\widehat{f}_c(\boldsymbol{T}_{r2015} \mid \overline{lGDPpc}_{rt}, \overline{CDD}_{r2015}, \overline{HDD}_{r2015})}_{\substack{\text{Temperature-induced energy consumption without climate change} \\ \text{(with income growth)}}} \\
= \; & \Big[\widehat{\boldsymbol{\beta}_c} \cdot \boldsymbol{T}_{rt} \\
& + [\widehat{\boldsymbol{\eta}_{1c}} \cdot \boldsymbol{T}_{rt}](\bar{I}_c - \overline{lGDPpc}_{rt})\mathbf{I}_{\overline{lGDPpc}_{rt} < \bar{I}_c} + [\widehat{\boldsymbol{\eta}_{2c}} \cdot \boldsymbol{T}_{rt}](\overline{lGDPpc}_{rt} - \bar{I}_c)\mathbf{I}_{\overline{lGDPpc}_{rt} \geq \bar{I}_c} \\
& + \underbrace{\sum_{k=1}^{2} \widehat{\gamma_{kc}} \overline{CDD}_{rt} \sum_{d \in t} (T_{rd}^k - 20^k)\mathbf{I}_{T_{rd} \geq 20} + \sum_{k=1}^{2} \widehat{\lambda_{kc}} \overline{HDD}_{rt} \sum_{d \in t} (20^k - T_{rd}^k)\mathbf{I}_{T_{rd} < 20}\Big]}_{\substack{\text{Temperature-induced energy consumption under climate change} \\ \text{(with income growth and climate-driven adaptation)}}} \\
& - \Big[\widehat{\boldsymbol{\beta}_c} \cdot \boldsymbol{T}_{r2015} \\
& + [\widehat{\boldsymbol{\eta}_{1c}} \cdot \boldsymbol{T}_{r2015}](\bar{I}_c - \overline{lGDPpc}_{rt})\mathbf{I}_{\overline{lGDPpc}_{rt} < \bar{I}_c} \\
& \qquad + [\widehat{\boldsymbol{\eta}_{2c}} \cdot \boldsymbol{T}_{r2015}](\overline{lGDPpc}_{rt} - \bar{I}_c)\mathbf{I}_{\overline{lGDPpc}_{rt} \geq \bar{I}_c} \qquad \text{(5.C.5)} \\
& + \underbrace{\sum_{k=1}^{2} \widehat{\gamma_{kc}} \overline{CDD}_{r2015} \sum_{d \in 2015} (T_{rd}^k - 20^k)\mathbf{I}_{T_{rd} \geq 20} + \sum_{k=1}^{2} \widehat{\lambda_{kc}} \overline{HDD}_{r2015} \sum_{d \in 2015} (20^k - T_{rd}^k)\mathbf{I}_{T_{rd} < 20}\Big]}_{\substack{\text{Temperature-induced energy consumption without climate change} \\ \text{(with income growth)}}}
\end{aligned}
$$

(5.C.6)

where the $\widehat{\ }$ denotes estimated objects from Equation 5.C.4, and $\overline{lGDPpc}_{rt}$, $\overline{CDD}_{rt}$, and $\overline{HDD}_{rt}$ denote 15-year moving averages of the covariates for impact region r in future year t. The object $ImpactOfClimateChange_{crt}$ represents the change in annual per-capita electricity or other fuels consumption due to a shift in the temperature distribution under climate change, accounting for the evolution of energy-temperature responses as locations warm and incomes rise. It isolates the *additional* impact of climate change net of other factors (e.g. income) that will change in the future. The no-climate change counterfactual is constructed to be identical in every way, except for the climate, and is therefore evaluated at the future level of income, but with $\overline{CDD}$ and $\overline{HDD}$ values drawn from the 2015 climate distribution.

We construct estimates of Equation 5.C.5 for all impact regions up to 2099 under emissions scenarios RCP8.5 and RCP4.5, using each of the 33 climate projections in the SMME (Figure 5.2 in the main text). In addition, to highlight the critical importance of income growth and climate-driven adaptation in shaping future energy-temperature responses, we also display a "no-adaptation" impact projection in which energy-temperature responses $\widehat{f}_c(\cdot)$ are determined solely by an impact region's income, $\overline{CDD}$, and $\overline{HDD}$ values in the year 2015 (*Methods* Equation 5.3).

5.C.5 Accounting for uncertainty in projected energy consumption impacts of climate change

An important feature of our analysis is to develop estimates of the energy consumption impacts of climate change that reflect the uncertainty inherent in these future projections. This uncertainty arises from two distinct sources– climatological and econometric. The confidence intervals displayed in Figure 5.2C in the main text, as well as the kernel density plots in Figure 5.3A and confidence intervals in Figure 5.3B, represent the combined uncertainty from both these sources.

To account for uncertainty in the climate system, we construct estimates of the energy consumption impacts of future climate change (Equation 5.2 in *Methods*) for each of 33 distinct climate projections in the surrogate model mixed ensemble.[16] Distributions of climate change impacts are provided by weighting across the 33 projections using the weights listed in Table 5.A.1. Furthermore, independent of physical uncertainty, an important second source of uncertainty arises from the econometric estimation of Equation 5.C.4. To account for this econometric uncertainty, we apply the delta method[55] to characterize the Gaussian distribution of impacts (*Methods* Equation 5.2) under each of the 33 climate projections. Finally, to characterize the full distribution of impacts across both climate and econometric uncertainty, we construct the mixture distribution of these 33 Gaussian distributions[17] using Newton's method.[63]

The confidence intervals for the global impact time series shown in Figure 5.2C and Figure 5.3B in the main text are derived from quantiles of such mixture distributions. Importantly, the same method to characterize uncertainty can also be applied to impacts for individual impact regions, as illustrated by the kernel density plots for selected impact regions in Figure 5.3A.

[16] Note that while the surrogate model mixed ensemble fully represents the tails of the climate sensitivity distribution (*Methods*; *Appendix* 5.A.2.3), there remain important sources of climate uncertainty that are not captured in our projections. These include some climate feedbacks that may amplify the increase of global mean surface temperature, as well as some factors affecting local climate that are poorly simulated by GCMs.

[17] Probability weights for each of the 33 climate projections are used in the construction of the mixture distribution (*Appendix* 5.A.2.3).[36]

5.D Valuing impacts

To monetize the projected impacts of climate change on energy consumption, we apply geographically-specific real prices for electricity and other fuels to the projected quantity impacts constructed from Equation 5.C.5, thus reflecting differential costs across geographies and fuels. We consider a range of future price scenarios, constructed either from direct extrapolation of present-day price statistics or from price projections generated by integrated assessment models (IAMs).

5.D.1 Extrapolating present-day prices

For scenarios based on direct extrapolation (Panels I and II in Table 5.4E in the main text), impacts on electricity consumption are valued using an average cost of electricity generation, which the IEA's *World Energy Outlook 2017* provides globally as of 2016 at the country or world region level.[18] Impacts on other fuels consumption are valued using residential and non-residential end-user prices excluding taxes, which the IEA's *Energy Prices and Taxes* dataset provides for coal, oil, and natural gas fuels in 55 countries as of 2012.[19] Countries that lack price data for a given fuel are assigned the global average price for that fuel. To obtain a price for the pooled, multi-fuel "other fuels" category, we weight the prices of the individual fuels according to their shares in a country's overall "other fuels" consumption as of 2012 (the most recent year for which consumption data are available).[20] Thus, each country receives unique prices at which its impacts on other fuels consumption are valued.[21] To extrapolate prices into the future, we consider three annual price growth trajectories: a moderate growth trajectory of 1.4% annual price growth (equal to the historical growth of US real energy prices[22]), a stagnant trajectory of 0% annual price growth, and a high growth trajectory of 3% annual price growth.[23]

5.D.2 Price projections from IAMs

As a complement to price scenarios extrapolated from present-day prices, we also consider price scenarios based on projections of multiple IAMs. IAMs simulate the entire energy system and determine energy prices by equilibrating supply and demand. We obtain price projections of IAMs from IIASA's Scenario Explorer database.[24] This database contains output from 416 IAM $\times$ scenario combinations, of which 155 include price projections for electricity and other fuels as part of their output. To monetize impacts under RCP8.5, we limit ourselves to scenarios that do not posit any policies to mitigate greenhouse gas

[18] Costs are specified for the following geographies: Japan, European Union, Korea, Brazil, Australia, Mexico, Southeast Asia, Middle East, India, Africa, United States, China, Canada, Russia. When a cost is not available specific to a particular geography we extend these costs based on UN world region classifications: Oceania receives the Australia cost, N., S., and W. Europe receive the EU cost, E. Europe receives the Russia cost, Central America/Caribbean receive the Mexico cost, S. America receives the Brazil cost, N. Africa receives the Middle East cost, and S. Asia receives the India cost.

[19] We take a weighted average of residential and non-residential prices, with a weight of 16% on residential and 84% on non-residential. These weights are determined based on the average share of consumption in these two sectors in the set of 55 countries where a sectoral breakdown is available.

[20] Although our consumption data do cover fuels besides coal, oil, and natural gas (including solar, geothermal, and biofuels), no price data are available for fuels other than coal, oil, and natural gas. We therefore extend the weighted average price to the remaining fuels.

[21] Even countries that are assigned the global average price for the individual fuels will differ in the shares of each fuel in their other fuels consumption mix.

[22] Price data from 1970-2015 taken from the *State Energy Data System* of the U.S. Energy Information Administration.

[23] These trajectories apply globally to prices for both electricity and other fuels.

[24] https://data.ene.iiasa.ac.at/iamc-1.5c-explorer/#/workspaces.

emissions (e.g. carbon price). There are 5 IAM × scenario combinations that satisfy this requirement. To monetize impacts under RCP4.5, we identify a counterpart scenario for each of these combinations that is broadly consistent with RCP4.5. Table 5.D.1 lists the IAM × scenario combinations whose price projections we utilize for monetizing impacts under RCP8.5 and RCP4.5.

IAM	Scenario for RCP8.5	Scenario for RCP4.5
WITCH-GLOBIOM 4.2	ADVANCE_NoPolicy	ADVANCE_2030_WB2C
REMIND 1.7	ADVANCE_NoPolicy	ADVANCE_2030_WB2C
REMIND 1.7	CEMICS-Ref	CEMICS-2.0-CDR8
MERGE-ETL 6.0	BAU	DAC2_66
REMIND-MAgPIE 1.7-3.0	SMP_REF_Def	SMP_2C_Def

Table 5.D.1: IAMs and scenarios used for monetizing impacts under RCP8.5 and RCP4.5. This table lists IAMs and associated scenarios whose price projections we use to monetize the projected impacts of climate change on energy consumption. IAMs and scenarios were selected from IIASA's Scenario Explorer database,[46] based on availability of future price projections and suitability for RCP8.5 and RCP4.5 emissions trajectories.

Every IAM × scenario reports electricity, oil, coal, natural gas, and biofuels prices for 6 world regions up to 2100.[25] To obtain a price for the "other fuels" category, we weight the prices of the individual fuels according to their shares in a region's overall "other fuels" consumption in each year, where the consumption levels for years up to 2100 are reported by the IAM × scenario.

The right panel of Figure 5.D.1 presents the time series of total global monetized impacts under RCP8.5 assuming various price trajectories, while the left panel does the same for RCP4.5. Regardless of the emissions scenario or assumed price trajectory, end-of-century damages (i.e. net savings) represent a minute fraction of the US $353 trillion end-of-century global GDP projected under SSP3.

[25]The 6 regions are OECD 1990 and EU, Eastern Europe and former Soviet Union, Middle East and Africa, Latin America, Asia, and Rest of the World. Prices are available at 5-year intervals and are linearly interpolated for the intervening years.

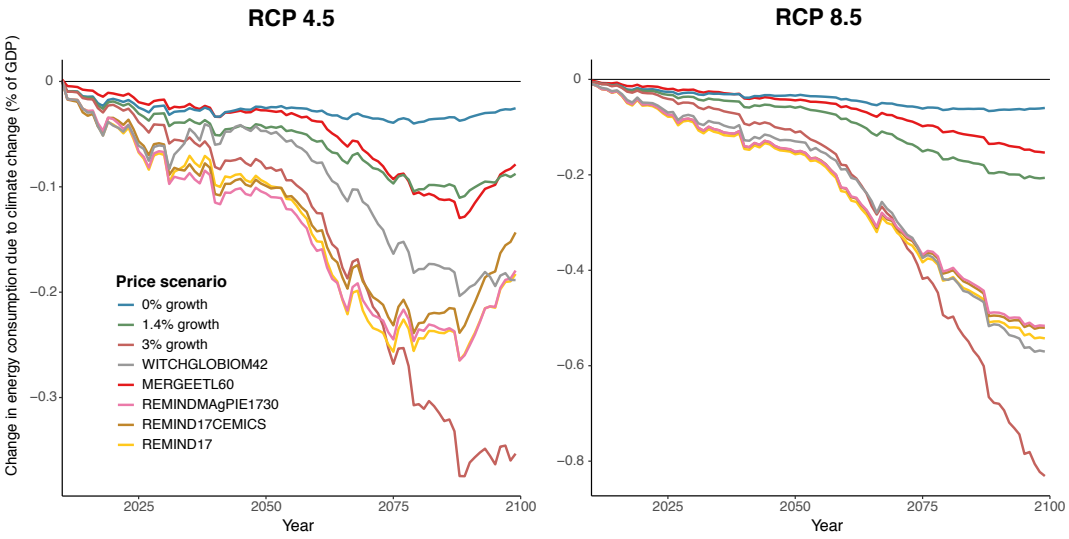

Figure 5.D.1: Valuing the impacts of climate change on energy consumption. Figure 5.D.1 displays the time series of total global energy consumption damages under the SSP3 socioeconomic scenario for high (RCP8.5) and moderate (RCP4.5) emissions scenarios, assuming various price trajectories. Damages in a given year are expressed as a percent of global GDP in that year. Aggregate global damages are obtained by monetizing and summing over the spatially disaggregated impacts across both electricity and other fuels.

5.E Damage function estimation

As described in *Methods*, we pool empirical estimates of climate change impacts constructed using Equation 5.C.5 to fit global energy damage functions. These damage functions express global energy consumption costs of climate change as a function of the change in global mean surface temperature relative to the 2001-2010 average level (ΔGMST).[1] Summarizing the economic costs of all impacts measured in the detailed empirical analysis, these functions can be differentiated everywhere, allowing for marginal costs of a CO_2 impulse to be computed for any global climate trajectory. Due to differences in the availability of climate and socioeconomic projections pre- and post-2100, there are some important differences in our approach for calculation of damage functions before and after 2100.

Computing damage functions through 2100 As detailed in *Methods*, we estimate time-varying damage functions for all years t prior to 2100 directly from the high-resolution climate change impact projections described in Appendix 5.C.4 and converted to dollar values as described in Appendix 5.D. To construct a damage function for year t, we pool all simulated damage estimates (globally summed across 24,378 impact regions and converted to dollar value) within a 5-year window of t and estimate the quadratic damage function shown in Equation 5.4 (*Methods*). Because the underlying population distribution and level of per capita income are evolving over time, the shape of our estimated damage functions change throughout the 21^{st} century (see Figure 5.E.1 below).

Computing damage functions after 2100 Data limitations in climate and socioeconomic projections beyond end-of-century necessitate an alternative approach to estimating post-2100 damage functions. Only 6 of the 21 GCMs that we use to build our surrogate model mixed ensemble 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 (*Appendix*, 5.A.2.2). Furthermore, the SSPs needed to project future incomes and demographics also end in 2100. Although one approach is to simply end economic cost calculations in 2100,[37] neglecting post-2100 damages is a substantial omission as a large fraction of costs, in net present value, are thought to occur after 2100 at 3% discount rates.[64]

To estimate post 2100-damages, we follow the method in ref.[23] to extrapolate changes in the damage function beyond 2100 using the observed evolution of damages in the second half of the 21^{st} century. To implement this extrapolation, we pool values D_{tlps} from 2085-2099 and estimate a quadratic model similar to Equation 5.4 (*Methods*), but interacting each term linearly with year t.[26] This allows us to estimate a damage surface as a parametric function of time. We then predict extrapolated damage functions for all years after 2100, smoothly transitioning from our flexible climate model-based damage functions prior to 2100.[27]

Figure 5.E.1 illustrates damage functions every 10 years prior to 2100, as well as extrapolated damage functions for the years 2150, 2200, 2250, and 2300. Extrapolated damages continue to become more steeply negative post-2100, as they did pre-2100.

[26]The specific interaction model we estimate is: $D(\Delta GMST, t)_{tlps} = \nu_0 + \nu_1 \Delta GMST_{tlp} \times t + \nu_2 \Delta GMST^2_{tlp} \times t + \varepsilon_{tlps}$.

[27]Furthemore, we also separately estimate a time-interacted regression for each of 19 quantiles (i.e. every 5^{th} percentile from the 5^{th} to 95^{th}), which we use to predict these quantiles of the damage functions for years after 2100. See *Methods* for details.

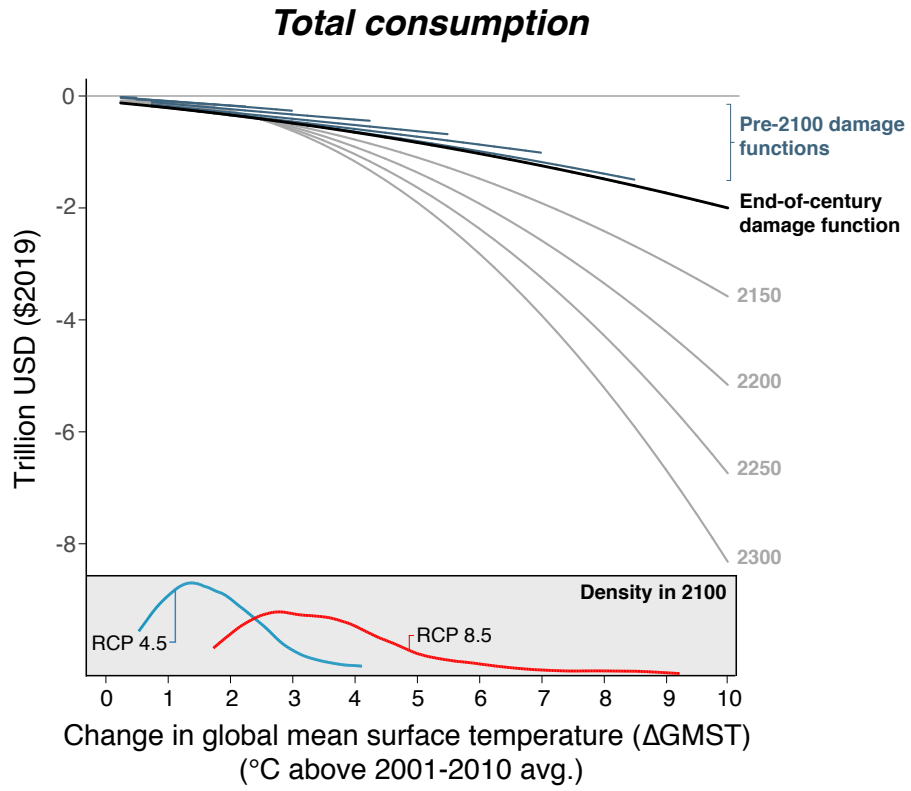

Figure 5.E.1: Empirically-derived damage functions by year. This figure shows damage functions relating empirically-derived total global energy consumption damages to anomalies in global mean surface temperature (ΔGMST). Energy consumption impacts of climate change used to estimate these damage functions are valued under a moderate price growth trajectory of 1.4% annual price growth. The bottom panel displays the distribution of ΔGMST at end of century under two emissions scenarios across the 33 climate projections in the SMME. In the top panel, the black line reproduces the end-of-century damage function from rightmost panel of Figure 5.3C in the main text. Additional damage functions are shown in blue for every 10 years pre-2100, each of which is estimated analogously to the end-of-century damage function, and in grey for every 50 years post-2100, each of which is extrapolated. Our projection results generate energy consumption damages only up to 2100, due to limited availability of climate and socioeconomic projections for years beyond that date. To capture impacts after 2100, we extrapolate observed changes in damages over the 21st century to generate time-varying damage functions through 2300.

5.F Calculation of an energy consumption partial social cost of carbon using a simple climate model

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

5.F.1 Set up of the climate module using a simple climate model

A core component of any analysis of the SCC is the climate module used to estimate both the baseline climate and the response of the climate system to a marginal change in greenhouse gas emissions. The Finite Amplitude Impulse Response (FAIR) model[22] satisfies key criteria for such a module, including those outlined by the National Academies of Sciences, Engineering, and Medicine.[3] In particular, the National Academies 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" [3, 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:

$$\frac{dT_i}{dt} = \frac{q_i F - T_i}{d_i}, i \in \{0, 1\}, \tag{5.F.1}$$

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 frac-

tions, R_j, representing different uptake timescales:

$$\frac{dR_j}{dt} = a_j E - \frac{R_j}{\alpha_j \tau_j}, j \in \{0, 1, 2, 3\} \tag{5.F.2}$$

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.[22]

We obtain the latest release of the FAIR model, which was version 1.3.2 at the time of computation, from its online repository.[28] As described below in Section 5.F.2, we develop a methodology to generate energy 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,[3] 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 5.F.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 5.4A-D; as described below and in *Methods*, this difference is used to construct energy partial SCC estimates.

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

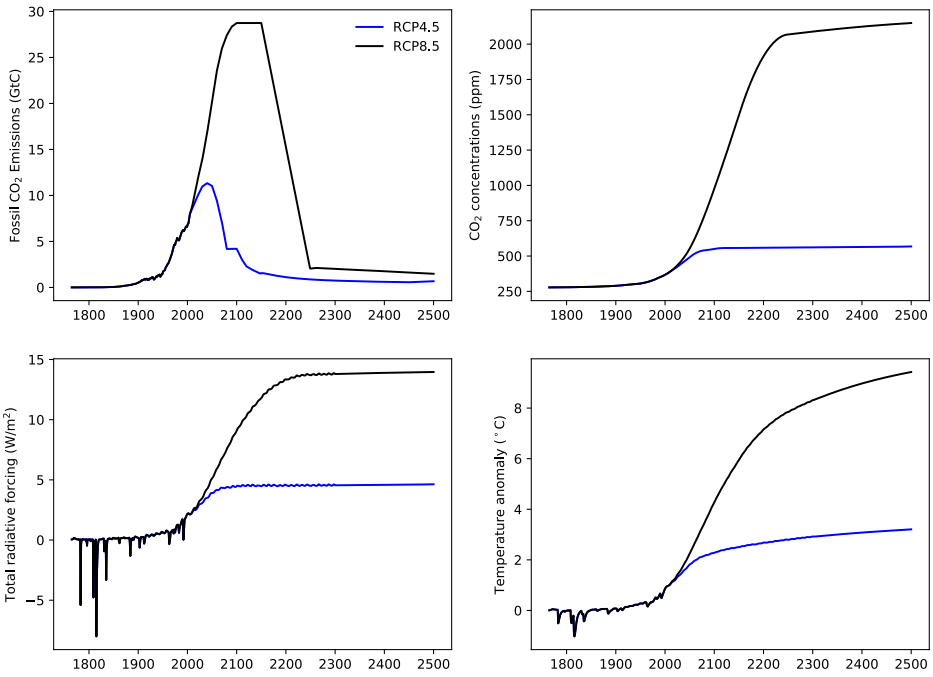

Figure 5.F.1: Behavior of key variables in the FAIR simple climate model under median climate parameters: Each panel shows the temporal trajectory of key variables in FAIR that are used in our calculation of the social cost of carbon. The trajectories shown arise under FAIR run with median climate parameter 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.

5.F.2 Methodology for capturing uncertainty in climate sensitivity within the simple climate model FAIR

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. However, a complete study of the energy partial SCC should represent the uncertainty in key model parameters, including the joint probability distribution of the ECS and TCR. We now discuss the development of such 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);[65] for d_2 and τ_3 we follow ref.,[22] based on refs.[66] and.[67]

In general, we produce initial distributions of these parameters based on constraints from the literature. 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[66,68] and those using the CMIP5 5–95% range of TCR and ECS as 5-95% input ranges to their models[22] 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 ref.[69] 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:** Ref.[65] concludes 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,[70] *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 a 17th to 83rd percentile range of 1.0-2.5 °C.

- **RWF:** An RWF likely range of 0.45 to 0.75 is approximately consistent with the ECS likely range of 1.5 – 4.5°C.[65] 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:** Ref.[65] concludes 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 a 17^{th}-83^{rd} percentile range of 1.60-4.65 °C.

- $\boldsymbol{d_2}$**:** The AR5 does not assess the range of d_2. 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.[22]

- $\boldsymbol{\tau_3}$**:** Ref.[66] 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 ref.[66] 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 5.F.1 and Figure 5.F.2.

Initial pulse-adjustment timescale (IPT) The National Academies highlights the IPT as a measure that is important for SCC computations, yet does not provide a clear, consistent definition.[3] 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." [3, 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 ref.[66] and other subsequent analysis[68] indicate that a peak in warming in response to a pulse emission occurs within about a decade after emission. In particular, ref.[68] 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. Ref.,[68] 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 ref.,[68] 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). TCRE is between 0.8 and 2.5°C per 1000 GtC with at least 66% probability.[65] 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 Coupled climate-carbon cycle experiments indicate that a majority (about 70% in the multimodel mean) of peak warming persists 500 years after emissions.[66] 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 5.F.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 5.F.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.[3] 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 5.4E in the main text.

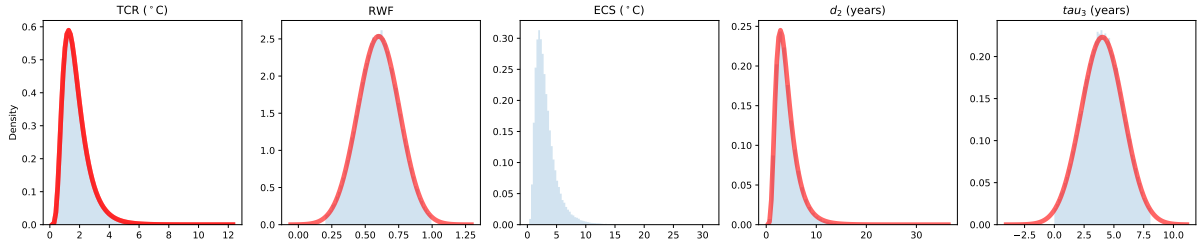

Figure 5.F.2: Distributions of key FAIR parameters for climate sensitivity uncertainty both before (red curve) and after (blue shading) applying constraints: Each panel indicates the distribution of a key parameter in the FAIR simple climate model, both before (in red) and after (in blue) the imposition of constraints described in the text. Distributions shown from left to right are: transient climate response (TCR); realized warming fraction (RWF) used to define ECS (=TCR / RWF); equilibrium climate sensitivity (ECS) shown only after applying constraints due to unrealistic values in the initial distribution occurring as RWF→ 0; short thermal adjustment time (d_2); time scale of rapid carbon uptake by the ocean mixed layer (τ_3).

Parameter	*Distribution from literature*	*Pre-IPT distribution*	*Post-IPT distribution*	Distribution	*Source*
TCR (C)	[1.00, 2.50]	[1.00, 2.49]	[1.00, 2.50]	Lognormal	AR5
RWF	[0.45, 0.75]	[0.45, 0.75]	N/A	Normal	NAS (2017)
ECS (C)	[1.5, 4.5]	[1.60, 4.65]	[1.61, 4.61]	Lognormal	AR5
d_2 (years)	(1.6, 8.4)	(1.6, 8.4)	(1.6, 8.3)	Lognormal	[22]
τ_3 (years)	[66]	4.04 (1.07, 6.96)	4.04 (1.25, 6.79)	Normal	[66]
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	[68]
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 5.F.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.[68]

Finally, we assess the reasonableness of the "handoff" between the SMME models, with which the damage function is estimated (Equation 5.4 in *Methods*), 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 energy damages derives from the SMME, while the uncertainty we proliferate through to the SCC relies on the simple climate model FAIR. Figure 5.F.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 row). SMME data are available until the year 2100; thus, the 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 would be expected based on the construction of the SMME).

5.F.3 Converting temperature scenarios to an energy consumption partial SCC

We convert the temperature scenarios developed in the climate module into estimates of energy-related damages using the global damage functions described in *Methods* and in Appendix 5.E. These damage functions characterize valued energy consumption impacts as a function of ΔGMST (changes in GMST relative to 2001-2010). The coefficients on these quadratic damage functions are constructed for each year from 2015 to 2300, as described in *Methods* and Appendix 5.E. We then generate annual estimates of temperature-related energy 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 energy damages associated with each scenario, we subtract each control scenario from the corresponding pulse scenario and divide by the pulse amount to estimate the marginal effect of the pulse. This time series is then discounted using 2.5%, 3% and 5% discount rates, and summed through time to create a net present value, following Equation 5.5 in *Methods*. This final value is the net present value of the full energy consumption impacts caused by a marginal emission of CO_2. A more robust estimate would make use of Ramsey-like discounting, accounting for the relationship between consumption growth and the discount rate, but we leave this for future study.

5.F.4 Uncertainty in the energy partial SCC

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

Energy partial SCC estimates accounting for both climate sensitivity and damage function uncertainty As described in *Methods*, damage functions are computed using estimates of the global monetized damages in each year generated from 33 climate models, two emissions scenarios, and a resampling of damage estimates that captures uncertainty in the estimation of Equation 5.1 (*Methods*). These multiple sim-

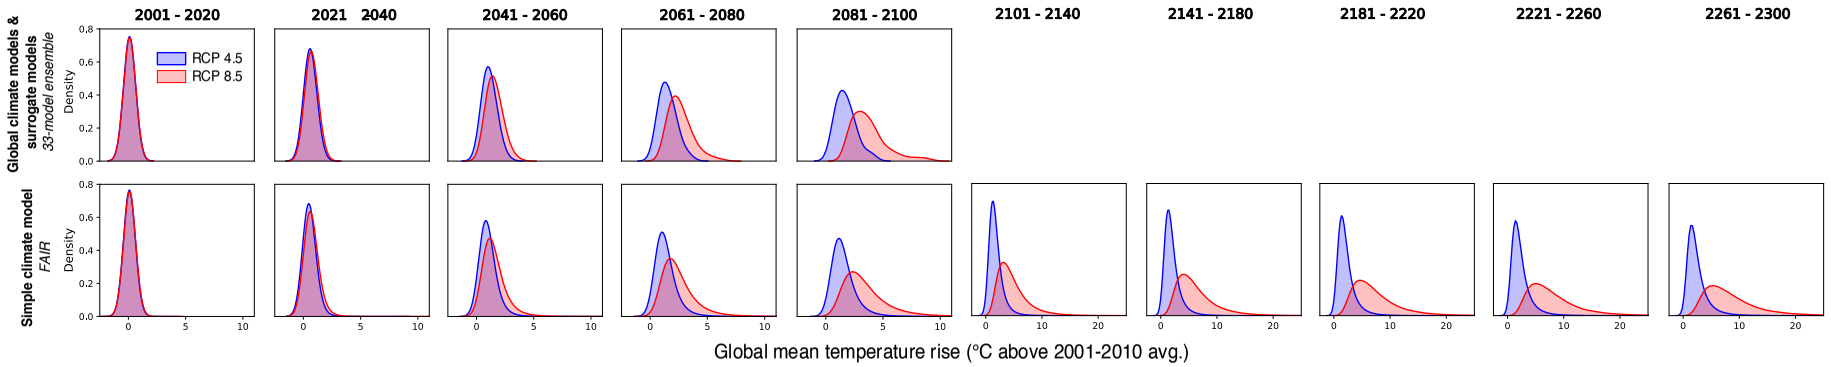

Figure 5.F.3: Distribution of changes in global mean surface temperature (ΔGMST) from an ensemble of global climate models and surrogate models (SMME) and from the simple climate model FAIR: Top row: Distribution of ΔGMST from 2001 to 2100, according to an ensemble of 33 GCMs and surrogate models that 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.

ulations (we draw 100 realizations of global damages for each climate model, emissions scenario, SSP trajectory, and year) give us an empirically-derived distribution of potential economic outcomes that are conditional on the ΔGMST value for the year, emissions scenario, and climate model used to generate that projection. To account for uncertainty in a single year's damage function, we pool these realizations for the associated 5-year window (see *Methods* and Appendix 5.E). 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 (*Appendix* 5.E). In this extrapolation, we allow each quantile 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 5^{th}-95^{th} percentile ranges are resolved from the resulting distribution of energy partial SCCs.

Energy partial SCC estimates accounting for climate sensitivity uncertainty only: To isolate uncertainty in the energy partial SCC that derives from climate sensitivity uncertainty, we run the mean damage function through each of the 96,306 sets of FAIR parameters. The corresponding SCC 5^{th}-95^{th} percentile range is resolved from the resulting distribution of energy partial SCCs.

Energy partial SCC estimates accounting for damage function uncertainty only: To isolate uncertainty in the energy 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 energy partial SCC estimates) and up-weight runs in order to reflect probability mass in the damage function uncertainty space. The corresponding SCC 5^{th}-95^{th} percentile range is resolved from the resulting distribution of energy partial SCCs.

Table 5.F.2 reproduces SCC estimates from Table 5.4E (main text), with interquartile ranges separately showing the influence of climate sensitivity and damage function uncertainty.

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
	I: 1.4% price growth		
RCP 8.5	-1.51	-1.16	-0.60
Climate sensitivity uncertainty	[-2.57,-0.96]	[-1.92,-0.74]	[-0.97,-0.38]
Damage function uncertainty	[-2.09,-0.97]	[-1.48,-0.80]	[-0.70,-0.51]
Full uncertainty	[-2.58,-0.66]	[-1.92,-0.55]	[-0.97,-0.35]
RCP 4.5	-1.37	-1.08	-0.58
Climate sensitivity uncertainty	[-2.24,-0.89]	[-1.74,-0.71]	[-0.91,-0.38]
Damage function uncertainty	[-1.96,-1.06]	[-1.44,-0.84]	[-0.65,-0.50]
Full uncertainty	[-2.45,-0.73]	[-1.83,-0.61]	[-0.90,-0.35]
	II: 0% price growth		
RCP 8.5	-0.72	-0.61	-0.39
Climate sensitivity uncertainty	[-1.19,-0.46]	[-0.99,-0.39]	[-0.62,-0.25]
Damage function uncertainty	[-0.87,-0.52]	[-0.72,-0.48]	[-0.43,-0.34]
Full uncertainty	[-1.12,-0.36]	[-0.95,-0.34]	[-0.62,-0.23]
RCP 4.5	-0.66	-0.57	-0.37
Climate sensitivity uncertainty	[-1.05,-0.43]	[-0.90,-0.38]	[-0.58,-0.24]
Damage function uncertainty	[-0.85,-0.55]	[-0.69,-0.49]	[-0.40,-0.33]
Full uncertainty	[-1.08,-0.39]	[-0.91,-0.35]	[-0.58,-0.23]
	III: MERGE-ETL 6.0 prices		
RCP 8.5	-1.12	-0.82	-0.39
Climate sensitivity uncertainty	[-1.82,-0.73]	[-1.32,-0.54]	[-0.63,-0.25]
Damage function uncertainty	[-1.43,-0.84]	[-0.99,-0.61]	[-0.45,-0.33]
Full uncertainty	[-1.66,-0.61]	[-1.26,-0.47]	[-0.63,-0.24]
RCP 4.5	-1.15	-0.83	-0.38
Climate sensitivity uncertainty	[-1.83,-0.77]	[-1.31,-0.56]	[-0.60,-0.25]
Damage function uncertainty	[-1.47,-1.04]	[-1.03,-0.72]	[-0.44,-0.33]
Full uncertainty	[-1.90,-0.76]	[-1.36,-0.54]	[-0.61,-0.24]

Table 5.F.2: Social cost of energy consumption due to climate change. This table displays estimates of a social cost of carbon for excess energy consumption costs under high (RCP8.5) and moderate (RCP4.5) emissions scenarios, assuming various discount rates and future energy price scenarios (1.4% annual price growth, 0% price growth, and prices projected by the MERGE-ETL 6.0 IAM). All estimates are computed using the socioeconomic scenario SSP3. Brackets indicate interquartile ranges that reflect climate sensitivity uncertainty, damage function uncertainty, or both (i.e. full uncertainty).

5.G Sensitivity of the energy consumption partial social cost of carbon

The partial SCC estimates shown in the main text and in Table 5.F.2 depend upon a set of future price scenarios, damage function extrapolation beyond 2100, and are reported for a particular socioeconomic scenario (SSP3). In this section, we provide a range of additional partial SCC estimates under alternative future price scenarios, alternative extrapolation approaches for the damage function, and multiple different socioeconomic scenarios. In all cases, we show estimates under multiple discount rates and baseline emissions trajectories.

5.G.1 Alternative future price scenarios

Table 5.G.1 presents partial SCC estimates under a range of future price scenarios. Three of these scenarios are based on direct extrapolation of present-day price statistics at either moderate (1.4%), stagnant (0%), or high (3%) annual growth rates (*Appendix* 5.D.1). In addition, we show partial SCC estimates based on price projections from five different integrated assessment models (*Appendix* 5.D.2). Estimates are qualitatively similar across all price scenarios.

5.G.2 Alternative approach to estimating post-2100 damages

We explore the importance of post-2100 extrapolation of the damage function (*Appendix* 5.E) by using an alternative approach to estimating post-2100 damages, in which we calculate partial SCC estimates using a damage function held fixed at its end-of-century shape for all years 2100-2300. With this alternative approach, our central estimate of the energy consumption partial SCC (1.4% price growth, 3% discount rate) changes from \$-1.16 to \$-1.08 for RCP8.5 (\$-1.08 to \$-1.07 for RCP4.5), indicating that extrapolation of the damage function has negligible impact on our partial SCC estimates, due in part to the important role of discounting (Table 5.G.2).

5.G.3 Alternative socioeconomic scenarios

In the main text, we display climate change impact projections and estimates of the partial social cost of carbon under the socioeconomic scenario SSP3. Each SSP scenario models a different possible pathway of economic development, population growth, and demographics. Here, we show estimates of the energy consumption partial social cost of carbon under two alternative scenarios (SSP2 and SSP4, alongside SSP3). Results from these alternative scenarios are similar in magnitude to those from SSP3.

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
		1.4% price growth	
RCP 8.5	-1.51 [-2.58,-0.66]	-1.16 [-1.92,-0.55]	-0.60 [-0.97,-0.35]
RCP 4.5	-1.37 [-2.45,-0.73]	-1.08 [-1.83,-0.61]	-0.58 [-0.90,-0.35]
		0% price growth	
RCP 8.5	-0.72 [-1.12,-0.36]	-0.61 [-0.95,-0.34]	-0.39 [-0.62,-0.23]
RCP 4.5	-0.66 [-1.08,-0.39]	-0.57 [-0.91,-0.35]	-0.37 [-0.58,-0.23]
		3% price growth	
RCP 8.5	-4.43 [-8.28,-1.52]	-2.98 [-5.36,-1.15]	-1.11 [-1.81,-0.59]
RCP 4.5	-4.00 [-7.74,-1.85]	-2.77 [-5.03,-1.37]	-1.06 [-1.70,-0.62]
		MERGE-ETL 6.0 prices	
RCP 8.5	-1.12 [-1.66,-0.61]	-0.82 [-1.26,-0.47]	-0.39 [-0.63,-0.24]
RCP 4.5	-1.15 [-1.90,-0.76]	-0.83 [-1.36,-0.54]	-0.38 [-0.61,-0.24]
		REMIND 1.7 (ADVANCE) prices	
RCP 8.5	-4.14 [-6.57,-2.28]	-3.15 [-4.97,-1.82]	-1.57 [-2.49,-0.96]
RCP 4.5	-4.10 [-6.56,-2.49]	-3.13 [-4.92,-1.93]	-1.55 [-2.38,-0.96]
		REMIND 1.7 (CEMICS) prices	
RCP 8.5	-4.01 [-6.74,-2.26]	-3.04 [-5.00,-1.17]	-1.51 [-2.41,-0.92]
RCP 4.5	-4.01 [-6.35,-2.26]	-3.05 [-4.75,-1.78]	-1.50 [-2.29,-0.92]
		REMIND-MAgPIE 1.7-3.0 prices	
RCP 8.5	-3.96 [-6.40,-2.24]	-3.02 [-4.80,-1.78]	-1.52 [-2.38,-0.94]
RCP 4.5	-3.95 [-6.30,-2.41]	-3.03 [-4.72,-1.88]	-1.51 [-2.29,-0.95]
		WITCH-GLOBIOM 4.2 prices	
RCP 8.5	-4.01 [-6.48,-2.01]	-2.98 [-4.83,-1.60]	-1.41 [-2.27,-0.84]
RCP 4.5	-3.83 [-6.40,-2.20]	-2.88 [-4.72,-1.67]	-1.37 [-2.17,-0.82]

[Brackets] indicate full uncertainty IQR. (See Appendix 5.F.4.)

Table 5.G.1: Social cost of energy consumption due to climate change under alternative future price scenarios. This table displays estimates of a partial Social Cost of Carbon for excess energy consumption costs, under the socioeconomic scenario SSP3. Costs are valued under various projected price trajectories.

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
		1.4% price growth	
RCP 8.5	-1.33	-1.08	-0.60
RCP 4.5	-1.34	-1.07	-0.57
		0% price growth	
RCP 8.5	-0.71	-0.60	-0.39
RCP 4.5	-0.70	-0.59	-0.37
		3% price growth	
RCP 8.5	-3.24	-2.42	-1.06
RCP 4.5	-3.30	-2.44	-1.04
		MERGE-ETL 6.0 prices	
RCP 8.5	-0.87	-0.70	-0.38
RCP 4.5	-0.92	-0.73	-0.37
		REMIND 1.7 (ADVANCE) prices	
RCP 8.5	-3.65	-2.91	-1.56
RCP 4.5	-3.79	-2.98	-1.53
		REMIND 1.7 (CEMICS) prices	
RCP 8.5	-3.55	-2.83	-1.50
RCP 4.5	-3.70	-2.90	-1.48
		REMIND-MAgPIE 1.7-3.0 prices	
RCP 8.5	-3.51	-2.80	-1.50
RCP 4.5	-3.66	-2.89	-1.50
		WITCH-GLOBIOM 4.2 prices	
RCP 8.5	-3.50	-2.74	-1.39
RCP 4.5	-3.55	-2.74	-1.36

Table 5.G.2: Social cost of energy consumption due to climate change under alternative approach to estimating post-2100 damages. This table displays estimates of a partial Social Cost of Carbon for excess energy consumption costs under the socioeconomic scenario SSP3. In contrast to the estimates in Tables 5.4E (main text) and 5.G.1, these estimates are calculated using a damage function held fixed at its end-of-century shape for all years 2100-2300.

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
		SSP2	
RCP 8.5	-1.08	-0.93	-0.58
RCP 4.5	-1.29	-1.06	-0.60
		SSP3	
RCP 8.5	-1.51	-1.16	-0.60
RCP 4.5	-1.37	-1.08	-0.58
		SSP4	
RCP 8.5	-1.36	-1.05	-0.55
RCP 4.5	-1.25	-0.98	-0.51

Table 5.G.3: Social cost of energy consumption due to climate change under alternative socioeconomic scenarios. This table displays estimates of a partial Social Cost of Carbon for excess energy consumption costs under various socioeconomic scenarios, for the 1.4% price annual growth trajectory.

5.H Incidence of future climate change impacts by present-day income deciles

Although we project that future climate change will lead to modest net savings in energy expenditures globally (Figure 5.3B in the main text; Figure 5.D.1), these savings are not borne equally across locations. The map in main text Figure 5.3A reveals heterogeneity in costs and savings at end-of-century by geography; Figure 5.H.1 demonstrates that this heterogeneity is systematically correlated with present-day income.

Over most of the present-day income distribution, we find that locations with higher incomes today are projected to experience larger overall net savings at end-of-century. This partly reflects the fact that today's richest locations tend to be in temperate climates, where energy savings from fewer cold days will more than offset increases in costs from more hot days. The smallest savings at end-of-century are projected to occur in the third and fourth deciles of the present-day income distribution, which is consistent with many of these locations being situated in the tropics and also attaining sufficiently high income levels at end-of-century to increase electricity consumption due to more hot days (Figure 5.2A in the main text).

The positive correlation between present-day income and net savings at end-of-century does not hold in the lower ranges of today's income distribution. Net savings in today's poorest deciles (i.e. first and second) are actually higher than in the third and fourth deciles, as many of the poorest locations are projected to remain too poor at end-of-century to increase electricity consumption on hot days.

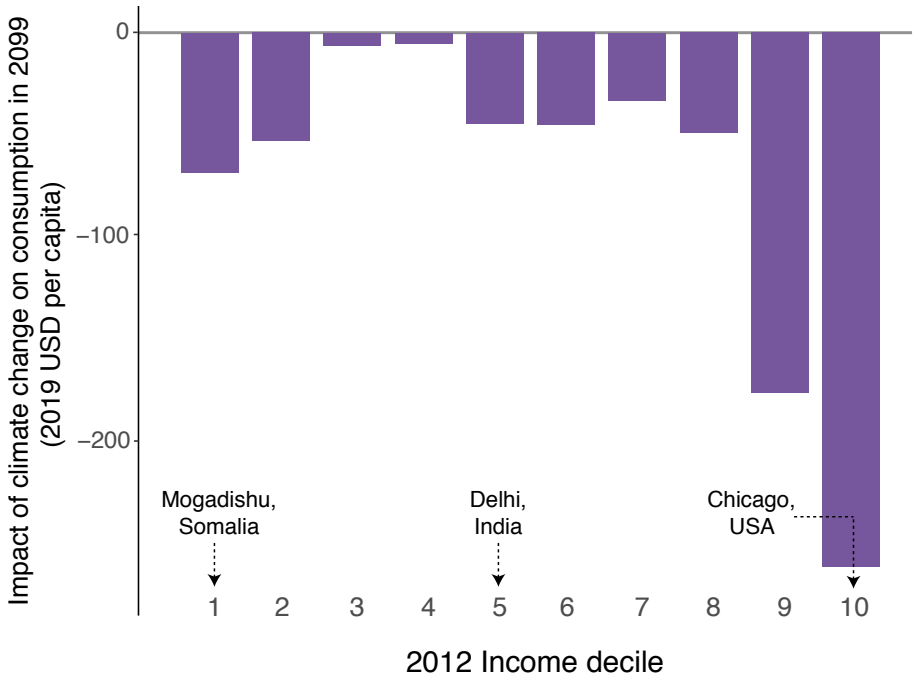

Figure 5.H.1: End-of-century energy consumption damages from climate change by present-day income deciles. The bar chart below depicts annual energy consumption damages from climate change at 2099 under a high emissions scenario (RCP8.5) and the SSP3 socioeconomic scenario, separately for each decile of 2012 per capita income. Income deciles are calculated across 24,378 global impact regions and are population weighted using 2012 population values; representative locations in selected deciles are indicated. Damages are calculated under a 1.4% annual price growth scenario and are expressed in 2019 USD per capita based each decile's projected 2099 population. Bars represent mean damage estimates across an ensemble of 33 climate models.

5.I Robustness and sensitivity checks

In this section, we explore how alternative assumptions about income growth and climate-driven adaptation affect our projected impacts. In addition, we conduct a robustness checks to address issues of data quality. Finally, we describe an alternative estimation approach, impact projection, and partial SCC calculation that reflect a particular scenario of persistent technological progress, where the cost of energy services continues to decline indefinitely based on historically measured rates, without corresponding gains in efficiency.

5.I.1 Alternative assumptions on the role of income growth and climate-driven adaptation

In our main results, we use the estimated coefficients from Equation 5.C.4 in combination with high-resolution data on future incomes ($\overline{LogGDPPC}$) and climates ($\overline{HDD}$, $\overline{CDD}$) to extrapolate energy-temperature responses over time, thus capturing future changes in the responses due to income growth and climate-driven adaptation (*Appendix* 5.C.3). In conducting this extrapolation, it is necessary to make an assumption regarding the rates at which energy-temperature responses evolve with changing incomes and climates. As discussed previously, our projections rely on 15-year moving averages of $\overline{LogGDPPC}$, $\overline{HDD}$, and $\overline{CDD}$ (*Appendix*, 5.C.3).

Here, we conduct a sensitivity analysis where the speed at which energy-temperature responses change with income and climate covariates is deterministically reduced by half. This exercise is used to understand how the impacts of future climate change differ if energy-temperature responses evolve more slowly with income and climate than is estimated in the historical data.[29]

In the main projection, income grows for each impact region r according to $GDPPC_{rt} = \rho_{jt} GDPPC_{r,t-1}$, where j indicates the country that region r falls into, and ρ_{jt} is a country- and year-specific growth rate given exogenously by the SSP scenarios. The moving average values of heating degree days for region r used in the main projection are specified by $\overline{HDD}_{rt} = \overline{HDD}_{r,t-1} + \Delta\overline{HDD}_{rt}$, and moving averages of cooling degree days are specified similarly. In this "slow adaptation" alternative approach, we replace income growth with $GDPPC_{rt} = \left(\frac{\rho_{jt}-1}{2} + 1\right) GDPPC_{r,t-1}$ after the year 2015, and we reduce linear change in heating degree days by replacing it with $\overline{HDD}_{rt} = \overline{HDD}_{r,t-1} + \frac{\Delta\overline{HDD}_{rt}}{2}$. The linear change in cooling degree days is similarly replaced. Note that both the main and slow-adaptation analyses generate identical incomes, heating degree days, and cooling degree days (and hence, energy-temperature responses) in 2015.

Figure 5.I.1 displays the time series of global per-capita electricity and other-fuels consumption impacts under a slow-adaptation scenario. Impacts under slow-adaptation are seen to be considerable lower in magnitude than our main estimates.

[29] Note that in our econometric estimation of Equation 5.C.4, income is modeled as a 15-year moving average, while climate is treated as a time-invariant average over the entire sample (*Appendix*, 5.C.3), due to limited in-sample variation in the average climate. Thus, the decision of how quickly the energy-temperature response evolves with climate in the future is less informed by historical data than is the case for income.

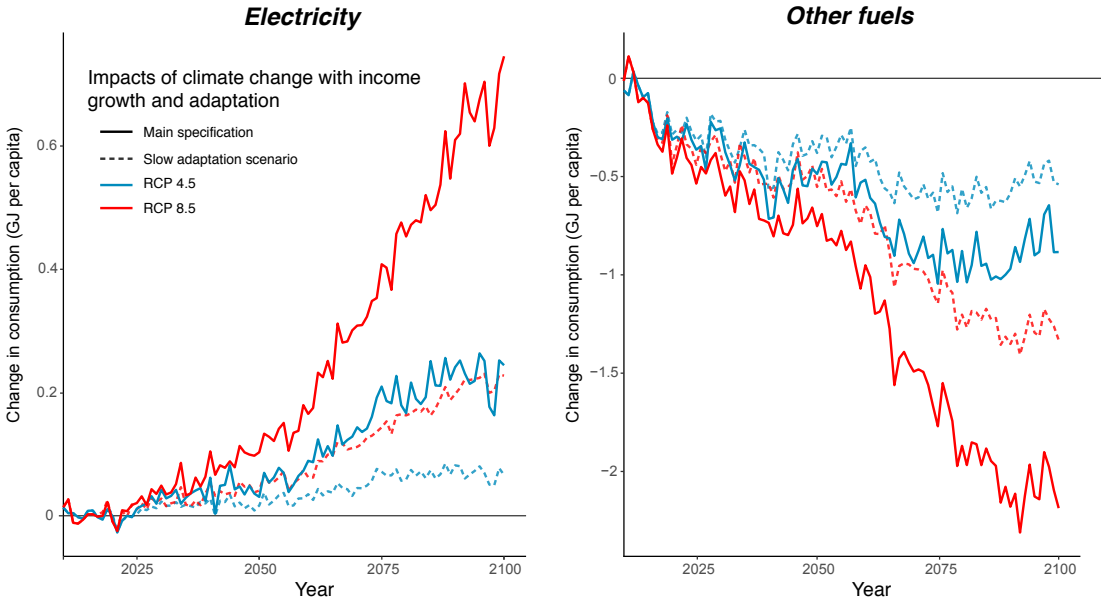

Figure 5.I.1: Projected impacts of climate change under a slow-adaptation scenario. The dashed lines present the time series of average global energy consumption impacts of climate change (in GJ per capita) under a slow-adaptation scenario, where the speed at which energy-temperature responses change with income and climate covariates is reduced by half. To contrast, the solid lines present projected impacts based on the estimates of the main model (Equation 5.C.4). All impacts are calculated under the SSP3 socioeconomic scenario using climate projections from a single climate model (CCSM4).

5.I.2 Data quality

The IEA data on energy consumption are of differential quality across countries as well as years within a country, and issues of quality and comparability are extensively noted in the documentation. All our specifications account for data quality and comparability concerns through inverse variance weighting, imposition of country-regime fixed effects, and dropping observations that fail to meet basic standards of comparability (*Appendix*, 5.A.1 and 5.C.1). To further ensure that data quality issues do not pose a threat to our core findings, here we conduct a robustness check in which we re-estimate Equation 5.C.4 after dropping additional observations in which the energy consumption data is obtained by imputation.[30] If the energy consumption data in such cases are themselves estimated as a function of income and climate variables, it is possible that our results simply reflect this imputation procedure rather than a real causal relationship. Figure 5.I.2 displays the matrices of energy-temperature responses when estimating Equation 5.C.4 including and not including observations with imputed energy consumption data (dark and light response curves respectively).[31] The predicted responses are similar, thus our results do not seem to be an artifact of imputed energy consumption data.

5.I.3 Modeling technological progress

A potential limitation of the energy-temperature response specified in Equation 5.C.4 is that it does not explicitly model the role of technological progress in altering how en-

[30] The documentation notes instances where the energy consumption data are derived solely from estimates by either the IEA or country government. For electricity consumption 145 additional country-years are dropped, while for other fuels consumption 465 additional country-years are dropped.

[31] The dark curves are identical to those from the matrices in Figure 5.1C (main text).

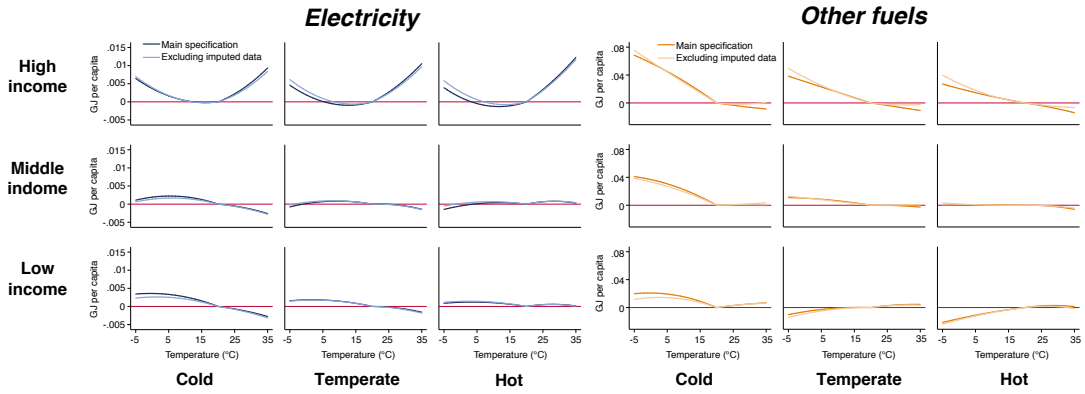

Figure 5.I.2: Energy-temperature response as a function of income and climate: Robustness to imputed data. This figure displays the results from an econometric specification that models heterogeneity in the energy-temperature response due to both income and long-run climate (Equation 5.C.4). The dark response curves are identical to those from Figure 5.1C (main text), while the light response curves are estimated using a subset of the data in which observations with imputed energy consumption data are dropped. Each cell within a matrix represents predicted energy-temperature responses at a point in the income × long-run climate covariate space within the full sample. Cells are ordered vertically by income terciles (increasing income from bottom to top) and horizontally by terciles of annual cooling degree-days (increasingly warm climate from left to right) (*Appendix* 5.C.3).

ergy consumption responds to temperature. By allowing energy-temperature responses to evolve into the future with incomes and climates, we do proxy for diffusion and advancement of existing technologies in accordance with these two factors. However, one may imagine scenarios in which new technologies reduce the cost of cooling/heating, leading to increasingly temperature-responsive energy consumption in the future, even after conditioning on income and climate.[32] Alternatively, new technologies may increase the energy efficiency of cooling/heating, leading to less temperature-responsive energy consumption in the future after conditioning on income and climate.

To accomodate such possibilities, we augment Equation 5.C.4 with a linear time trend

[32] A historical example of this is the declining cost and consequent proliferation of air-conditioning units in the decades after they first entered the marketplace.[25,71]

in the effect of temperature on energy consumption that flexibly varies by income:

$$
\begin{aligned}
E_{jtc} = \quad & \underbrace{\boldsymbol{\beta_c} \cdot \boldsymbol{T}_{jt}}_{\text{Effect of Temperature}} + \underbrace{[\boldsymbol{\tau}_c \cdot \boldsymbol{T}_{jt}]t}_{\substack{\textbf{Linear Time Trend} \\ \textbf{in Energy-Temperature Response } (A)}} \\
& + \underbrace{[\boldsymbol{\eta}_{1c} \cdot \boldsymbol{T}_{jt}](\bar{I}_c - \overline{lGDPpc}_{jt})\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c} + [\boldsymbol{\eta}_{2c} \cdot \boldsymbol{T}_{jt}](\overline{lGDPpc}_{jt} - \bar{I}_c)\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}}_{\text{Effect of Income-driven Adaptation on Energy-Temperature Response}} \\
& + \underbrace{[\boldsymbol{\psi}_{1c} \cdot \boldsymbol{T}_{jt}]t(\bar{I}_c - \overline{lGDPpc}_{jt})\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c} + [\boldsymbol{\psi}_{2c} \cdot \boldsymbol{T}_{jt}]t(\overline{lGDPpc}_{jt} - \bar{I}_c)\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}}_{\textbf{Linear Time Trend in Energy-Temperature Response: Varying by Income } (B)} \\
& + \underbrace{\sum_{k=1}^{2} \gamma_{kc}\overline{CDD}_j \sum_{d\in t}(T_{jd}^k - 20^k)\mathbf{I}_{T_{jd}\geq 20}}_{\substack{\text{Climate Adaptation on Energy-Temperature Response} \\ \text{for days} \geq 20^\circ \text{ C}}} + \underbrace{\sum_{k=1}^{2} \lambda_{kc}\overline{HDD}_j \sum_{d\in t}(20^k - T_{jd}^k)\mathbf{I}_{T_{jd}<20}}_{\substack{\text{Climate Adaptation on Energy-Temperature Response} \\ \text{for days} < 20^\circ \text{ C}}} \\
& + \underbrace{\left[\kappa_{1c}\overline{lGDPpc}_{jt} + \phi_1\right]\mathbf{I}_{\overline{lGDPpc}_{jt}<\bar{I}_c} + \left[\kappa_{2c}\overline{lGDPpc}_{jt} + \phi_2\right]\mathbf{I}_{\overline{lGDPpc}_{jt}\geq\bar{I}_c}}_{\text{Direct Effect of Income on Energy Consumption}} \\
& + \underbrace{\boldsymbol{\theta}_c \cdot \boldsymbol{P}_{jt}}_{\text{Precipitation Controls}} + \underbrace{\alpha_{jic} + \delta_{wtc}}_{\text{Fixed Effects}} + \underbrace{\varepsilon_{jtc}}_{\text{Error Term}} . \qquad (5.I.1)
\end{aligned}
$$

All terms in Equation 5.I.1 are identical to those of Equation 5.C.4, with the exception of terms A and B, which specify a time trend in the effect of temperature on energy consumption that varies by income. These additional terms capture historical technological progress that altered the cost and/or efficiency of energy services (e.g. introduction and proliferation of air-conditioning) within the sample time span. Furthermore, the additional terms allow for linear extrapolation of such progress when predicting responses into the future.

Formally, for $\overline{LogGDPPC} = I$, the per year change in the energy-temperature response at T° C (relative to 20° C) is $\left(\boldsymbol{\tau}_c + \boldsymbol{\psi}_{1c}(\bar{I}_c - I)\mathbf{I}_{I<\bar{I}_c} + \boldsymbol{\psi}_{2c}(I - \bar{I}_c)\mathbf{I}_{I\geq\bar{I}_c}\right) \cdot \left(\boldsymbol{T} - \mathbf{20}\right)$, where the vector $\boldsymbol{T} \equiv (T, T^2)$ and the vector $\mathbf{20} \equiv (20, 20^2)$. The Figure 5.I.3A plots this per-year change evaluated at mean incomes of the three income terciles and demonstrates that electricity-temperature responses and other fuels-temperature responses at higher income levels are becoming more steeply sloped over time. For instance, the additional electricity consumption on a 35°C day relative to a 20°C day is increasing annually by .00067 GJ per capita at the richest income tercile, -.00018 GJ per capita at the middle income tercile, and 6.1e-06 GJ per capita at the poorest income tercile.

Extrapolating these rates of change linearly into the future dramatically amplifies the predicted energy-temperature responses across the income-climate space. To illustrate this, Figure 5.I.3B depicts matrices of energy-temperature responses in-sample using Equation 5.C.4 (dark response curves), as compared to the year 2099 using Equation 5.I.1 (light response curves).[33]

Figure 5.I.3C illustrates the consequences of incorporating time trends in energy-temperature responses (in addition to changing income and climate covariates) into projections of climate change impacts. Incorporating the time trends results in dramatically larger impacts by end-of-century, with global average electricity consumption predicted

[33]The dark curves are identical to those from the matrices in Figure 5.1C (main text).

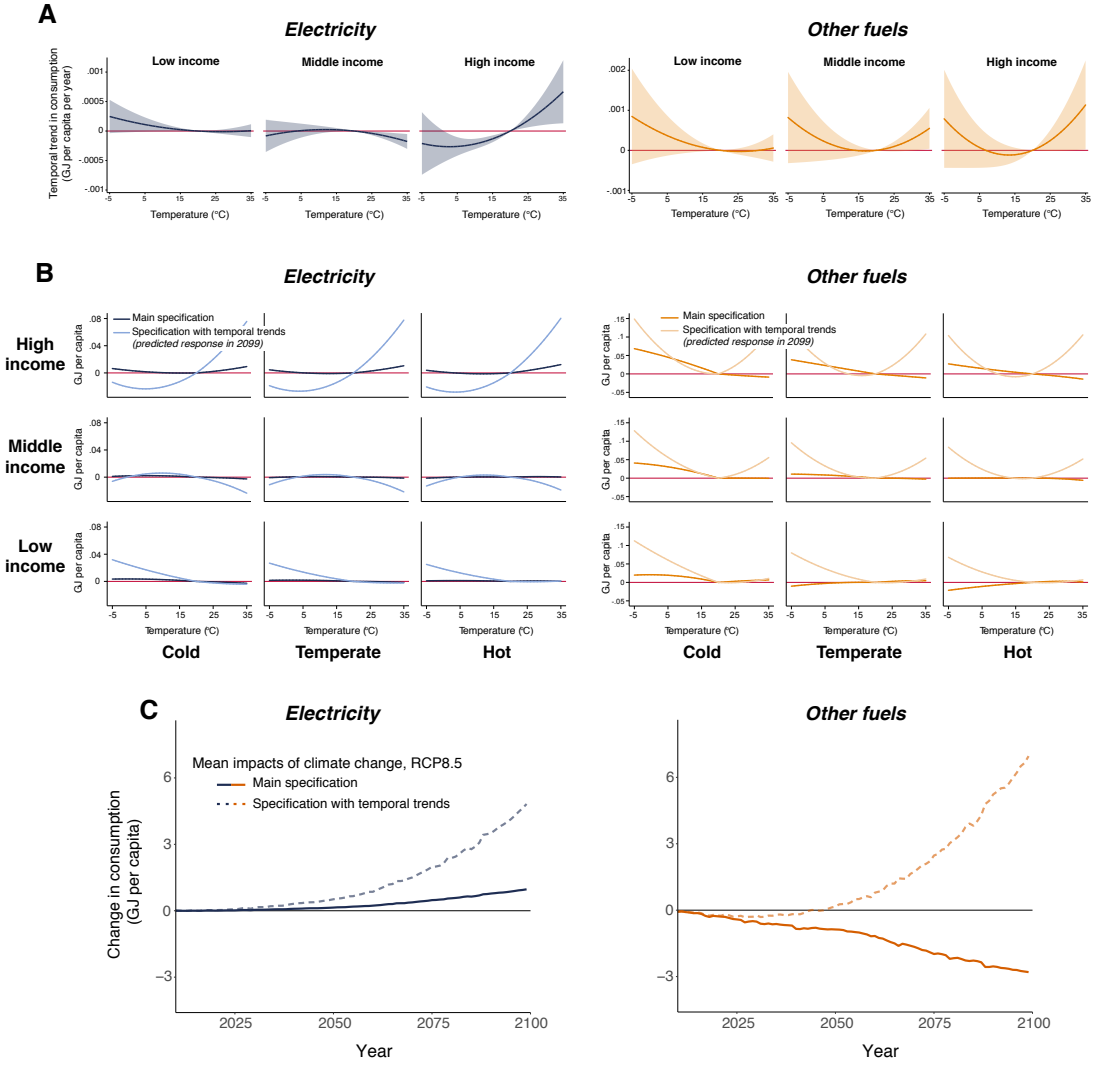

Figure 5.I.3: Modeling technology trends. This figure displays results from an econometric specification that models heterogeneity in the energy-temperature response due to both income and long-run climate, while also modeling a linear time trend in the effect of temperature on energy consumption that flexibly varies by income and fuel category (Equation 5.I.1). **(A)** The change in the electricity- and other fuels-temperature responses per year for each income tercile. **(B)** The predicted temperature response functions at various points in the income × long-run climate covariate space within our data sample. Light curves illustrate predicted responses at 2099 that reflect a linear extrapolation of the annual rates of change shown in (A). As a benchmark, the dark response curves are estimated from Equation 5.C.4 and are identical to those from Figure 5.1C (main text). Cells are ordered vertically by income terciles (increasing income from bottom to top) and horizontally by terciles of annual cooling degree-days (increasingly warm climate from left to right) (*Appendix* 5.C.3). **(C)** The time series of average global impacts in GJ per capita under a high emissions scenario (RCP8.5) and SSP3 socioeconomic scenario, when allowing for energy-temperature responses in the future to reflect linear time trends in addition to changing incomes and climates (dashed lines). To contrast, the solid lines indicate the analogous time series based on the main model, without linear time trends (Equation 5.C.4). All impact projections represent mean estimates across an ensemble of 33 climate models.

to rise by 4.82 GJ per capita (in contrast to 0.97 GJ per capita shown in Figure 5.2 in the main text) and other fuels consumption by 6.96 GJ per capita (in contrast to -2.8 GJ per capita shown in Figure 5.2) due to climate change.

Constructing damage functions and calculating a partial SCC with these projected impacts results in a substantially larger partial SCC for a given discount rate (Table 5.I.1, Panel II). Additionally, we calculate an SCC based on a scenario in which historical trends not only continue unabated, but are accelerated as future climate change spurs increases in the relative rate of technological progress and cost declines in cooling/heating (without efficiency gains).[40] To implement such a scenario, we deterministically double the historically estimated time trends in energy-temperature responses and extrapolate these linearly into the future. This results in an even higher partial SCC (Table 5.I.1, Panel III).

Discount rate:	$\delta = 2.5\%$	$\delta = 3\%$	$\delta = 5\%$
		Main model	
RCP 8.5	-1.51	-1.16	-0.60
RCP 4.5	-1.37	-1.08	-0.58
		I: Extrapolating trends	
RCP 8.5	9.33	5.67	1.24
RCP 4.5	9.96	5.88	1.20
		II: Extrapolating 2x trends	
RCP 8.5	18.95	11.49	2.52
RCP 4.5	20.06	11.83	2.43

Table 5.I.1: Social cost of energy consumption due to climate change (technological trends scenarios). This table displays estimates of a partial Social Cost of Carbon for excess energy consumption costs (under the socioeconomic scenario SSP3) based on scenarios of indefinitely falling costs of energy services without efficiency gains. These are contrasted with partial SCC estimates from the main model. **(I)** As a benchmark, the estimates resulting from the main model (Equation 5.C.4) under SSP3 (reproducing Table 5.4E, Panel I). **(II)** The estimates resulting from direct linear extrapolation of the energy-temperature response time trends estimated in Equation 5.I.1. **(III)** The estimates resulting from linear extrapolation at double the rate of the estimated time trends. In all panels, costs are valued under the 1.4% price growth trajectory.

It should be noted that linearly extrapolating historical trends in energy temperature responses (beyond those associated with income and climate) indefinitely into the future not only represents an extreme assumption about ongoing declines in the cost of energy services relative to other goods, but also assumes that these cost declines will not be accompanied by efficiency gains (which would lower fuel use). Climate-change driven acceleration (i.e. doubling) of these trends represents an even more extreme assumption. It is unclear that such arbitrary assumptions are likely to hold. Future research should focus on better understanding the specific ways in which a changing climate and climate policy (or lack thereof) might feed back into technological innovation.

5.J Comparisons to other studies

In this section, we consider prior econometric and process model-based estimates of climate change impacts on energy consumption, and compare them to their nearest counterparts from among our estimates. Although not perfectly comparable in their scope or their intent, we generally find that estimates from prior studies are qualitatively consistent with those from this analysis (Table 5.J.1).

Ref.[27] estimate the effect of temperature on electricity consumption using daily data from European countries, and use these estimates to project the impacts of climate change under one of the 33 climate projections (i.e. GFDL-ESM2M) in our SMME. Their analysis encompasses electricity consumption in all sectors, including transportation (which we exclude).[34] Moreover, they do not account for the effect of income differences or climate-driven adaptation on the electricity-temperature response. Nonetheless, their end-of-century impact projections under RCP8.5 for major European countries are comparable in magnitude and direction to ours (Table 5.J.1A).

Using a combination of econometrically estimated energy-temperature responses along with price, supply-side, and macroeconomic feedbacks determined through the NEMS energy model, ref.[37] project the impact of climate change on total energy expenditures in the United States. Calculating a mean impact across the full set of climate projections in the SMME, they project a 12.1% increase in annual energy expenditures at end-of-century (relative to 2012) due to climate change under RCP8.5, while we project a 9.0% decrease (Table 5.J.1B). However, the two studies have important differences in scope and approach. Most fundamentally, the analysis in ref.[37] is limited to the residential and commercial sectors (excluding industrial and agricultural consumption, both of which we include) and also does not account for the effect of income differences or climate-driven adaptation on energy-temperature responses. Furthermore, the estimates in ref.[37] are specific to the endogenously determined price trajectories of NEMS.

Lastly, ref.[32] use a process model-based approach to project the impact of climate change on global cooling and heating expenditures under the SSP2 socioeconomic scenario. Using one of the 33 climate projections (i.e. CESM1-BGC) in our SMME, they project that at end-of-century, climate change under RCP8.5 will lead to a net increase in global expenditures representing 0.3% of 2100 global GDP (Table 5.J.1C). For the same socioeconomic scenario, climate projection, and RCP, we project a similarly modest change in global expenditures (a net decline representing 0.03% of 2100 global GDP), despite using a purely econometric approach. Although the magnitudes are similar, it should be noted that the two studies do not purport to estimate the same object. While the authors explicitly model changes to cooling and heating demands, our estimates encompass changes from all end uses, not necessarily limited to cooling and heating. Another important distinction is that the estimates in ref.[32] are specific to the endogenously determined price trajectories of the GCAM energy model.

[34] As of 2013, transportation accounted for less than 3% of electricity consumption in the EU. See https://www.eea.europa.eu/ds_resolveuid/8d7f4a83fc3a4936be07b8d40c24352c.

Geography	Study Estimate	Our Estimate
	A. Wenz et al., 2017	
	Change in annual per-capita electricity consumption at end-of-century under RCP8.5, as percent of 2012 consumption.	
France	0.9%	0.4%
Germany	-0.8%	-1.4%
Greece	7.3%	5.4%
Italy	1.3%	2.7%
Spain	5.2%	4.0%
United Kingdom	-1.7%	-1.4%
	Notes: Both estimates are based on climate projections from the GFDL-ESM2M climate model. Wenz et al.'s estimates include electricity consumption in the transportation sector, which we exclude.	
	B. Hsiang et al., 2017	
	Change in annual per-capita energy expenditure at 2090 under RCP8.5, as percent of 2012 expenditure.	
United States	12.1%	-9.0%
	Notes: Both estimates represent mean impacts across all climate projections in the SMME (*Appendix* 5.A.2.3). Hsiang et al.'s estimates are for the period 2080-2099 and do not include energy consumption in the industrial sector, which we include. Our estimates determine future expenditures under the 1.4% price growth scenario, while those of Hsiang et al. use prices determined endogenously through the NEMS (AEO 2013) energy model.	
	C. Clarke et al., 2018	
	Change in total annual energy expenditures at end-of-century under RCP8.5 and SSP2 socioeconomic scenario, as percent 2100 global GDP.	
Global	0.3%	-0.03%
	Notes: Both estimates are based on climate projections from the CESM1-BGC climate model. Clarke et al.'s estimates explicitly model changes to cooling and heating demands, while our estimates are encompass changes from all end uses, not limited to cooling and heating. Our estimates determine future expenditures under the 1.4% price growth scenario, while those of Clarke et al. use prices determined endogenously through the GCAM energy model.	

Table 5.J.1: Comparison of our climate change impact projections with those of other studies. This table compares our projected impacts of climate change on energy consumption (right column) with those of previous studies that focus on specific geographical regions and fuels (middle column). Our projected impacts are derived from the main model (Equation 5.C.4). Caveats to comparability are noted below each study comparison.

5.K Feedback effects of climate change-induced energy consumption on emissions

Changing energy consumption patterns due to climate change have the potential to alter the trajectory of CO_2 emissions via changes to emissions-intensive energy production. However, given the small size of our estimated impacts, these feedbacks are plausibly negligible. To examine whether this is the case, in Table 5.K.1, we calculate the additional emissions resulting from our projected impacts to electricity and other fuels consumption at 2099 under the RCP8.5 emissions scenario and SSP3 socioeconomic scenario. Assuming all climate change impacts to electricity consumption are powered by a combined cycle natural gas plant with 46 % efficiency (i.e. the average 2015 efficiency of US natural gas-fired combined-cycle technology),[72] our projected increases in electricity consumption would lead to 1.3 additional Gt CO_2e of annual emissions in 2099. Assuming all climate change impacts to other fuels consumption are due to changes in natural gas consumption, our projected decreases in other fuels consumption would reduce annual emissions by 1.8 Gt CO_2e in 2099.[35] Together, these feedbacks are negligible when compared to total 2099 global emissions of 100 Gt CO_2 under RCP 8.5.[73]

	Electricity	Other Fuels
Per capita impact (GJ)	1.0	-2.8
Total impact (billion GJ)	12.1	-34.9
Emissions factor (t CO_2e per GJ)	0.11	0.05
Additional emissions (Gt CO_2e)	1.3	-1.8

Notes: Global average per capita impacts at 2099 to electricity and other fuels consumption (Row 1) are taken from projections displayed in Figure 5.2C (main text), and are converted to global total impacts (Row 2) by multiplying by the projected world population in 2099 under SSP3. Multiplying total electricity impacts by the emissions factor for natural gas[74] scaled by 46% efficiency, and multiplying total other fuels impacts by the emissions factor for natural gas (Row 3) yields additional emissions from projected climate-change induced electricity and other fuels consumption (Row 4).

Table 5.K.1: Feedback effects of climate change-induced energy consumption on emissions. This table provides a calculation of additional emissions in 2099 resulting from projected impacts of climate change on energy consumption under the RCP8.5 emissions scenario and SSP3 socioeconomic scenario. The calculation assumes that all climate change impacts to electricity consumption are powered by a combined cycle natural gas plant with 46 % efficiency (i.e. the average 2015 efficiency of US natural gas-fired combined-cycle technology), and all climate change impacts to other fuels consumption are due to changes in natural gas consumption.

[35] These assumptions likely lead to an upper bound in the magnitude of emissions feedbacks.

Chapter 6: Conclusion

Understanding the economic impacts of climate change and appropriately pricing the greenhouse gas externality is critically important to the welfare of future generations (Interagency Working Group on Socal Cost of Carbon, 2010; National Academies of Sciences, Engineering, and Medicine, 2017). Relying on the Ricardian approach, the seminal work of Mendelsohn, Nordhaus, and Shaw (1994) found that climate impacts on U.S. agricultural profits, net of adaptation, would likely only be modestly negative and possibly beneficial. Using fixed effects estimators more robust to omitted variables bias (Deschênes and Greenstone, 2007), the subsequent literature has found negative climate impacts across a range of outcomes, though estimates are generally unable to account for adaptation (Auffhammer and Schlenker, 2014; Carleton and Hsiang, 2016).

The studies presented here combine elements of both the Ricardian and within-estimator approaches: we use cross-sectional variation in long-run climate to uncover heterogeneity in the response to short-run (within-estimator) weather shocks. In so doing, we find that adaptation does indeed mitigate the effects of weather relative to the effect of an unanticipated shock.

However, for the first time we also uncover the implied costs of these adaptive actions. In contrast to Mendelsohn, Nordhaus, and Shaw 1994, we find that accounting for both adaptation and its costs, negative effects of climate change persist and are widespread. Severe ($\sim$40%) yield losses persist for major grain growing regions of the world; the costly migration of cropped areas is substantial, but only offsets $\sim$10% of producer losses; human mortality effects persist and are valued at 3.2% of end-of-century global GDP. Only in the case of global energy consumption do we find gains from climate change – in this case increases in the consumption of electricity on hot days as locations become hotter on average is offset by the decreased consumption of heating fuels as cold places both warm and experience fewer cold days.

Several conclusions follow. First, accounting for both adaptation and its costs is clearly important for grounding an accurate estimate of the social cost of carbon (SCC). Second, performing these analyses at a global scale and with high spatial resolution reveals important sources of heterogeneity – both in impacts as well as in the magnitude of adaptation – that regionally focused or geographically coarse studies might miss. Third, results such as these may be useful inputs to the Integrated Assessment Modeling community, who may be able to incorporate estimates like these into their damage functions, improving their empirical grounding.

References

Auffhammer, Maximilian and Wolfram Schlenker. 2014. "Empirical studies on agricultural impacts and adaptation." *Energy Economics* 46:555–561.

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

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

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

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.

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.