Title
The economics of environmental change: Essays on climate and water

Author
Carleton, Tamma

Publication Date
2018

Peer reviewed|Thesis/dissertation

The economics of environmental change:
Essays on climate and water

by

Tamma Anne Carleton

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 Solomon M. Hsiang, Co-chair
Professor Maximilian Auffhammer, Co-chair
Assistant Professor James M. Sallee
Assistant Professor Benjamin Faber

Spring 2018

The economics of environmental change:
Essays on climate and water

Abstract

The economics of environmental change:
Essays on climate and water

by

Tamma Anne Carleton

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Solomon M. Hsiang, Co-chair

Professor Maximilian Auffhammer, Co-chair

Many questions of importance for sustained human progress on a planet of finite resources remain unanswered. As environmental change accelerates, there is growing urgency to build an understanding of how the environment and society influence one another. This dissertation mobilizes novel, large-scale datasets in combination with methodological advances in casual inference and original theoretical contributions to study the interdependence between societies, economies, and two of the world's most vital natural endowments: the climate and freshwater. Throughout these essays, I estimate causal, policy-relevant relationships between global-scale processes of environmental change and human wellbeing. The first three essays provide new estimates of the climate's influence over market and non-market outcomes, generating insights in climate impact attribution, climate impact mechanisms, and the estimation of adaptation costs and benefits under anthropogenic climate change. The final essay studies the other direction of the society-environment relationship, measuring how changes in economic systems influence freshwater. This chapter uses satellite data to uncover links between agricultural policies and water depletion at a scale that has been impossible to investigate with standard water monitoring tools. Together, these essays aim to demonstrate that combination of data and methods across physical and social sciences can yield valuable answers to questions that have been discussed, debated, and left unresolved for generations.

In memory of my father, Vincent, whose example taught me to be curious and concerned about the world.

Contents

List of Figures

List of Tables

Acknowledgments

I am deeply grateful for the many peers, mentors, friends, and family who made this dissertation possible and gave me the incredible opportunity to pursue a PhD.

First and foremost, I would like to thank my advisor, mentor, and friend, Solomon Hsiang. Sol's guidance, impossibly high standards, and unwavering dedication to answering humanity's biggest questions have completely transformed my experience as a student and aspiring academic. His irrational confidence in my abilities has encouraged me to constantly grow, and I will forever be grateful for this challenge.

Many other faculty have defined my experience at Berkeley. I thank Max Auffhammer for believing in and supporting me from the very beginning, and for helping me make the hardest decisions both in research and in life. I thank Michael Greenstone for making me a better economist with every conversation, and for setting such an extraordinary example of what it means to be both a researcher and an influencer of policy. I thank Jim Sallee for guiding my research projects, Ted Miguel for inspiring and encouraging my ideas, and Marshall Burke for helping me navigate an interdisciplinary path through an economics degree.

I have the great fortune of being part of two unique research communities, the Global Policy Lab and the Climate Impact Lab, and I am grateful to my colleagues in both of these collaborations. In particular, I thank Jonathan Proctor for hours of stimulating conversation, Amir Jina and James Rising for leading by outstanding example, and Ashwin Rode, Ishan Nath, Michael Delgado, and Andy Hultgren for pursuing creative problem-solving with me. I thank the unparalleled research assistants at EPIC, as well as Terin Mayer and Megan Landin for exceptional project management.

Deirdre Sutula, Hilary Soldati, and, in particular, Jonathan Kadish, have been pillars of my time at Berkeley, and I thank them for their friendship and support. I am grateful for conversations with and guidance from Fran Moore, Ceren Baysan, Fiona Burlig, Patrick Baylis, Felipe Gonzalez, John Loeser, and Ed Rubin. And for many miles of hilly running and always-insightful conversation, I thank Megan Lang.

Thank you to my mother, Carolyn, for teaching me to care for and cultivate the land, and for embodying a standard of dedication to work and to community that I will always strive toward. I thank my brother, Alex, the rest of my Elk community, and my Oxford family of friends, for supporting the path that brought me here.

For seemingly infinite patience, generosity, and optimism, I thank my incredible partner, Scott. The happiness you bring me makes everything else possible.

Chapter 1

Overview

Scholars have long considered the coupled nature of human and environmental systems. For centuries, academics across disciplines have hypothesized about and attempted to quantify the degree to which the environment governs human wellbeing. Creating a climate classification system, Aristotle deemed the tropics completely uninhabitable due to their heat, while in *The Spirit of Laws* the French philosopher Montesquieu argued that the climate was a fundamental determinant of social and economic progress. Scholarly debates continue today regarding the importance of natural endowments in economic development, and the ability of social institutions to constrain or overcome environmental influences on human life (Gallup, Sachs, and Mellinger, 1999; Acemoglu, Johnson, and Robinson, 2001). A parallel line of questioning has concerned the other direction of the society-environment relationship, with scholars, activists and policymakers examining the ways in which social and economic changes influence the environment. Recent academic work across physical and social sciences has begun to systematically document the many ways human activities degrade natural resources and reshape environmental processes (Rockström et al., 2009).

While both sides of the interdependence between society and the environment have been discussed, analyzed, and politically acted upon for decades, many questions of importance for sustained human progress on a planet of finite resources remain unanswered. For example, there is an enduring debate regarding whether and how geographic characteristics determine processes of economic development (Acemoglu, Johnson, and Robinson, 2012). Similarly, while some scholars argue that the services provided to society by diverse plant and animal life generate enormous value (Costanza et al., 1997; Chichilnisky and Heal, 1998; Heal, 2000), the concept of this "natural capital" — an environmental counterpart to the physical and human capital building blocks of economic models — remains outside standard macroeconomics. In particular, many aspects of the society-environment relationship remain qualitatively discussed but unmeasured, despite pressing need for evidence to inform rapidly evolving policy landscapes in climate change, biodiversity conservation, freshwater depletion, and many other large-scale processes of environmental change. This dissertation mobilizes datasets from both physical and social sciences along with methodological innovations in causal inference to estimate policy-relevant relationships between processes of global environmental change

and human wellbeing. These four essays together address both directions of the society-environment relationship, focusing on two intricately linked natural endowments that are relied upon by populations across the globe: the climate, and water resources.

The first three chapters concern the climate and the unfolding impacts of anthropogenic climate change. While our understanding of how human activities regulate the atmosphere and our climate is well constrained by both theory and empirical evidence, the many complex ways in which the climate governs social and economic outcomes remains an area of ongoing research and debate. In recent years, an explosion of studies has uncovered causal links between changes in the climate and a wide range of social outcomes, from economic output to crime to human health. These studies take a quantitative approach, mixing data and methods from the climate, social, and statistical sciences to reveal the nature of climate's influence over society. Building on these innovations, the first three essays in this dissertation leverage existing methodologies to study new contexts, while generating original insights through employing new analytic approaches in three understudied areas: climate impact attribution, climate impact mechanisms, and adaptation costs and benefits. In the first chapter, my co-author Solomon Hsiang and I synthesize the new and rapidly growing climate impacts literature, while conducting novel analysis in climate impact attribution to demonstrate the magnitude of climate's influence over human society, both today and in the future. In the second chapter, I show that changes in the climate in India have grave consequences for the risk of suicide, focusing attention on identification of the mechanism through which this link manifests. In the third chapter, co-authored with an interdisciplinary team, we derive a new theoretical framework that enables us to empirically estimate both the costs and the benefits of adaptation to climate change in the context of the global mortality consequences of anthropogenic warming.

The final essay in this dissertation seeks to explore how changes in economic and social structures influence natural resources. This chapter takes a first step toward answering this question for freshwater, one of the world's most vital natural endowments. Human pressures on freshwater resources are accelerating, posing threats to populations and economies around the world. However, partly due to data constraints, most research on water resource management is local in scope. In this essay, I use a globally-comprehensive, satellite-based measure of total water storage to quantify water depletion at a scale that has been impossible with standard water monitoring tools. I show that agricultural policies, which shape humanity's most water-intensive enterprise, have substantial implications for water resources that are detectable in satellite records. The empirical findings shown here represent the first identification of the human footprint in freshwater resources at global scale.

Together, the essays in this dissertation aim to demonstrate that combination of data and methodologies across physical and social sciences can yield valuable answers to questions that have been discussed, debated, and left unresolved for generations.

1.1 Chapter summaries

Social and economic impacts of climate

In Chapter 2, my co-author Solomon Hsiang and I synthesize over 100 quantitative empirical studies contributing to the evolution of research on climate impacts, draw new insights that cut across subfields of the literature, and conduct new analysis to demonstrate the magnitude of climate's influence over human society. While our study serves as a reference for state-of-the-art methods and results in climate impact estimation, it also demonstrates the techniques behind and results from climate impact attribution. We conduct impact attribution exercises across a range of climate impact categories to point out that society may benefit substantially from addressing ongoing impacts of climate in the present. For example, we find that temperature today lowers U.S. maize yields by approximately 48%, and that current temperature climatologies slow global economic growth roughly 0.25 percentage points per year. In comparison, future warming may slow global economic growth rates by 0.28 percentage points per year by end of century, and lower maize yields in sub-Saharan Africa by 22% by 2050. In general, we estimate that the economic and social burden of the current climate tends to be comparable in magnitude to the additional projected impact caused by end-of-century anthropogenic climate changes. This finding reflects persistent "adaptation gaps" — differences between a hypothetical fully-adapted state, and the observed damages we quantify under current climatic conditions. We emphasize that future research that can uncover when, where, and why adaptation is or is not successful will generate major social benefits, today and in the future. Our findings point to climate as an important influence on the historical evolution of the global economy, and we argue that these impacts should inform how we respond to modern climatic conditions, as well as guide how we predict the consequences of future climate changes. These conclusions set the stage for the essay in Chapter 4, where my co-authors and I document adaptation gaps in the human health response to climate and develop a new method to explain the drivers of those adaptation gaps at global scale.

Climate and suicide in India

In Chapter 3, I study the suicide epidemic in India, where suicide rates have doubled since 1980. Using state-level nationally comprehensive panel data between 1967 and 2013, I demonstrate that fluctuations in the climate, particularly temperature, significantly influence suicide rates. For temperatures above 20°C, a 1°C increase in a single day's temperature causes approximately 70 suicides, on average. I use a suite of mechanism tests, including investigation of spatial, temporal, and seasonal heterogeneity in the temperature response, pattern matching of temperature responses across crops and suicide rates, and testing for the presence of lagged effects, to show that this effect likely occurs through temperature damages to crop yields. This finding suggests that growing season climate damages manifest as economically-motivated suicides. The main contributions of this study are to demonstrate that the climate has influence over

a key indicator of human hardship and to identify the mechanism through which that effect is likely to occur. However, I also contribute to literatures on climate change adaptation and climate impact attribution. Employing four distinct approaches to test for the presence of adaptive behavior, I find no evidence that acclimatization, rising incomes, or other unobserved drivers of adaptation to warming temperatures are occurring. I execute a counterfactual simulation following the technique demonstrated in Chapter 2 for climate impact attribution, quantifying the number of suicides that can be attributed to already-observed warming trends throughout India. I find that warming since 1980 is responsible for 59,300 suicides across India, accounting for 6.8% of the total upward trend experienced over this period.

The global mortality consequences of climate change

In Chapter 4, I work with an interdisciplinary group of co-authors to construct global estimates for the full value of changing mortality risk due to climate change. To the best of our knowledge, our study is the first high-resolution, probabilistic, and globally comprehensive valuation derived from empirical estimates that are plausibly causal. It is also the first empirical estimate for any type of non-market climate damages to take into account both the benefits and costs of adaptations that populations will likely undertake to protect themselves against higher temperatures. We assemble the most exhaustive micro-data set on mortality and climate to date (accounting for 56% of the global population) and develop an approach that allows us both to predict dose-response functions in regions of the world where mortality data are unavailable, while simultaneously modeling adaptation costs and benefits. We find a clear nonlinear global relationship between mortality and temperature, with hot and cold days both leading to excess mortality. However, we find that the temperature sensitivity of mortality is substantially lower in richer locations and in populations that experience these temperatures more frequently, indicating that the adaptation gaps described in Chapter 2 are at least partially driven by budget constraints and acclimatization in the case of mortality and temperature. We combine these findings with standard projections of income and population along with novel probabilistic climate change projections to estimate the full damages of excess mortality risk imposed by warming. These projections account for both the benefits of adaptations as well as their costs, which we estimate using a revealed preference framework. Using our preferred assumptions for the value of a statistical life (VSL), we estimate the total mortality related costs of climate change range from 1.7% (RCP4.5, low emissions scenario) to 6.8% (RCP8.5, high emissions scenario) of projected global GDP at end of century.

The implicit global water market

Human pressures on freshwater resources are accelerating, posing threats to populations and economies around the world. An extensive literature in economics has informed efforts to manage this resource, but attempts to empirically ground theoretical findings have been data-constrained, and the focus has been local—water is rarely studied at

the global scale. However, water is progressively an international resource, as it is embedded in tradable goods that increasingly cross country lines, linking production and consumption decisions in one location to water depletion in distant trading partners. To contribute to filling this research gap, in Chapter 5 I use a globally comprehensive dataset to show that human footprints can be uncovered at the global scale in this metric of changes in total water availability. My data come from the Gravity Recovery and Climate Experiment (GRACE), a satellite mission that uses the time-varying distance between two satellites to uncover changes to Earth's gravitational field, which in turn generate estimates of changes in water mass at all points on Earth. I develop a conceptual framework that links the explicit international market for water-intensive agricultural goods to an implicit global market for water. I show conceptually, and verify empirically, that policies which intervene in the explicit market have substantial impacts on water resources in the implicit market. I find that policies which increase the wedge between domestic and global prices of water-intensive crops cause measurable increases in water depletion. My estimates imply that when the difference between domestic agricultural producer prices and global prices increases by 100 percentage points for staple crops, annual water loss at a single fully cropped $1^\circ \times 1^\circ$ grid cell amounts to about 3% of agriculture's global annual water footprint. These effects are largest when water-intensive crops are subsidized, and are particularly severe in locations suitable to these crops. To put these impacts in perspective, I compare the magnitudes of these effects with variation in the climate. I find that on average, the effect of this price increase has roughly the same impact on water loss as changing monthly average temperature from 30°C to 40°C. To my knowledge, this chapter provides the first global-scale empirical assessment of the impact of agricultural policies on freshwater resources.

Chapter 2

Social and economic impacts of climate

2.1 A brief history of thought*

Does climate affect our society? Or does human willpower and ingenuity render climate largely irrelevant to our affairs, as we overcome environmental challenges with resilience and innovation? If climate affects our lives, how much does it matter and why? Thinkers have asked these questions for generations, wondering whether climatic differences between regions could be partially responsible for differences in politics, economies, and culture, and whether large-scale social transformations, such as the rise of golden ages and the fall of empires, could be triggered by climatic changes. Over the last decade, an innovative community of researchers has taken a rigorous quantitative approach to these questions—mixing data and methods from the climate, social, and statistical sciences—making unprecedented and exciting progress. In this article we review recent advances, findings, and open questions in this emerging interdisciplinary field.

Our focus is recent progress, but consideration of the social impact of climate is as old as the academy. Aristotle developed a climate classification system where the tropics were described as an uninhabitable "torrid zone" (Lee et al., 1952) and Montesquieu argued that climate played a fundamental causal role in determining the structure and prosperity of different societies (De Montesquieu, 2011). In the late 19th century, theories on the impact of climate and other geographical factors led to a collection of ideas known as "environmental determinism," the notion that environmental conditions played the primary role in shaping social, economic and political outcomes, with little scope for leadership, innovation, institutions or social will to alter societal trajectories. Some of these hypotheses were invoked to justify European colonialism as responsible paternalism—colonial advocates argued that climatically-caused "morally

*The material from this chapter is co-authored with Solomon Hsiang and was published in *Science* in September, 2016. The published version can be found here: `https://doi.org/10.1126/science.aad9837`.

inferior" character traits could be remedied through oversight by "advanced" societies that had already matured in more conducive climes (Howe, 2002).

The association of environmental determinism with colonial ambition had a chilling effect on this line of research in much of the social sciences during the late 20th century. Nonetheless, research continued among engineers and ergonomists interested in optimizing military and industrial performance using laboratory experiments to test the effects of environmental conditions on human performance (Mackworth, 1946; Froom et al., 1993).

Beginning in the 1970's, concern over booming populations led to a blossoming of theoretical work in resource economics. A key realization was that environmental conditions might influence economic performance and could be modeled as "natural capital," analogous to physical capital (e.g. machines) or human capital (e.g. education), and could be similarly developed or degraded (Solow, 1991).

At the turn of the 21st century, this economic approach, supported by advances in computing, led to the development of theoretical-numerical "integrated assessment models" that provide insights into how the global climate might be managed to maximize future "global welfare" under different assumptions (Nordhaus, 1993; Waldhoff et al., 2011; Stern, 2006). At the core of these models are theoretical "damage functions" that describe how global mean temperature translates into economic and social costs (Revesz et al., 2014). Because these models are now used to design global policies (Interagency Working Group on Socal Cost of Carbon, 2010; Department for Environment, Food and Rural Affairs, 2005), much of the current empirical research summarized here is framed as providing an empirical basis for global climate policy calculations (Kopp, Hsiang, and Oppenheimer, 2013; Burke et al., 2016).

A research agenda running parallel to climate change policy design is aimed at understanding how *current* climatic events, such as droughts or tropical cyclones, shape social outcomes today, irrespective of possible future climatic changes. This strand of work aims to minimize current social costs of climate events and promote economic development (Peduzzi, De Bono, and Herold, 2015; Kunreuther and Michel-Kerjan, 2009), either by identifying cost-effective risk-management strategies or minimizing harm through reactive instruments or policies, such as weather index insurance (Bryla-Tressler et al., 2011). As with climate change management, success in this arena depends critically on our quantitative understanding of the causal effect that climatic conditions have on populations.

2.2 Quantifying climatic influence on societies and economies

Recent advances in empirically measuring the effect of climate on society have been rapid, catalyzed by growth in computing power, access to data, and advances in the statistical theory of causal inference for non-experimental studies (Holland, 1986). Progress has been particularly explosive over the last decade, with exponential growth in publi-

cation volume due to innovations specific to studying the climate-human system, such as new methods to map climatic data onto social data and the development of spatiotemporal statistical models. For an in-depth treatment of the following techniques and innovations, we refer readers to Hsiang (2016).

Breaking down the problem

Climate is the joint probability distribution over several weather parameters, such as temperature or wind speed, that can be expected to occur at a given location during a specific interval of time (Fig. 2.1a-b). To understand how alterations in this distribution affect populations, modern approaches separate the influence of climate into two pathways: through information regarding what environmental conditions might occur and through directly altering what actually happens (Hsiang, 2016). The "informational" pathway operates because individuals' expectations about their climate (Fig. 2.1a) may change how they act; for example, individuals who believe they live in a rainy climate may purchase umbrellas. The "direct" pathway operates because any change in the probability distribution of weather events must generate a change in the distribution of events that individuals actually experience (Fig. 2.1b); for example, individuals who live in a rainy climate will face rain more often. Informational effects result from the fact that individuals prepare for the distribution of weather events and corresponding direct effects that they expect. These adaptations may alter the overall direct effect of specific weather events (Fig. 2.1c)—for example, individuals who own umbrellas may use them to stay drier when an actual rainstorm occurs—a distinction that can be accounted for when examining these relationships empirically.

Figure 2.1 depicts these two ways that climate and social outcomes are linked. Weather events (panel b) are drawn from the probability distribution that defines the climate (panel a). Each event generates some direct effect on a population, where these direct effects can be described by a *dose-response* function $f(X)$ where specific "doses" of a weather parameter X (e.g. rain) generate "responses" within the population (e.g. getting wet; see dashed arrow from panels b to d). This sequence of direct effects combine with non-climatic influences on the social outcome to produce the distribution of observed social data (panels d-e). If the climate shifts (pink in panel a), this will alter the distribution of weather (panel b) and its corresponding social impacts. A direct effect (e.g. experiencing more rainfall) occurs, but the information effect (e.g. buying umbrellas), may also cause populations to adapt such that the structure of the dose-response function changes (panel c), leading to a shift in the distribution of outcomes that is a combination and interaction of these two effects (panel e). The core of the empirical challenge is to credibly reconstruct the dose-response function for pairings of weather variables and social outcomes, while simultaneously accounting for the possibility that adaptations alter this relationship.

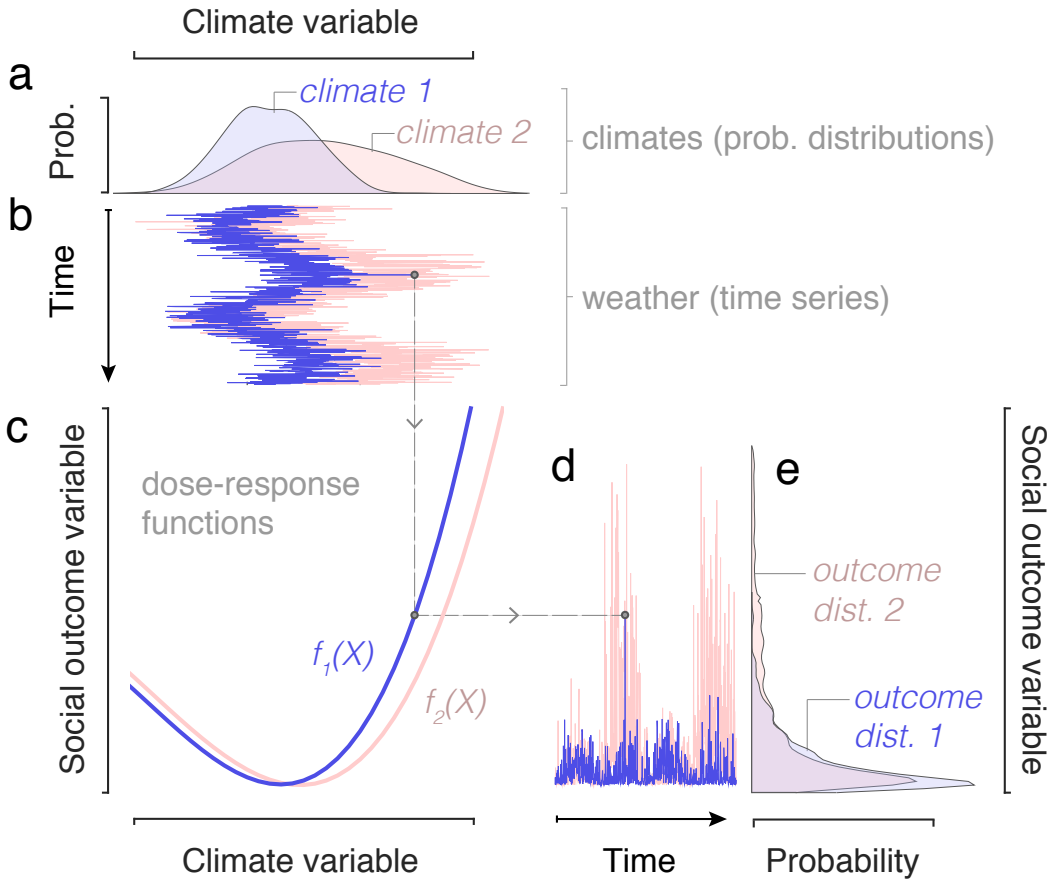

Figure 2.1: Breaking down the influence of climate into analytical components

Climate affects the distribution of social outcomes by altering the distribution of weather events and how populations prepare and respond to these events. (a) Climate is defined as a probability distribution over weather events, such as the distribution *climate_1* (blue) characterizing the probability of the event $climate\_variable = X$, e.g. the likelihood of a rainy day. *climate_2* (pink) characterizes a climate distribution that is shifted to the right and more variable. (b) Weather events over time are realized from each climate, experienced by individuals on the ground, and observed as time series. (c) Statistical analysis recovers "dose-response" functions $f(X)$ that describe social outcomes as a response to each weather "dosage." If populations adapt to their climates (*climate_1* and *climate_2*), then they may respond differently to physically similar weather events, producing dose-response functions that differ (blue=$f_1(X)$, pink=$f_2(X)$)—e.g. if individuals in rainy climates own umbrellas, they may get less wet than populations in normally dry climes (who own few umbrellas) when both populations experience a day with $rainfall = X$. (d) Mapping a sequence of weather events through dose response functions (dashed grey line) generates time series of social outcomes attributable to climatological conditions, accounting both for different distributions of weather events and corresponding adaptations. Signals in an outcome resulting from political, economic, cultural, and other drivers of outcomes might be superimposed on these time series (not shown). (e) Different distributions of expected social outcomes can then be attributable to the two climates (*outcome_distribution_1* and *outcome_distribution_2*), e.g. how much individuals in each climate were soaked by rain over the course of a year.

Mapping climate data onto societies and economies

The first step in analysis is to collapse large quantities of high-dimensional climate data into measures that efficiently summarize the dimensions of climate that are influential on specific aspects of populations. This procedure is challenging because most weather data is collected by physical scientists with the goal of answering physical science questions, so existing structures used to organize these spatially and temporally varying data do not map directly onto social systems. Often, devising a suitable approach for "translating" physical data into a socially-meaningful measure X is the critical innovation that allows researchers to study an entire class of phenomena (Hsiang, 2016). For example, the construction of data describing extreme heat-hours, measured in units of "degree days" and properly aggregated across space, led to strikingly consistent measurement of the surprising effect of temperature on crop yields (Schlenker and Lobell, 2010; Schlenker and Roberts, 2009; Schlenker, Hanemann, and Fisher, 2007; Burke and Emerick, 2016) and electricity demand (Auffhammer and Aroonruengsawat, 2011; Auffhammer, 2018). In another example, tropical cyclone track data were converted into surface wind-exposure of populations to understand the human and economic damage of these storms (Hsiang, 2010). In other work, researchers gain insight from developing new measures of human exposure to the El Niño-Southern Oscillation (Hsiang, Meng, and Cane, 2011), drought indices (Harari and La Ferrara, 2013; Bastos, Busso, and Miller, 2013; Hoddinott and Kinsey, 2001), daily temperature distributions (Deschênes and Greenstone, 2011; Deryugina and Hsiang, 2017), rainfall variability (Fishman, 2016), crop exposure to vapor pressure deficients (Lobell et al., 2013), and trade or neighbor network exposure to multiple variables (Colmer, 2016; Roberts and Schlenker, 2013; Munshi, 2003).

Using research design to identify causal effects

Once societally-relevant measures of climate exposure X are constructed, measuring the causal effect of a weather event on a societal outcome requires that we compare what actually occurred to a counterfactual outcome that would have occurred had the weather been different (Holland, 1986; Hsiang, 2016). For example, simply observing that ten individuals are admitted to a hospital on a hot day does not imply all ten admissions were caused by the heat; it might be the case that nine of those individuals would have gone to the hospital anyway, regardless of the temperature.

In an ideal experiment designed to measure the effect of climate on a social outcome, we would take two populations that are identical in every way and expose one to a "control" climate while exposing the other to a "treatment" climate. The control population serves as the counterfactual for the treatment population and the difference in outcomes would be the effect of the climate treatment. In general, this experiment is infeasible, forcing researchers to rely on "natural experiments" or quasi-experiments.

Early researchers, stretching back to Montesquieu, tried to approximate this ideal experiment, implementing *cross-sectional analyses* where different populations inhabiting different climates are compared to one another and their differences are attributed

to their climates. For example, a researcher might observe that Nigeria has higher crime rates and is hotter than Norway, concluding that higher temperatures lead to crime. This comparison and conclusion are flawed as there are numerous dimensions along which Norway and Nigeria differ—such as geography, history, culture, politics, social institutions—which make Nigeria an unsuitable "treatment" comparison for a Norwegian "control." Some researchers have tried to adjust their analyses to account for important factors known to influence their outcome of interest, but for many complex social outcomes, such as economic growth or civil conflict, it is impossible to know if all relevant factors have been accounted for, and thus unknowable whether a result is credibly causal.

Recent work recognizes this weaknesses of cross-sectional analysis and does not compare different populations to one another. Instead, it leverages the insight that the most comparable group for a certain population is itself, at a moment earlier or later in time. Thus, these *longitudinal studies* follow individual populations over time and examine how they respond to changes in the climatic conditions that they face. When using this approach, researchers have confidence that fundamental factors that influence societies, such as geography and political institutions, are "held fixed" since the population is not changing. In essence, a population just before an event serves as the "control" for that same population right after the event "treatment." Comparing outcomes before and after the climatic event, while accounting for secular trends, provides insight into its effect.

In practice, this approach is complicated by the fact that there are a multiplicity of states for weather and climate, and because societies experience constant variation in both (as suggested by Figure 2.1) it is sometimes difficult to determine if an observed social outcome is the result of current conditions or of climatic events in the past. This challenge is solved by deconvolution of the outcome as a series of responses to continuous climatic conditions. Having observed time series of climatic events or "impulses" (Figure 2.2a) and resulting outcomes (panels b-d), one can search for the characteristic *impulse-response function* that best fits how a single climatic event of unit "dosage" (panel e) generates a response in the outcome (formally, the impulse-response function describes inter-temporal structure of the dose-response function). Figure 2.2f-h displays the characteristic responses that would have been recovered from the different types of outcome data in panels b-d (these simulated responses have been constructed to illustrate three types of real behavior recovered by previous studies; see Figure 2.4).

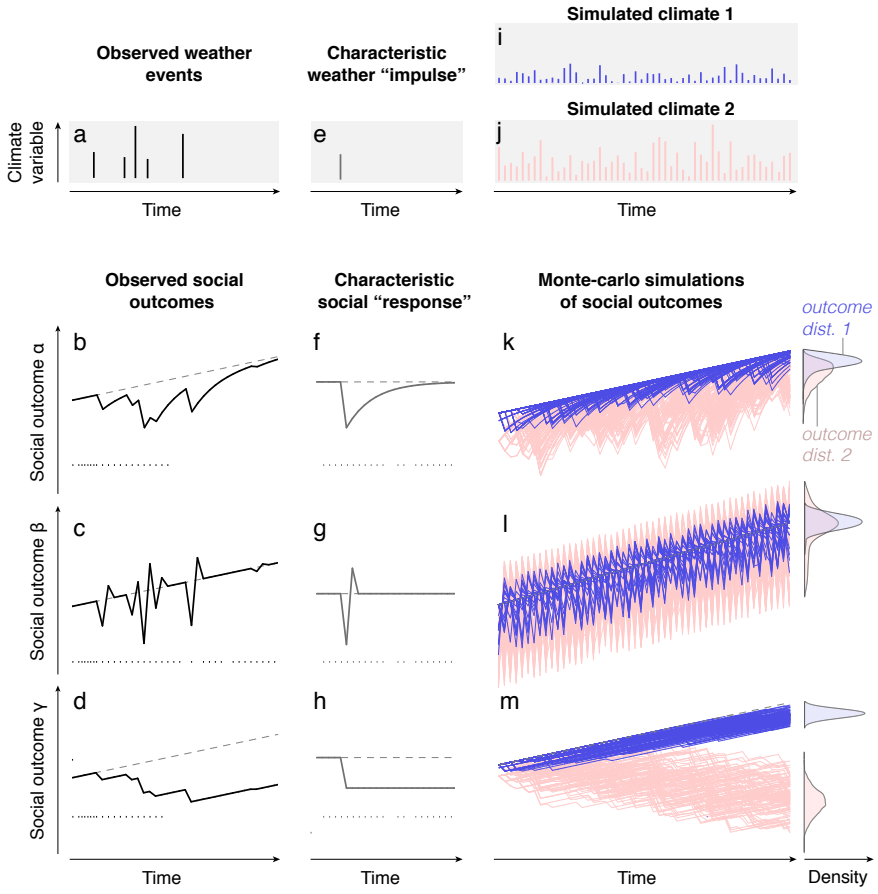

Figure 2.2: The dynamics of societal responses to climate determine how alterations to a climate influence social outcomes

Modern approaches "hold non-climatic factors fixed" by studying a single population over time and identifying social responses to sequential climatic events. Because societal responses may persist (or reverse) after a climatic event ends, continuing through another event that generates another overlapping response, a characteristic impulse-response function can only be recovered from the original data by deconvolution. (a) Time-series of a single population's exposure to weather events each period of magnitude X, indicated as the height of bars (analogous to Figure 2.1b). (b-d) Example time series of three different social outcomes (solid line) that vary relative to baseline trends (dashed line) in response to weather events in (a). (e) A characteristic single weather "impulse" of normalized magnitude. (f-h) Characteristic impulse-response functions describing how each social outcome responds to the weather impulse in (e), recovered from deconvolving data in (a-d). Impulse-responses illustrate different classes of behavior: (f) persistent but decaying effects (e.g. cold-related mortality, see Fig. 2.4a), (g) "temporal displacement" or "harvesting" where delayed responses partially compensate for initial responses (e.g. heat effect on births, see Fig. 2.4c), and (h) permanent effects (e.g. cyclone effects on GDP, see Fig. 2.4d). (i-j) Simulations of weather drawn from two distinct climate distributions. (k-m) Monte-Carlo simulations of social outcomes based on sampling weather from climate distribution 1 (blue, from (i)) and climate distribution 2 (pink, from (j)) and convolving these impulses with the characteristic impulse-response of of each social outcome from (f-h). Distributions of social outcomes under each simulated climate are shown to the right of each panel (analogous to Figure 2.1 e).

Considering the different structures of these responses is important for understanding the response of social systems to different types of climatic factors. For example it has been shown that extreme heat reduces the number of children born exactly nine months later but elevates births eleven and twelve months later, as some of the successful conceptions that would have occurred during the hot period, but did not, end up occurring in the near future (Barreca, Deschenes, and Guldi, 2015). In these cases, where climatic events simply displace the timing of societal outcomes (a pattern illustrated in Figure 2.2g), changes in the distribution of climatic events may have a smaller net effect than one would predict if this dynamic response were not accounted for. Although we do not illustrate it here, it is worth noting that different locations in the dose-response function (Figure 2.1c) may have different dynamics over time (Figure 2.2)—for example cold days cause delayed excess mortality by causing individuals to become ill (analogous to Figure 2.2f), while hot days generate essentially all excess mortality immediately (Deschênes and Moretti, 2009).

Using statistical results to translate climate into outcomes

Once the full structure of a dose-response function, along with its dynamic properties, is identified for a specific population across different weather and climate conditions, researchers can simulate how a population might respond to distributions of weather events that differ slightly from historically experienced distributions (Figure 2.2i-j)—with repeated simulations enabling probabilistic assessment (panels k-m). Gradually distorting the climatological distribution of weather events in such calculations, while adjusting response functions to account for measured patterns of adaptation, allows us to estimate how a shift in the climate may translate into a shift in the distribution of expected social outcomes (Hsiang, 2016).

2.3 Effects of climate on societies

Recent application of the tools described above demonstrate that societies are influenced by the climate in numerous dimensions and at many scales. Individuals face conditions that compromise personal health while entire trade networks or countries can be weakened under adverse climate variation. The linkages between individuals within societal groups can themselves even be fractured by climatic conditions, triggering violence or migrant flows, for example. We review major findings at all these scales, examining effects on human health, economic conditions, social interactions (including violence), and demographic responses (including migration).

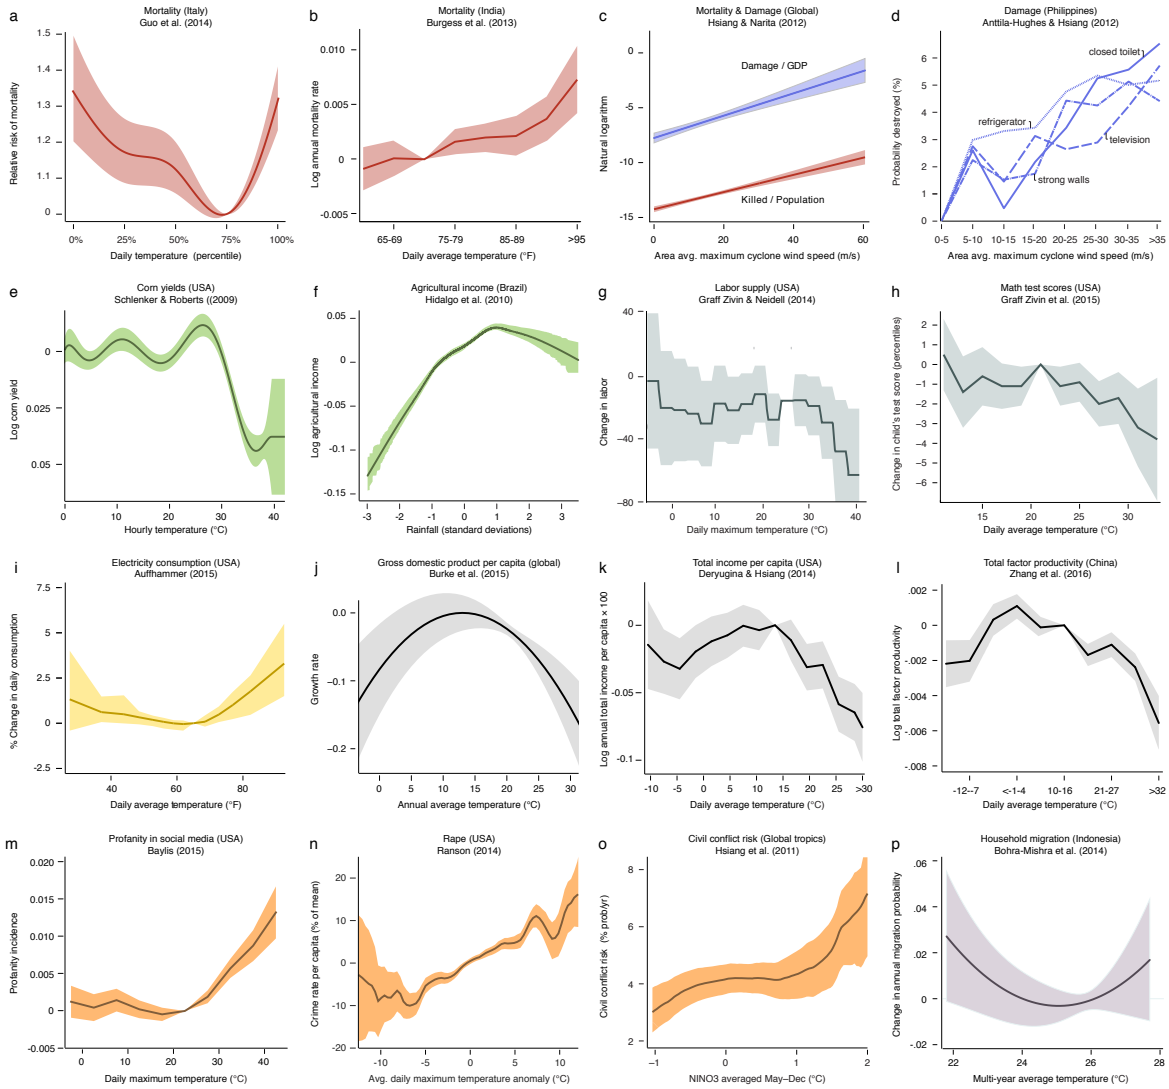

Figure 2.3: Empirical studies demonstrate that climate variables impact social and economic outcomes in many sectors and contexts

(a-p) Examples of dose-response functions estimating the causal effect of climatological events on various social outcomes. Reproduced from authors' original estimation, titles list the outcome variable and location studied. Colors indicate categories of outcome variables: red, mortality (Guo et al., 2014; Burgess et al., 2014); blue, cyclone damage to assets (Hsiang and Narita, 2012; Hsiang and Jina, 2014); green, agriculture (Schlenker and Roberts, 2009; Hidalgo et al., 2010); teal, labor productivity (Graff Zivin and Neidell, 2014; Graff Zivin, Hsiang, and Neidell, 2018); yellow, electricity (Auffhammer, 2018); grey, aggregate economic indicators (Burke, Hsiang, and Miguel, 2015b; Deryugina and Hsiang, 2017; Zhang et al., 2016); orange, aggression, violence, and conflict (Baylis, 2015; Ranson, 2014; Hsiang, Meng, and Cane, 2011; Hsiang, Burke, and Miguel, 2013); purple, migration (Bohra-Mishra, Oppenheimer, and Hsiang, 2014). Climate variables differ by study, but include temperature, cyclone wind speed, rainfall anomalies and ENSO measures. Response functions only identify relative changes and are either normalized to "zero effect" at a designated climatic event, such as a minimum valued outcome, or the sample mean of an outcome. Shaded areas are confidence intervals, as computed by original authors.

Health impacts: Mortality

As individuals, each of us is constantly exposed to temperature, and under extreme heat or cold our bodies struggle to successfully thermoregulate, sometimes leading to severe cardiovascular, respiratory and cerebrovascular effects that can result in death (Basu and Samet, 2002; Deschenes, 2014). Both hot and cold environmental temperatures increase death rates (Figure 2.3a-b): in Delhi, deaths increase by 3.2% per °C above 20°C (Hajat et al., 2005b), and in the U.S. days above 90°F and below 20°F increase male mortality rates by 2% and 1.4%, respectively (Deschênes and Moretti, 2009). Effects of high temperature are rapid and acute but decay quickly, sometimes depressing mortality in following days, as some of the same individuals would have died in subsequent days had an extreme heat event not occurred (Deschênes and Moretti, 2009) (red line in Figure 2.4a). In contrast, cold days have delayed and smaller—albeit enduring—effects lasting up to a month as some individuals become ill, such as contracting influenza, and fail to recover (blue line in Figure 2.4a).

Evidence suggests that adaptations moderate these direct mortality effects. For example, in the USA, mortality from extreme heat declined 80% over the course of the 20th century as air conditioner adoption soared (Barreca et al., 2016) (Figure 2.5d). Remarkably, mortality responses are highly consistent across contexts, when "hot" and "cold" conditions are defined relative to what populations are accustomed to (Guo et al., 2014), suggesting that populations cope with regional climates in a consistent way. Anthropogenic climate change is projected to increase heat-related mortality but decrease cold-related mortality, redistributing mortality rates across locations (Houser et al., 2015), but with an overall net increase in total mortality rates (Burgess et al., 2014; Deschênes and Greenstone, 2011). In a cost-analysis of climate change in the USA, these deaths accounted for the largest share of losses across all impacts (Houser et al., 2015). Effects of humidity exacerbate these patterns (Barreca, 2012; Houser et al., 2015) and mortality in poor agricultural contexts are more extreme (Burgess et al., 2014).

Climatic factors other than temperature also influence mortality. Tropical cyclones directly cause mortality, for example through trauma or drowning, with immediate deaths in storms increasing exponentially with wind speed exposure (Hsiang and Narita, 2012) (Figure 2.3c). Populations regularly exposed to storms appear to adapt somewhat, as their mortality rates are lower than more naive populations when both experience physically comparable events (Hsiang and Narita, 2012) (Figure 2.5a-b). However, these immediate deaths may be minor in magnitude compared to "economic" deaths that occur in the wake of a cyclone (Anttila-Hughes and Hsiang, 2012; Hsiang and Narita, 2012). For example, in the Philippines, changing economic conditions in the years after a cyclone lowers incomes and corresponding spending on food and healthcare, causing mortality among female infants roughly fifteen times higher than direct mortality across all age groups (Anttila-Hughes and Hsiang, 2012) (Figure 2.4b). Extreme rainfall events outside tropical storms also influence mortality – in agriculturally-dependent contexts, infants born in arid areas face elevated risk of death when exposed to droughts (Kudamatsu, Persson, and Strömberg, 2012; Burgess et al., 2014), while

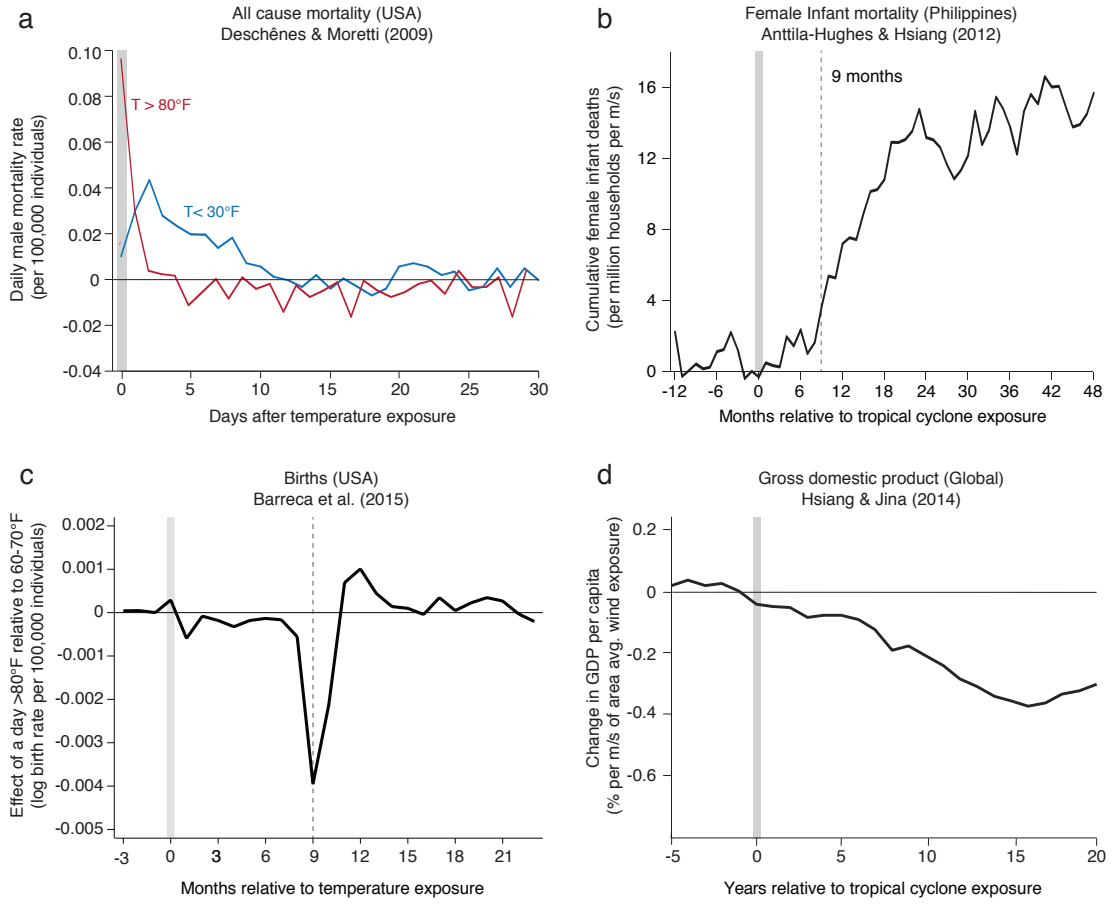

Figure 2.4: Distinct dynamic characteristics of impulse-response functions uncovered in empirical studies

Examples of impulse-response functions from studies identifying dynamic relationships between climate variables and social outcomes, as illustrated heuristically in Figure 2.2e-g. Vertical grey shaded bars indicates the timing of a unit climate "impulse." (a) Male mortality rates in the US increase on both hot and cold days, but hot-day responses rapidly decay and tend to be small and negative for multiple weeks—indicating temporal displacement—while cold days generate a more gradual and enduring mortality effect (Deschênes and Moretti, 2009). (b) Tropical cyclones increase female infant deaths, but with a delayed effect that grows rapidly roughly a year after exposure (Anttila-Hughes and Hsiang, 2012). (c) Birth rates in the US fall 8-10 months after a hot day, but this decline is partially compensated for by an increase during months 11-13 (Barreca, Deschenes, and Guldi, 2015). (d) GDP in countries exposed to tropical cyclones falls gradually but persistently during the 15 years following the exposure (Hsiang and Jina, 2014).

flooding has been linked to death throughout Europe (Hajat et al., 2005a).

Health impacts: Morbidity

Many injuries to human health caused by climate are non-fatal. One means of detecting these effects is to measure the impact of climatic events on hospital admissions. Admissions for respiratory and cardiovascular diseases respond to temperature similarly to mortality, with impacts at both high and low daily temperatures (Ziebarth, Schmitt, and Karlsson, 2014; Kovats, Hajat, and Wilkinson, 2004). The precise spatial and temporal resolution of these hospital- or city-level studies allows authors to account for key temperature correlates, such as air pollution and humidity, which also influence hospitalizations. This adjustment is important, as failing to account for particulate matter and ozone may exaggerate the effect of temperature by up to a factor of two (Ziebarth, Schmitt, and Karlsson, 2014; O'Neill et al., 2005). Even without hospital-level data, evidence using cause-of-death records can illuminate key morbidity effects; for example, humidity is an important driver of influenza, a significant cause of hospitalization and mortality in temperate climates (Barreca and Shimshack, 2012).

A major component of morbidity affected by the climate is vector-borne disease. For example, malaria and dengue fever infect approximately 200 million and 50 million people globally each year (World Health Organization, 2014), respectively, and the life-cycles of mosquito vectors transmitting these illnesses are strongly influenced by climate. Temperature nonlinearly influences the reproduction of parasites, extreme temperatures lower mosquito survival rates, and open water critical for mosquito breeding is constrained by rainfall (Craig, Snow, and Le Sueur, 1999; Gething et al., 2011). These climatic factors affect the intensity of infection in areas where malaria and dengue are already endemic (Zhou et al., 2004) as well as affecting where the disease may spread to (Bhatt et al., 2013). These dynamics make measurement of climate-disease interaction challenging: some studies aim to recover incidence as nonlinear functions of temperature and rainfall (Barreca, 2010; Colón-González et al., 2013), while others parameterize ecological models of vector transmission, using model output as indices to predict cases with data (McCord, 2016) or simulation (Small, Goetz, and Hay, 2003). Anthropogenic climate change is likely to shift disease ranges and increase exposure globally, but changing temperatures, rainfall, and intervention strategies complicate projections (Gething et al., 2010; Bhatt et al., 2013); more research in this area is needed to link climate, ecological models, and social data.

Health impacts: Early life

Climatic conditions experienced during early stages of life can have outsized impact because altered early development affects long-run health and well-being (Almond and Currie, 2011). For example, in-utero exposure to high temperatures can lower birth weight (Deschênes, Greenstone, and Guryan, 2009) and exposure to tropical cyclones leads to a variety of birth complications (Currie and Rossin-Slater, 2013). Mechanisms explaining these in-utero effects remain elusive, as it is challenging to separate effects on gestational length and nutrient accumulation (Deschênes, Greenstone, and Guryan, 2009), and because climate shocks occurring at different points in the gestational period

likely operate through distinct channels. For example, high temperatures at conception lead to fetal losses that, through selection, improve outcomes for babies who do survive (Wilde et al., 2014), while high temperatures in the third trimester have unambiguously negative impacts (Isen, Rossin-Slater, and Walker, 2017).

Regardless of mechanism, in-utero health insults have later-life economic consequences, such as lowered income (Fishman, 2016; Isen, Rossin-Slater, and Walker, 2017). In developing country contexts, adverse rainfall in the year of birth lowers adult female health outcomes and educational attainment (Maccini and Yang, 2009) and droughts experienced as toddlers lowers childhood growth and education (Hoddinott and Kinsey, 2001; Alderman, Hoddinott, and Kinsey, 2006). These rainfall-related impacts likely operate through agricultural income loss and lowered nutrition; however, our understanding of these channels is generally weak and work is needed to parse out direct physiological impacts from economic factors and behavioral responses.

Economic impacts: Agricultural yields

Study of the direct effect of climate on economic outcomes began in agriculture, where the importance of climatic factors is clearest (Auffhammer and Schlenker, 2014). Despite centuries of agricultural experience, a surprising recent finding is the importance of temperature, often dominating rainfall, in the production of staple crops (Lobell and Burke, 2008; Schlenker, Hanemann, and Fisher, 2005; Schlenker and Roberts, 2009; Auffhammer, Ramanathan, and Vincent, 2012). Highly nonlinear yield losses on the hottest days drives much of this effect (Schlenker and Roberts, 2009) (Figure 2.3e), a relationship recovered in the US (Schlenker and Roberts, 2009), Africa (Schlenker and Lobell, 2010), Europe (Moore and Lobell, 2015), Southeast Asia (Welch et al., 2010), and India (Guiteras, 2009; Burgess et al., 2014). Crops are most sensitive to temperatures during specific phases of the growth cycle (Welch et al., 2010; Auffhammer, Ramanathan, and Vincent, 2012). While temperature impacts generally outweigh those of rainfall, low and very high total seasonal rainfall levels do damage yields in many contexts (Schlenker and Roberts, 2009; Fishman, 2016) (Figure 2.3f), an effect that is partially attenuated when water storage and irrigation are widely available (Auffhammer, Ramanathan, and Vincent, 2012; Fishman, 2011; Welch et al., 2010; Duflo and Pande, 2007). Similarly, within a single growing season, farms that experience a small number of extremely rainy days suffer damaged yields, relative to the same quantity of rain distributed evenly across growing days (Fishman, 2016). These various dose-response functions have been recovered and replicated for major global crops like maize, rice, soy, and wheat, but less is known about effects on regional crops like millet and cassava—which can be critical in poor rural regions—and specialty crops like fruits and vegetables, with some notable exceptions (Schlenker and Lobell, 2010; Wineman and Mulenga, 2014; Lobell et al., 2008). A body of research in dairy science has established that both temperature and humidity nonlinearly impact milk yields (André et al., 2011; Bryant et al., 2007; Key, Sneeringer, and Marquardt, 2014) while linearly lowering cattle pregnancy rates (Amundson et al., 2005), but little is known outside of highly managed livestock operations in industrialized countries.

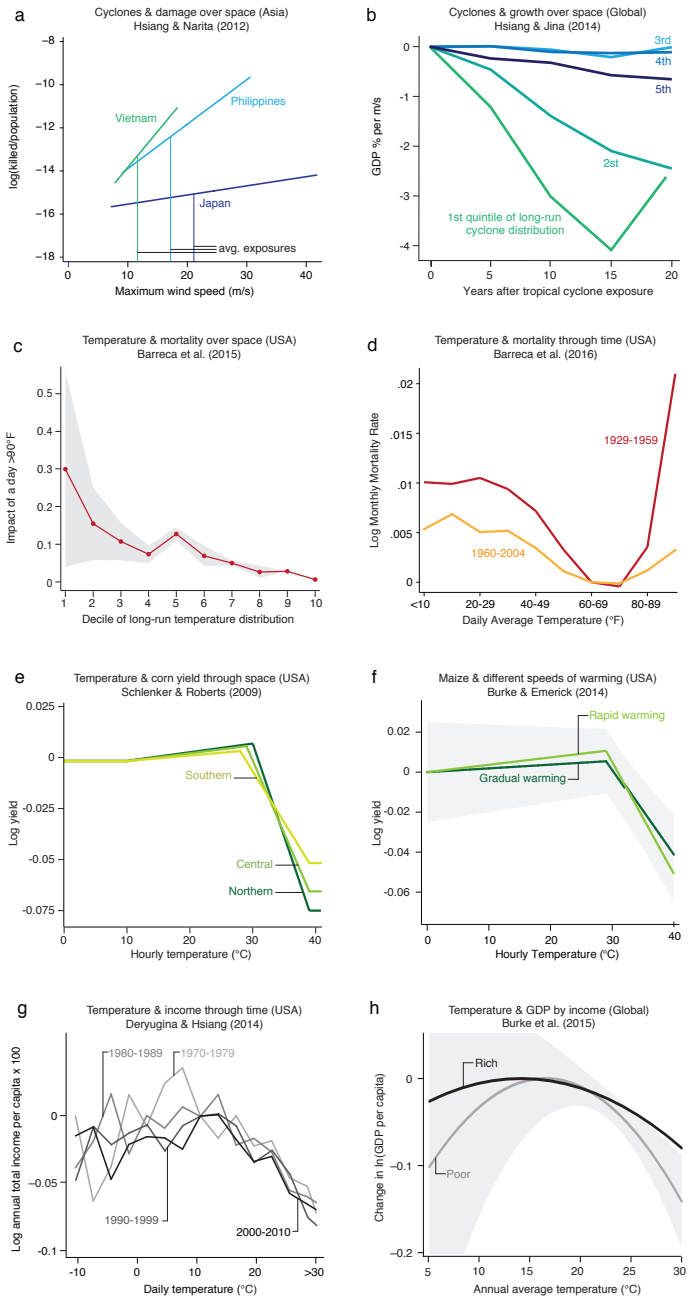

Figure 2.5: Different responses to physically similar events indicates the presence or absence of effective adaptations

Comparison of response functions over time and across space can indicate where populations successfully adapted and where an "adaptation gap" might persist. Global cyclone losses indicate adaptation: (a) mortality rates increase with cyclone intensity faster in countries where average exposure (thin vertical line) is lower (Hsiang and Narita, 2012) and (b) effects of cyclones on GDP over time are most negative in countries with the lowest levels of historical experience—rank indicates quintile of exposure (Hsiang and Jina, 2014). Temperature-induced mortality in the USA exhibits adaptation: (c) locations that with hotter long-run climates have smaller effects (Barreca et al., 2015) and (d) sensitivities have declined significantly over time (Barreca et al., 2016). Maize yields in the USA indicate limited adaptation: (e) hot and cool climates exhibit similar effects of heat (Schlenker and Roberts, 2009), and (f) yields are equally impacted by rapid and slow changes in temperature (Burke and Emerick, 2016). Aggregate income exhibits limited adaptation: (g) county-level losses from high temperatures have not changed over time in the USA (Deryugina and Hsiang, 2017), and (h) country-level GDP reductions are only slightly less severe in rich nations than poor countries (Burke, Hsiang, and Miguel, 2015b). Shaded areas are confidence intervals, as computed by original authors.

Effective adaptation to climate in agriculture appears modest, as dose-response functions change little across time and space (Schlenker and Roberts, 2009; Schlenker, Roberts, and Lobell, 2013) (Figure 2.5e), even when warming effects are gradual (Burke and Emerick, 2016) (Figure 2.5f). Furthermore, large but temporary climate events, like the US Dust Bowl, have had persistent multi-decadal impacts on farm values (Hornbeck, 2012). These findings contrast with historical narratives of farmer adaptability, such as the 200-year-long spread of agriculture into previously non-arable land (Olmstead and Rhode, 2011a,b) and adjustment of cultivars in response to drought (Sutch, 2011). These two views of agriculture adaptability remain unreconciled and identifying obstacles to adaptation, such as poor incentives (Annan and Schlenker, 2015) or high adaptation costs (Schlenker, Roberts, and Lobell, 2013), are a critical area for future research.

Economic impacts: Labor supply and productivity

Agricultural effects cannot explain many patterns in the overall economic response to climate, leading to the hypothesis that effects on labor are another important channel of influence (Hsiang, 2010). A growing body of evidence now supports this theory (Heal and Park, 2015). Heat stress can lower work intensity (Seppanen, Fisk, and Lei, 2006a), reduce cognitive performance (Graff Zivin, Hsiang, and Neidell, 2018), and voluntarily shorten work hours in sectors of the economy most exposed to outdoor temperature, such as construction and agriculture (Graff Zivin and Neidell, 2014) (Figure 2.3g-h). Impacts on manufacturing production have been identified in both high and low income contexts (Somanathan et al., 2015; Adhvaryu, Kala, and Nyshadham, 2014), although understanding the full impact of this effect is made challenging by reallocation of labor within an economy (Colmer, 2016). Patterns in the overall macro-economic responses to temperature (discussed below) are consistent with labor effects playing an important role (Hsiang, 2010; Deryugina and Hsiang, 2017; Burke, Hsiang, and Miguel, 2015b) (Figure 2.3j-l), where individuals are each affected modestly but a large number of affected individuals might generate substantive aggregate impacts on output, and possibly on growth (Burke, Hsiang, and Miguel, 2015b). Theory suggests that labor productivity losses might be exacerbated by market reactions that reduce the intensity of labor used in economic activities (Heal and Park, 2013) and slow downstream production (Wenz and Levermann, 2016). Investments in climate control for work environments can offset some of these labor productivity effects (Somanathan et al., 2015), but at substantial cost, such as expenditures on energy.

Economic impacts: Energy supply & demand

The relationship between climate and energy is unique. Energy systems are directly impacted by climate—high temperatures provoke demand surges while straining supply and transmission—and they also serve a critical role supporting adaptation by enabling cooling, heating, irrigation, trade, etc.; while simultaneously being the the largest contributor to anthropogenic climate change.

The effect of temperature on energy demand is highly nonlinear. Households and firms use energy heavily for indoor climate, based on the weather and available infrastructure (Auffhammer and Mansur, 2014; Cian, Wing et al., 2014). Almost universally, energy demands fall with rising cool temperatures and increase steeply at high temperatures, leading to a U-shaped relationship (Figure 2.3i) (Deschênes and Greenstone, 2011; Auffhammer and Aroonruengsawat, 2011; Davis and Gertler, 2015). Investments in new energy-intensive infrastructure, such as heaters, may respond to climate more slowly as households and industry adopt expensive technology based on their beliefs about their climates. Evidence from the US (Sailor and Pavlova, 2003), Mexico (Davis and Gertler, 2015), and China (Auffhammer, 2014) indicates that electricity demands on hot days rise fastest in locations that tend to be hot, presumably because more buildings in these locations have air conditioners that are all utilized simultaneously on hot days.

Engineering models and simple thermodynamics suggest electricity supply and transmission systems should suffer efficiency losses at high temperatures (Jaglom et al., 2014), but these effects are empirically challenging to measure in the presence of fluctuating demand. Evidence indicates river water temperatures can influence electricity prices (McDermott and Nilsen, 2014), nuclear power capacity utilization may fall with high temperature (Linnerud, Mideksa, and Eskeland, 2011), and droughts can shift generation away from hydropower and toward carbon-intensive fuel sources (Muñoz and Sailor, 1998; Eyer and Wichman, 2014), but it is unclear whether these findings generalize.

Projections under climate change generally indicate energy demand will grow on net, even though fewer days will require energy for heating. Sensitivity to high temperatures will likely grow as air conditioner use expands due to improvements in technology, rising incomes, and investments specifically motivated by warming (Auffhammer, 2014; Davis and Gertler, 2015). These investments may affect energy prices by substantially elevating peak demand (Houser et al., 2015), but better understanding of these issues is required to support long-term energy planning.

Economic impacts: Trade

The current structure of the global economy represents a spatial equilibrium in which the location of populations and sites of economic production are all determined by the functioning and friction of markets through which individuals trade with one another and the factors that make locations more or less productive. Analyses of climatic influence on migration can be interpreted as a reallocation of labor across these locations, perhaps in response to changing economic conditions, which we discuss below. Yet given an approximately fixed distribution of populations across locations, climate may also affect how populations decide to trade with one another. For example, global wind patterns and ocean currents have strongly influenced patterns of trade historically because of the role these factors play in the cost of shipping along different routes (Feyrer and Sacerdote, 2009; Kaluza et al., 2010). High temperatures that reduce productivity lower the quantity of goods exported from a country, both in agriculture (Roberts and Schlenker, 2013) and manufacturing (Jones and Olken, 2010), and cyclone strikes

that lower national incomes tend to reduce imports (Hsiang and Jina, 2014). In large integrated trade networks, the spatial distribution of climatic conditions can affect market prices (Kazianga and Udry, 2006; Houser et al., 2015), presumably through effects on both supply costs and demand, and should theoretically determine the location of different economic activities (Desmet and Rossi-Hansberg, 2015; Costinot, Donaldson, and Smith, 2016).

These reallocations across space and time can, in some contexts, mitigate the direct damages of climate. For example, outdoor labor supply shifts to cooler hours of the day during heat shocks (Graff Zivin and Neidell, 2014), water storage weakens the link between rainfall and agricultural productivity (Duflo and Pande, 2007), unskilled labor moves from agriculture to manufacturing when crops are hit by high temperatures (Colmer, 2016), and grain inventories adjust to smooth weather impacts on farm profits (Roberts and Schlenker, 2013). However, these adjustments may be limited – historical evidence of inter-temporal substitution is minimal for aggregate incomes (Deryugina and Hsiang, 2017; Burke, Hsiang, and Miguel, 2015b) and cyclone damages (Hsiang and Jina, 2014), and in the future sequential periods of similar extreme conditions may make such reallocations over time more difficult. Reallocation across space may also be constrained in the future – current simulations disagree as to whether adjustment of trade patterns under climate change will dampen or amplify its overall social costs (Costinot, Donaldson, and Smith, 2016; d'Amour et al., 2016; Houser et al., 2015). Investigation of substitution patterns across both space and time is a key area for future work.

Economic impacts: Economy-wide effects

Rather than examining individual or sectoral responses to climate, an alternative "top down" approach examines how the macro-economy as a whole responds to climatic conditions. This approach is usually implemented by examining total income or Gross Domestic Product (GDP) per capita as the outcome of interest. Recent work has shown that low rainfall slows national incomes greatly in Africa (Miguel, Satyanath, and Sergenti, 2004; Barrios, Bertinelli, and Strobl, 2010), ENSO modulates a see-saw-like oscillation in total agricultural income between tropical and temperate countries (Hsiang and Meng, 2015), tropical cyclone strikes slow GDP growth for roughly 15 yrs in proportion to the intensity of the storm (Hsiang and Jina, 2014) (Figure 2.4d), and temperatures have a nonlinear effect on economic production, such that output is maximized around 13°C (Burke, Hsiang, and Miguel, 2015b) (Figure 2.3j-l). The roughly linear effects of cyclones and nonlinear effects of temperature at the macro-level are fully consistent with the structure of effects measured in micro-level analyses (Anttila-Hughes and Hsiang, 2012; Deryugina and Hsiang, 2017; Colacito, Hoffmann, and Phan, 2014; Zhang et al., 2016). Determining the persistence of these GDP losses is important because enduring losses may accumulate and compound, leading to larger long-run losses (Burke, Hsiang, and Miguel, 2015b; Hsiang and Jina, 2014)—this could occur if climatic events alter investment behavior (Burke, Hsiang, and Miguel, 2015b) or capital depreciation (Houser et al., 2015; Hsiang and Jina, 2015). However, existing

data and approaches have had difficulty constraining the overall persistence of these effects (Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015b).

Perhaps remarkably, effects of temperature and cyclones are globally generalizable in the sense that they have be recovered using subsamples of data from around the world, including both rich and poor countries (Hsiang and Jina, 2014; Burke, Hsiang, and Miguel, 2015b; Deryugina and Hsiang, 2017; Zhang et al., 2016). Early analyses focused on large negative effects of temperature on GDP in poor countries (Dell, Jones, and Olken, 2012; Hsiang, 2010), although later studies demonstrated that almost identical responses appeared in rich countries as well (Burke, Hsiang, and Miguel, 2015b; Deryugina and Hsiang, 2017) (Figure 2.5h). This finding—in conjunction with the results that the effects of temperature on income in the USA remained essentially unchanged from 1960 to 2010 (Burke, Hsiang, and Miguel, 2015b; Deryugina and Hsiang, 2017) (Figure 2.5g) and gradual warming has identical effects as short-lived warming (Dell, Jones, and Olken, 2012)—leads naturally to the conclusion that effective adaptation to temperature, at the macro-level, is limited. Across a variety of contexts, once temperatures are higher than the optimum, each increase in temperatures by 1°C lowers economic production roughly 1-1.7%. The single finding that suggests some effective adaptation at the macro-level is that cyclone-prone countries experience GDP losses (per cyclone) much smaller than countries where storms are infrequent (Hsiang and Jina, 2014).

Social interactions: Women and girls

Under economic pressure from climate, the terms and bargaining positions in personal relationships may change. These bargaining interactions are often gender-based, causing women and girls to experience these changes differently. For examples, in Sub-Saharan Africa, evidence suggests some women suffering income shortfalls during drought engage in "transactional" intercourse, leading to increased probability of HIV infection (Burke, Gong, and Jones, 2015); in the Philippines, female infants conceived after a tropical cyclone have elevated risk of mortality (Figure 2.4b), particularly if they have older brothers (Anttila-Hughes and Hsiang, 2012); and in Indonesia, girls born in drought years exhibit lower long-run to health and education, as diminished family resources are more often allocated towards investment in boys (Maccini and Yang, 2009).

Social interactions: Interpersonal violence and aggression

Evidence from numerous contexts repeatedly finds that interpersonal violence increases with temperatures and sometimes low rainfall (Hsiang, Burke, and Miguel, 2013; Burke, Hsiang, and Miguel, 2015a). This response manifests in low-level aggression, such as horn honking (Kenrick and Macfarlane, 1986), antisocial behavior toward service employees (Kolb, Gockel, and Werth, 2012), the use of profanity in social media (Baylis, 2015) (Figure 2.3m), as well as in outright violence, such as retaliation in sports (Larrick et al., 2011) and violent crimes: rape, murder, robbery and assault (Ranson, 2014;

Jacob, Lefgren, and Moretti, 2007; Basyan et al., 2014; Blakeslee and Fishman, 2017) (Figure 2.3n). The effect of temperature is strikingly linear with almost no delay, suggesting it might be driven by a physiological mechanism (Shaun Morrison, 2008; Ray et al., 2011; Seo, Patrick, and Kennealy, 2008). Effects of rainfall on interpersonal violence appear primarily in some poor agricultural contexts, such as rural India (Blakeslee and Fishman, 2017; Iyer and Topalova, 2014; Sekhri and Storeygard, 2012) and Tanzania (Miguel, 2005), suggesting that damage to agricultural yields may be a mediating factor.

Social interactions: Intergroup violence

Climatic conditions also influence relationships between groups, changing the risk of large-scale conflict (Hsiang, Burke, and Miguel, 2013; Burke, Hsiang, and Miguel, 2015a). Cold events during cold epochs, such as feudal Europe and dynastic China (Zhang et al., 2006, 2007; Tol and Wagner, 2010; Anderson, Johnson, and Koyama, 2013), or periods of low rainfall (Jia, 2014; Kung and Ma, 2012; Bai and Kung, 2010), produced instability and upheaval—probably related to crop failures. During the modern warm period, hotter conditions increase collective violence in settings as diverse as insurgency in India (Fetzer, 2014), land invasions in Brazil (Hidalgo et al., 2010), and civil war intensity in Somalia (Maystadt and Ecker, 2014). This relationship is linear, with violence rising roughly 11% per standard deviation in temperature, exhibits some evidence of adaptation through rising incomes (Carleton, Hsiang, and Burke, 2016), and has an unknown mechanism (Burke, Hsiang, and Miguel, 2015a). Rainfall extremes also increase intergroup conflict in agricultural contexts (Hidalgo et al., 2010; Miguel, Satyanath, and Sergenti, 2004; Harari and La Ferrara, 2013; Ralston, 2015), as does El Niño (Hsiang, Meng, and Cane, 2011) (Figure 2.3o).

Social interactions: Institutional breakdown & state failure

Governing institutions may falter under sufficiently strong climatological stress. Patterns such as the forcible removal of rulers (Kim, 2016; Burke and Leigh, 2010; Burke, 2012; Chaney, 2013) can be tied to fluctuations in climate, but attributing societal collapse to climate is more difficult because there are fewer events. Nonetheless, several historical cases are compelling, such as the collapse of the Akkadian (Cullen et al., 2000), Mayan (Haug et al., 2003), and Angkor (Buckley et al., 2010) empires, dynastic changes in China (Yancheva et al., 2007), and major transitions in Europe (Büntgen et al., 2011).

Demographic effects: Migration

Human mobility is likely an important strategy to cope with climatic changes, but it is challenging to characterize as climate appears to have two opposing influences: deteriorating economic conditions and safety motivate migration while simultaneously undercutting household resources needed to migrate (Kleemans, 2014; Cattaneo and

Peri, 2016). Net effects are mixed; for example, urbanization and outmigration from agriculturally-dependent areas may increase as temperatures hit crop-damaging levels and moisture declines (Hornbeck, 2012; Feng, Krueger, and Oppenheimer, 2010; Feng, Oppenheimer, and Schlenker, 2012; Henderson, Storeygard, and Deichmann, 2014; Bohra-Mishra, Oppenheimer, and Hsiang, 2014; Cai et al., 2014) (Figure 2.3p), but non-agricultural workers in Mexico move in response to temperature more rapidly than farm laborers (Jessoe et al., 2014) and some of the poorest countries show no emigration response(Cattaneo and Peri, 2016). In Africa, flows from urban to foreign locations appear responsive to weather (Marchiori, Maystadt, and Schumacher, 2012) but US-bound migration from urban Mexico is unaffected by heat waves (Nawrotzki et al., 2015). Climatological natural disasters that influence incomes, such as hurricanes and flooding, appear to have limited impact on total migration in low income contexts (Bohra-Mishra, Oppenheimer, and Hsiang, 2014; Drabo and Mbaye, 2011; Gray and Mueller, 2012) and cause simultaneous inflow and outflow of migrants in the US (Strobl, 2011; Deryugina, 2017). Overall, the wide-ranging climatic effects on migration are not well understood and remain an area of active investigation.

Demographic effects: Population structure & growth

Because climatic events affect subgroups within a population differently, such as women or the poor, it is thought that repeated exposure of the population may gradually distort its demographic structure. For example, recent findings suggest that male fetuses are less likely to survive challenging climatic events, such as extreme heat, leading to disproportionately female cohorts of surviving infants born just following hot years (Fishman, Russ, and Carrillo, 2015; Wilde et al., 2014). Demographic distortions may also occur through nonfatal mechanisms, such as the disproportionate migration of wealthy older individuals away from US counties struck by cyclones simultaneous with the movement of young and low income individuals into these same counties (Strobl, 2011; Deryugina, 2017). These seemingly small individual effects might grow to be substantial following repeated exposure, but the full scale and scope of climatological influence on equilibrium demographic structure remains unknown.

New findings also suggest that overall population growth may be directly influenced by the climate through altering sexual behavior or fertility rates. Birth rates are abnormally lower nine months following extreme heat events in both Sub-Saharan Africa (Wilde et al., 2014) and the US (Barreca, Deschenes, and Guldi, 2015) (Figure 2.4c), although identifying the mechanism driving this effect is challenging. Remarkably, these results appear to explain a large fraction of birth seasonality across climates and projections for the US suggest warming will reduce birth rates 3% (Barreca, Deschenes, and Guldi, 2015).

2.4 Attributing current and future effects of climate

The results above describe the structure of the dose-response functions that govern how populations respond to individual climatic events, where these relationships were isolated from data containing overlapping signals of numerous sequential climatic events. By mapping distributions of multiple climatic events back onto these empirically recovered dose-response functions, we can reconstruct distributions of predicted outcomes attributable to these weather distributions (as illustrated in Figures 2.1-2.2). Comparison of outcome distributions resulting from different climatologies allows us to estimate the first-order effects of any arbitrary change in the climate (Hsiang, 2016). In principle, with sufficient information on patterns of adaptation to climate, i.e. the "informational" channel that caused the dose-response function in Figure 2.1c to change, these comparisons can account for the full range of adaptations observed in the real world; although in practice such adjustments tend to be relatively minor (Hsiang and Narita, 2012; Schlenker, Roberts, and Lobell, 2013; Houser et al., 2015; Auffhammer, 2018), due in part to the fact that they are mathematically second-order (Hsiang and Narita, 2012; Hsiang, 2016) and in part to the observation that the informational effect tends to be modest in magnitude across numerous contexts (Ranson, 2014; Deryugina, 2017; Lobell et al., 2014; Burke and Emerick, 2016), especially once the costs of adaptive adjustments are accounted for (Schlenker, Roberts, and Lobell, 2013; Hsiang, 2016).

By using this approach to "reconstitute" distributions of impacts from climate, researchers are now beginning to provide first-order answers to three questions that originally motivated this research agenda: How much does the current climate affect outcomes that we observe in the current world? How much has recent warming affected outcomes? And how are projected changes in the climate expected to alter social outcomes?

Study	Social impact	Sample region	Sample period	Effects of current climate distribution	Effects of climate trends to date	Future impacts of climate change
Agriculture						
Auffhammer et al. 2012	Rice yield	India	1966-2002		Between 1966 and 2002, trends in temperature, monsoon characteristics and rainfall lowered yields by **5.7%** on average	
Lobell & Field 2007	Major crop yields	Global	1961-2002		By 2002, trends in temperature since 1981 caused annual losses of **40 megatons or $5 billion**	
Lobell et al. 2011	Major crop yields	Global	1960-2008		Between 1980 and 2002, trends in temperature and precipitation lowered maize and wheat yields by **3.8% and 5.5%**; rice and soy were unaffected	
Schlenker & Lobell 2010	Major crop yields	Sub-Saharan Africa	1961-2007			Predicted climate change$^{\diamond}$ by 2050 lowers annual yields by **22%** for maize, **17%** for sorghum and millet, **18%** for groundnut, and **17%** for cassava
Schlenker & Roberts 2009	Maize yield	Eastern USA	1950-2008	Relative to an optimal season at 29°C, realized temperatures lower annual yields by **48%** on average ($*$)		Predicted climate change$^{\times}$ by 2100 lowers annual yields by **63-82**%
Welch et al. 2010	Rice yield	South Asia	1979-2004		Between 1979 and 2004, trends in temperature and solar radiation lowered yield growth by **0-0.76%**	

Continued on next page...

Table 2.1: Attribution of climate impacts

Study	Social impact	Sample region	Sample period	Effects of current climate distribution	Effects of climate trends to date	Future impacts of climate change
Income						
Burke et al. 2015	Income	Global	1960-2010	Relative to each country's optimal annual temperature, realized temperatures lower the annual global growth rate by **0.25 percentage points** on average (∗)	Between 1980 and 2010, trends in temperature lowered the annual global growth rate by **0.002 percentage points** on average (∗)	Predicted climate change† by 2100 lowers global GDP by **23%** and between 2010 and 2100 lowers the global annual growth rate by **0.28 percentage points** on average
Deryugina & Hsiang 2015	Income	USA	1969-2011	Relative to each county's optimal annual temperature, realized county temperatures lowered the U.S. growth rate between 1970 and 2011 by **1.69 percentage points** on average		Predicted climate change† by 2100 lowers the U.S. annual growth rate by **0.06-0.16 percentage points**
Hsiang & Jina 2014	GDP growth	Global	1950-2008	Relative to a world without cyclones, realized cyclones lowered the global annual growth rate between 1970 and 2008 by **1.27 percentage points**		Predicted climate change◇ by 2090 induces damages of **$9.7 billion**
Zhang et al. 2016	Total factor productivity	China	1998-2007	Relative to a full year at 50-60°F, realized temperatures lower TFP by **31%** on average (∗)		Predicted climate change× by 2050 lowers annual TFP by **4.18%**

Continued on next page...

Table 2.1: Attribution of climate impacts (continued)

Study	Social impact	Sample region	Sample period	Effects of current climate distribution	Effects of climate trends to date	Future impacts of climate change
Health						
Anttila-Hughes & Hsiang 2012	Mortality rate, total deaths	Philippines	1950-2008	Realized typhoon-induced "economic" deaths account for **13%** of the overall infant mortality rate		
Burke et al. 2015	HIV rate	Sub-Saharan Africa	2003-2009	Rainfall shocks account for **14-21%** of cross-country variation in HIV prevalence		
Deschenes & Greenstone 2011	Mortality rate, energy use	USA	1968-2002	Relative to a full year at 50-60°F, realized temperatures increase mortality rates by **11.2%** and energy use by **29%** on average (∗)		Predicted climate change$^{\times}$ by 2100 increases annual mortality rates by **3%** and energy use by **11%**
Conflict						
Burke et al. 2009	Civil conflict	Sub-Saharan Africa	1981-2002	Relative to each country's optimal annual temperature, realized temperatures increase annual incidence of war by **43.6%** on average (∗)	Between 1981 and 2006, trends in temperature increased the annual incidence of war by **15.9%** on average (∗)	Predicted climate change$^{\diamond}$ by 2030 increases annual incidence of war by **54%**
Hsiang et al. 2011	Civil conflict	Global	1950-2004	Relative to the optimal state, realized ENSO conditions had a role in **21%** of all civil conflicts between 1950 and 2004		
Ranson 2014	Violent crime	USA	1980-2009	Relative to each county's optimal monthly temperature, realized temperatures increase crime rates by **6.1%** for rape, **2.4%** for murder, and **3.6%** for aggravated assault on average (∗)		Predicted climate change$^{\diamond}$ between 2010 and 2099 increases total crime cases by **180,000** for rape, **22,000** for murder, and **2.3 million** for aggravated assault

(∗) New calculation generated either from reanalysis of the authors' data, or from analysis of summary statistics provided in the authors' paper. See Supplementary Materials for detailed descriptions of each calculation. Climate change impacts are predicted using the Intergovernmental Panel on Climate Change A1F1 ($^{\times}$), A1B ($^{\diamond}$), or RCP 8.5 (†) future climate change scenarios.

Table 2.1: Attribution of climate impacts (continued)

The current climate

Most analyses do not explicitly report how much the distribution of a social outcome examined is driven by climatic factors, but such results are implicitly computed and relied upon in every deconvolution or regression analysis, and estimating the total effect of current climate distributions provides perspective on the overall magnitude of contemporary social impacts of climate. In column 5 of Table 2.1 we tabulate estimates from studies that do report such results as well as compute some new estimates based on reported values and available data. To compute the total effect of the current climate, one can use the sample of data analyzed and the empirical relationship recovered by the analysis to (i) compute the distribution of outcomes predicted by the current distribution of climatic events, and (ii) compare this to the distribution of outcomes obtained if the same population were exposed to their best possible environmental conditions continuously, where "best possible" is based on the nature of the estimated empirical relationship (see Appendix A for details). Essentially, to create this benchmark we imagine a world in which climate could be managed as other inputs to societies and economies are, such as the allocation of law enforcement or capital investments. For example, in their analysis of the effect of ENSO on civil conflict, Hsiang, Meng, and Cane (2011) estimate average conflict rates predicted by historical ENSO conditions and compare them to conflict rates that would be predicted if the world were to experience La Niña-like conditions, the climate state with least conflict, continuously. This thought-experiment is clearly unrealistic, as societies cannot uniformly be exposed to an optimal climate; however, it is a useful and precisely defined benchmark for considering the overall magnitude of effects resulting from observed climates.

In general, modern climates have substantial influence on social and economic outcomes. For example, historical temperatures in the United States are estimated to currently suppress maize yields by roughly 48% relative to ideal growing conditions (Schlenker and Roberts, 2009); raise average murder rates 2.4% and assault rates 3.6% relative to the coolest conditions experienced in each county and month (Ranson, 2014); increase residential energy consumption 29% and elevate mortality rates 11% on net (Deschênes and Greenstone, 2011); and reduce GDP growth by roughly 1.7 percentage points yr^{-1} (Deryugina, 2017). Temperatures contribute to 44% of civil conflicts in Sub-Saharan Africa (Burke et al., 2009), and 13% of infant mortality in the Philippines is attributable to typhoons (Anttila-Hughes and Hsiang, 2012). Globally, ENSO has elevated civil conflict rates 21% relative to constant low-conflict conditions (Hsiang, Meng, and Cane, 2011), while temperature and tropical cyclones reduce global economic growth by roughly 0.25 and 1.3 percentage points yr^{-1}, respectively (Burke, Hsiang, and Miguel, 2015b; Hsiang and Jina, 2014).

Climate change to date

Only a few agricultural studies estimate the social effect of recent already-observed anthropogenic climate trends. In Table 2.1 column 6 we show that, relative to an unchanged climate, trends in various climatic variables that occurred at the end of the

twentieth century have lowered rice yield growth rates in South Asia by up to 0.76% annually (Welch et al., 2010) and reduced global maize and wheat production 3.8% and 5.5%, respectively, whereas global gains and losses for soy and rice roughly balance one another out (Lobell, Schlenker, and Costa-Roberts, 2011). Based on calculations using data from Burke, Hsiang, and Miguel (2015b) and Burke et al. (2009), we estimate that warming trends since 1980 have slowed global average GDP growth 0.002 percentage points per year, and increased the incidence of civil conflict in Sub-Saharan Africa by approximately 16% (see Appendix A for details).

Future climate change

Projected impacts of future "business-as-usual" climate changes, relative to a counterfactual of no climate change, are generally much larger than impacts of warming that have already occurred and tend to be comparable to the baseline impact of climate on social and economic outcomes today (Table 2.1 column 7). For example, crop yields in Africa are likely to decline 17-22% for maize, sorghum, millet and groundnuts by 2050 (Schlenker and Lobell, 2010), yields for major crops in the USA are likely to decline 15-20% by 2050 (Schlenker and Roberts, 2009; Burke and Emerick, 2016; Houser et al., 2015) and 63-82% by 2100 (Schlenker and Roberts, 2009; Houser et al., 2015), although accounting for estimated effects of CO_2 fertilization may keep expected losses nearer to 15% (Houser et al., 2015). Projected estimates suggest armed conflict in Africa may rise roughly 50% by 2030 (Burke et al., 2009) while violent and property crimes in the USA may increase roughly 3% and 1% respectively (Ranson, 2014; Houser et al., 2015). Warming by end of century is projected to increase US mortality rates 3-9% and electricity consumption 11% (Deschênes and Greenstone, 2011; Houser et al., 2015). The growth rate of overall economic production is projected to fall roughly 0.12 percentage points yr^{-1} in the US (Deryugina, 2017) and 0.28 percentage pionts yr^{-1} globally (Burke, Hsiang, and Miguel, 2015b) during the next century due to the effects of rising temperature, with additional projected losses due to cyclones costing roughly $9.7 trillion dollars in present discounted value (Hsiang and Jina, 2014). Importantly, these impact projections are all constructed based on historically observed responses to environmental conditions, and the actual impact of future changes might be less disruptive if, for example, adaptive technologies improve dramatically in the future. Alternatively, future impacts could be worse than described here if current adaptive strategies, such as irrigation using fossil aquifers, are unsustainable or societal responses become highly nonlinear once the environment shifts to conditions beyond recent experience.

2.5 Critical challenge: understanding "adaptation gaps"

Overall, new empirical measurements suggest that current climatic conditions impose substantial economic and social burdens on modern populations, and that future climate change will further increase these ongoing costs considerably. These losses could be

avoided, in theory, if populations could costlessly and fully adapt to these dimensions of their climate—why this has not occurred to date remains an important open question, with potentially large gains for both present and future populations should it be solved.

Given information on the climatically determined probability distribution of potential weather events, populations may take actions or make investments that will reduce the influence of these events when they actually occur. As depicted in Figure 2.1, this adaptation can be detected implicitly by observing how the dose-response function linking climate variables to outcomes changes. More highly adapted populations will have flatter responses (Hsiang and Narita, 2012), such that changes in climatic variables have less influence on an outcome. An alternative approach to detecting adaptation is explicit measurement of outcomes that are themselves thought to be adaptations, such as investing in crop switching following a drought (Hornbeck, 2012). Notably, measurement of adaptation using either approach is made possible by the use of intertemporal changes in climatic variables, whether over short time-scales—e.g. days—or long time-scales (Burke and Emerick, 2016)—e.g. decades—and we note that in contrast to widely cited heuristics, short term weather variation can be used to exactly measure the influence of long-term climate changes under the right conditions, even when populations adapt to knowledge of their climate (Hsiang, 2016).

Comparison of adaptation results across different sectors reveals striking dissimilarities: in some cases adaptation appears remarkably effective at minimizing damages whereas in other cases we observe essentially no adaptation, leading to seemingly costly "adaptation gaps." For example, populations regularly exposed to cyclones experience substantially smaller losses than naive populations when exposed to physically similar events (Hsiang and Narita, 2012; Hsiang and Jina, 2014; Anttila-Hughes and Hsiang, 2012) (Figure 2.5 a-b). Similarly, mortality on hot days in hot climates is lower than in similar populations in cooler climates (Barreca et al., 2015; Guo et al., 2014) (Figure 2.5 c) and heat-related mortality has declined over time with rising availability of air-conditioning and other technologies (Barreca et al., 2016) (Figure 2.5 d). In sharp contrast, violence and crop yields in hot and cool locations respond almost identically to temperature in the US (Ranson, 2014; Schlenker and Roberts, 2009) (Figure 2.5 e), and the temperature sensitivity of agriculture (Burke and Emerick, 2016), crime (Ranson, 2014) and economic productivity (Deryugina, 2017; Burke, Hsiang, and Miguel, 2015b) has changed little over multiple decades, even though populations are presumably innovating and adjusting to climate over this time period (Figure 2.5 f-g). At a global scale, it has been widely hypothesized that wealthy populations will adapt effectively to future climate changes because they have greater resources, access to wider arrays of technology, and tend to have stronger governments (Waldhoff et al., 2011; Lomborg, 2004; Kahn, 2005), but data from the present largely suggests that overall economic activity in wealthy countries actually responds to temperature (in percentage terms) similarly to economic activity in poor countries (Figure 2.5 h, also g)—although there is suggestive but statistically insignificant evidence that wealthy countries might be adapting slightly more effectively. In puzzling incongruity, wealthy countries appear substantially more adapted than poor countries, in terms of some outcomes, to tropical cyclones (Hsiang and Narita, 2012) and ENSO (Hsiang, Meng, and Cane, 2011).

To date, it is not well understood why populations adapt so effectively in some dimensions with respect to certain aspects of the climate while entirely failing to adapt in other contexts. Existing evidence suggests that high costs of adaptation (Hsiang and Narita, 2012; Schlenker, Roberts, and Lobell, 2013; Davis and Gertler, 2015), incentives to adapt (Annan and Schlenker, 2015; Hsiang and Narita, 2012), limited access to credit for financing adaptations (Burgess et al., 2014), limited rationality when planning for future risks (Kunreuther and Michel-Kerjan, 2009; Deryugina, 2013), incorrect or limited information about the benefits of adaptation (Hornbeck, 2012; Shrader, 2016), perverse political incentives (Garrett and Sobel, 2003; Healy and Malhotra, 2009) or weak government institutions (Besley and Burgess, 2002; Kahn, 2005), constraints to sharing risk among individuals and groups (Townsend, 1995), and access to technologies (Barreca et al., 2015; Olmstead and Rhode, 2011a) might play substantial roles, although existing evidence is primarily suggestive as it relies on cross-sectional associations. To better understand what constrains adaptation, future work will likely need to exploit natural experiments where specific potential constraints (or costs) are exogenously eliminated; if the link between an outcome and climate disappears, it can be more confidently inferred that the altered constraint was playing a critical role in limiting adaptation (Fetzer, 2014; Barreca et al., 2016).

It is theoretically possible that existing adaptation gaps are "economically optimal" in the sense that the costs of additional adaptive actions and investments exactly balances their benefits, which are avoided climate-related social losses (Hsiang and Narita, 2012). Many patterns of adaptation described above seem qualitatively consistent with this notion of optimality, for example cyclone-prone locations benefit more from investments in cyclone shelters because they are utilized more often, so cost-benefit analyses would predict more shelters in locations that are more cyclone-prone. However, many patterns seem inconsistent with optimality, such as the persistant sensitivity of crop yields to temperature (Lobell et al., 2014; Burke and Emerick, 2016), but could be reconciled as optimal if adaptation technologies are extraordinarily costly. In general, there is no quantitative evidence that allows us to determine how closely current adaptation gaps reflect optimal investments or are bound at sub-optimal levels by the market-failures and other constraints described above.

Because the persistence of adaptation gaps has such large impacts on current and future wellbeing around the globe, understanding its cause is likely the most pressing current research question. Identifying the causes of these gaps and determining whether they are optimal is critical for designing policies that can support and accelerate adaptation in the numerous contexts where it lags. For example, if current adaptation gaps are optimal then policy should focus on improving the cost-effectiveness of adaptation technologies (Hsiang and Narita, 2012) rather than on correcting market failures. Such policies, if carefully designed and effectively implemented, could both substantially benefit current generations that presently suffer large economic and social burdens from the modern climate, as well as benefit future generations that would otherwise continue bearing these burdens along with all additional costs of climate change.

2.6 Discussion

The endeavor to understand the impact of climate on society is unlocking promise. Climate has imposed varied environmental constraints on humanity for millennia and new understanding provides insight into the role of climate in global historical development. More urgently, current climatic conditions and variations are constantly shaping and reshaping human wellbeing today, thus understanding these processes allow us to better prepare for and respond to the climate that we experience now. Finally, designing effective, efficient, and fair policies to manage anthropogenic climate change requires, critically, that we develop a quantitative grasp of how different investments today may impact economic and social possibilities in the future.

Advances in data, computing, and methods have triggered rapid progress in our ability to empirically measure how climatic conditions affect human wellbeing and productivity around the world. While climate is clearly not the only factor that affects social and economic outcomes, quantitative measurements reveal that it is a major factor, often with first-order consequences. Notably, these results suggest that the magnitude of influence that current climatic conditions have on social outcomes is generally comparable (and sometimes larger) than the projected effects of future warming. Collectively, these findings suggest that both local climatic conditions and the state of the global climate can be thought of as forms of "natural capital" that play an important role in supporting human welfare and are inputs to economic production.

An insight that emerges from these findings is the notion that current climate patterns may be an important source of inequality. Populations endowed with different distributions of climatic conditions face different environmental constraints that may lead to different distributions of outcomes. In a thought experiment where we hold all other factors constant, these recent findings directly suggest that hotter locations with more extreme rainfall patterns and more major disturbances, such as ENSO and tropical cyclones, will generally face additional health costs, lower productivity and additional economic costs, greater population movement, and higher rates of violence. To first order, this idea is broadly consistent with cross-sectional patterns (Nordhaus, 2006), however, as described earlier, it is not yet possible to ensure the "all other factors constant" assumption holds when comparing outcomes across different populations, so we cannot directly test these cross-sectional predictions empirically. Nonetheless, such inferences, with important repercussions for present and future inequality, would follow logically from these results.

Projections of climate changes based on these empirical results also inform questions of inequality, as predicted future impacts are highly unevenly distributed across locations, often because the effects of climate are nonlinear and different populations have different baseline climates, such that incremental warming has heterogeneous effects. For example, warming is expected to increase productivity in cool locations while decreasing productivity in warm locations, leading to projections where current patterns of inequality increase, sometimes dramatically (Houser et al., 2015; Deryugina and Hsiang, 2017; Burke, Hsiang, and Miguel, 2015b).

Recent advances in this literature point towards two areas of future work with

important policy consequences. First, "cracking the code" on when, where and why adaptation is or is not successful promises major social benefits. New evidence suggests that (i) there are some cases where populations are able to adapt such that they partially neutralize the effects of climate; (ii) there are many cases where adaptation does not occur; (iii) the social and economic benefits of successful, low cost, and widespread adaptation are potentially very large for both current and future populations, especially in many low income countries. Understanding what causes this "adaptation gap" can help policy address it; for example, if adaptation technologies are expensive (Hsiang and Narita, 2012; Schlenker, Roberts, and Lobell, 2013), then policy should focus on their research and development. Secondly, models used to understand the costs and benefits of different global climate change policies take as inputs various "damage functions" that describe how social and economic losses accrue under different future climate change scenarios (Interagency Working Group on Socal Cost of Carbon, 2010). Historically, these damage functions were theoretical constructs whose structure was based on modeling intuition informed by some data (Revesz et al., 2014; Pindyck, 2013), but the recent explosion of empirical work suggests that these global policy models can now be calibrated to real world relationships that characterize the many social impacts of climate (Houser et al., 2015; Moore and Diaz, 2015).

Chapter 3

Climate and suicide in India

3.1 Introduction*

Each year, over 130,000 lives are lost to self-harm in India (Patel et al., 2012). The causes of these deaths are poorly understood; drivers of suicidal behavior remain disputed across scientific disciplines, and nearly all evidence comes from developed country contexts (Mann et al., 2005; Stack, 2000; Rehkopf and Buka, 2006). Despite lack of substantiation, public debate in India has centered around one possible cause of rapidly rising suicide rates: increasing variability of agricultural income (Patel, 2007; Sengupta, 2006). Drought and heat feature prominently in these claims; climate events are argued to damage crop yields, deepening farmers' debt burdens and inducing some to commit suicide in response. With more than half of India's working population employed in agriculture, one-third lying below the international poverty line, and nearly all experiencing rising temperatures due to anthropogenic climate change, these arguments appear plausible. However, the relationship between economic shocks and suicide is controversial (Durkheim, 1951; Hamermesh and Soss, 1974; Stack, 2000; Rehkopf and Buka, 2006; Becker and Posner, 2004), and in India the effect of income-damaging climate variation on suicide rates is unknown. While the national government has recently announced a $1.3 billion climate-based crop insurance scheme motivated as suicide prevention policy (Agence-France Presse, 2016), evidence to support such an intervention is lacking.

Existing work has found that agricultural yields in India rely heavily on growing season temperature and precipitation (Auffhammer, Ramanathan, and Vincent, 2006, 2012), but it is unclear to what extent, if any, this sensitivity to climate influences suicide rates. Previous studies of income variability affecting suicide in India are anecdotal (Patel, 2007) or qualitative (Mohanty, 2005; Herring, 2008; Deshpande, 2002; Rao and Gopalappa, 2004; Sarma, 2004) and none attempt to identify and synthesize quantitative relationships between climate, crops, and suicides. To fill this knowledge gap, I

*The material from this chapter was published in the *Proceedings of the National Academy of Sciences* in August, 2017. The published version can be found here: `https://doi.org/10.1073/pnas.1701354114`.

use a dataset from India's National Crime Records Bureau (NCRB) which contains the universe of reported suicides in the country from 1967 to 2013. I pair these data with information on agricultural crop yields and high resolution climate data to identify the effect of climatic shifts on suicide rates, and to test whether agricultural yields are a mechanism through which these effects materialize. While my analysis is most directly applicable to India, it is also the first large-scale estimate of the effect of climate on suicide in the developing world.

My empirical strategy relies on a simple thought experiment in which I observe two identical populations, alter the climate in one, and compare suicide rates in this "treatment" population to those in an unaltered "control." In the absence of such an experiment, I emulate this comparison by observing a population within India under different climate realizations over time, allowing the same population to function as both "treatment" and "control." After accounting for secular trends, year-to-year changes in the climate are plausibly random, and amount to many ongoing approximations of my ideal experiment (Carleton and Hsiang, 2016). Because this approach isolates random variation in climate, other common factors associated with both suicide and the climate are unlikely to confound the analysis. Therefore, a causal interpretation of estimated regression coefficients is reasonable, even though the climate itself was not experimentally manipulated.

I analyze the relationship between annual suicide rates, measured at each of India's 32 states and union territories, and cumulative exposure to temperature and rainfall using a regression model that accounts for time-invariant differences across states in unobservable determinants of suicide rates, such as religion or history, as well as regional time trends in suicide rates which may derive from shifting cultural norms or suicide contagion effects, among many other possible forces. Under my estimation strategy, two key empirical concerns remain. First, the functional form of the relationship between suicide rates and climate variables has minimal precedent in existing literature. I therefore use a flexible nonlinear model and show robustness of my results to alternative functional form assumptions. Second, the channels through which adverse climate conditions may affect suicide rates are not immediately discernible, yet are of central policy relevance. To this end, I distinguish between climate conditions that damage crops and those that have no effect on agricultural yields. I do so by estimating differential impacts of climate during growing and non-growing seasons, using the arrival and departure of the southwest summer monsoon to define seasonality (see Appendix B for details). In additional mechanisms tests, I use spatial heterogeneity and temporal lags to assess the mediating factors between climate and suicide.

3.2 Temperature's influence on suicide rates

I find that temperature during India's main agricultural growing season has a strong positive effect on annual suicide rates (Figure 3.1a-3.1b, Table 3.1, and Table B.3). For days above 20°C, a 1°C increase in a single day's temperature during the growing season increases annual suicides by 0.008 per 100,000 people, causing an additional 67

deaths on average. This is a 3.5% increase in the suicide rate per standard deviation (σ) increase in temperature exposure. In contrast, temperatures in the non-growing season have no identifiable impact on suicide rates. This finding is robust to: inclusion of state-specific trends and national-level shocks to the suicide rate (Tables B.3, B.8); distinct methods for averaging gridded climate data across pixels within a state (Table B.5); alternative degree day cutoff values (Table B.6); controlling for irrigation (Table B.10); and alternative definitions of the growing season (Table B.11).

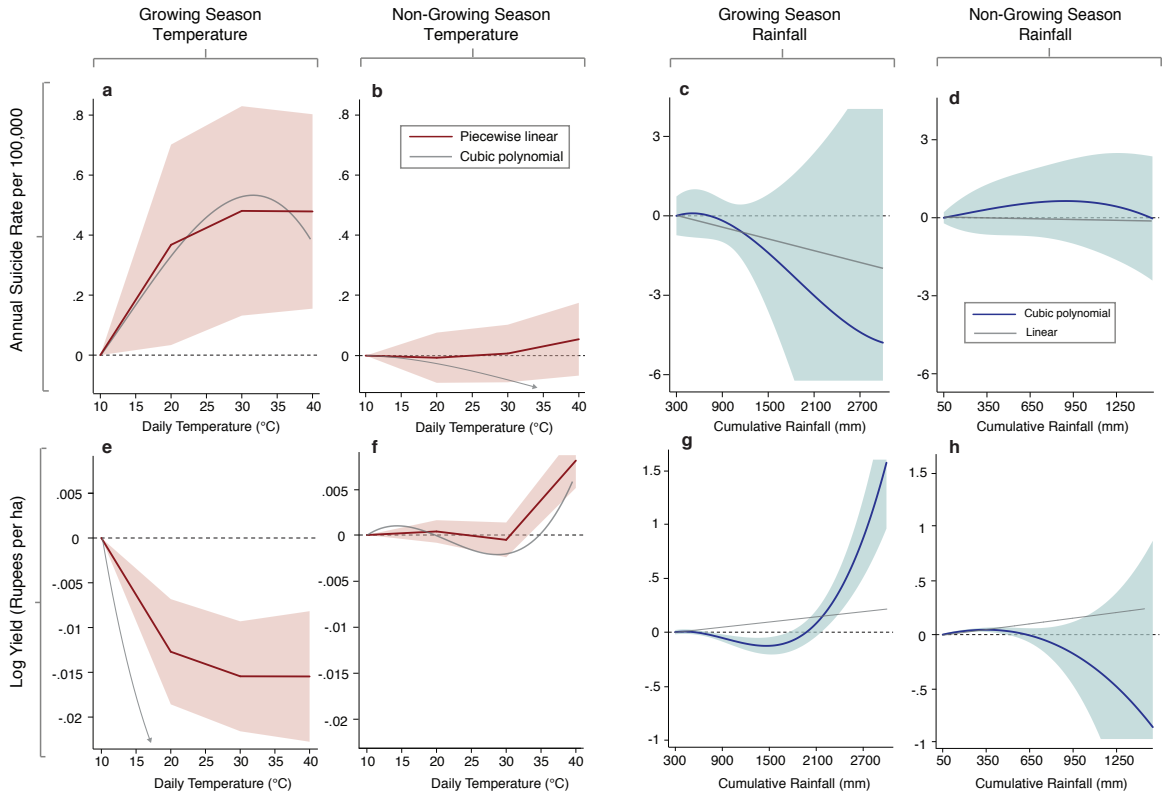

Figure 3.1: Nonlinear relationships between temperature, precipitation, suicide rates, and crop yield

The response of annual suicides rates (deaths per 100,000 people) to (**a**) growing season, and (**b**) non-growing season temperatures. Response of annual suicide rates to cumulative (**c**) growing season and (**d**) non-growing season rainfall. Analogous relationships for log annual yield, valued in Rupees per hectare (**e** − **h**). The slopes of the temperature figures panels (**a**), (**b**), (**e**), and (**f**) can be interpreted as the change in the annual suicide rate or log yield caused by one day's temperature rising by 1°C. The slopes of the precipitation figures panels (**c**), (**d**), (**g**), and (**h**) can be interpreted as the change in the annual suicide rate or log yield caused by one additional millimeter of rainfall. All graphs are centered at zero.

The differential response of suicide to temperature in the growing and non-growing seasons is consistent with an agricultural channel in which heat damages crops, placing economic pressure on farming households who may respond to such hardship with

suicide. These crop losses may also permeate throughout the economy, causing both farming and non-farming populations to face distress as food prices rise and agricultural labor demand falls. To further test this mechanism, I use district-level yield data covering 13 Indian states from 1956 to 2000 to estimate an identical regression model to that described above, now measuring the response of crop yields to variations in the climate. I find that yields mirror suicides in their response to temperature, falling with rising growing season temperatures but reacting minimally to non-growing season heat (Figure 3.1e-3.1f), a result identified in many other parts of the world (Auffhammer, Ramanathan, and Vincent, 2012; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010). For growing season days above 20°C, annual yields fall by 1.3%/σ. This finding is robust to the same specification checks listed above for suicide (Tables B.4, B.5, B.7, B.11). The striking similarity between the responses of suicide and yield to temperature suggests that variations in temperature affect suicide rates through their influence over agricultural output.

	Suicides per 100,000			*100×Log yield (Rupees/ha)*		
	State trends	Year fixed effects	State trends + year fixed effects	State trends	Year fixed effects	State trends + year fixed effects
Growing season						
Degree days below threshold	0.003***	0.000	0.004***	0.013	-0.019	-0.003
	(0.001)	(0.001)	(0.001)	(0.009)	(0.018)	(0.013)
Degree days above threshold	0.007***	0.009**	0.008**	-0.017***	-0.020*	-0.019*
	(0.002)	(0.004)	(0.003)	(0.006)	(0.010)	(0.010)
Nongrowing season						
Degree days below threshold	-0.001	-0.009*	-0.003*	0.002	0.007	0.001
	(0.001)	(0.004)	(0.002)	(0.003)	(0.005)	(0.004)
Degree days above threshold	-0.002*	0.002	0.001	0.010***	0.018***	0.010*
	(0.001)	(0.003)	(0.003)	(0.004)	(0.006)	(0.006)

Table 3.1: Effect of heat exposure on suicide rates and yield values, by agricultural season

Coefficients represent the effect of one day becoming 1°C warmer on the annual suicide rate (suicide deaths per 100,000 people) or annual yield (log Rupees/ha), where the degree day threshold is 20°C. All regressions include a cubic polynomial of seasonal precipitation (coefficients not shown). Columns (1)–(3) include state fixed effects and report standard errors clustered at the state level. Columns (4)–(6) include district fixed effects and report standard errors clustered at the district level.
*** p<0.01, ** p<0.05, * p<0.1.

India's agriculture is predominately rain-fed and dependent on the timing and duration of the monsoon, making growing season rainfall critical for crop growth (Burgess et al., 2014), as well as a potential driver of suicide. As expected, growing season precipitation positively impacts yields with an effect of 1.9%/σ, while non-growing season rainfall (of which there is little) has no statistically distinguishable effect (Figure 3.1c-3.1d and 3.1g-3.1h). These yield gains again mirror the response of suicides to climate — suicide rates fall as growing season rainfall increases — although the relationship is statistically insignificant across most robustness checks (Tables B.3-B.11). Despite

statistical uncertainty, the yield and suicide response functions with respect to rainfall also match in the non-growing season, where a flat relationship is estimated in both cases.

Imprecision in these rainfall estimates for suicide may be due to measurement error introduced by my need to characterize monsoon rainfall at the state level, as there can be important within-state differences in monsoon arrival and withdrawal (Burgess et al., 2014). The district-level agricultural data, in contrast, do not suffer from this problem. Consistent with measurement error, a less parametric estimate of rainfall's effect on suicide separately during each month of the year demonstrates that rain during all growing season months negatively influences suicide rates, but with high uncertainty (Figure B.7). Moreover, results from an alternative empirical model measuring impacts of longer-run trends in climate demonstrate a robust and substantial negative effect of growing season rainfall on suicide rates (Table B.9). Under this approach, I find that increasing growing season rainfall by 1cm is associated with a decrease of approximately 0.8 deaths per 100,000, lowering the suicide rate by 7% on average. Together, these results suggest that rainfall may mitigate suicide rates in India, plausibly through an agricultural channel.

3.3 The agricultural mechanism

I further examine the agricultural mechanism by including lagged effects in the regression model. If suicides are affected by climate variation through negative agricultural income shocks, there may be delayed impacts: poor harvests in one year may make subsequent conditions more unbearable, as households draw on stored crops or deplete monetary savings. In contrast, if these climate variables influence suicide prevalence purely through direct channels, such as the hypothesized neurological effects of heat exposure on aggressive behavior (Seo, Patrick, and Kennealy, 2008; Lövheim, 2012), delayed effects should not materialize. A model that includes lagged climate variables reveals that past growing season temperatures strongly influence suicide rates, with effects that last for approximately 5 years (Figure 3.2). Similarly, high precipitation years have a strong lagged effect in which heavy rainfall today causes lower suicide rates in 2 to 3 years; this beneficial yield shock may enable individuals to save crops and income, making future suicides less likely (Figure 3.2). Interestingly, drought appears to have no effect on suicide rates, either contemporaneously or in lagged form (Figure B.8).

Geographic heterogeneity in both suicide and crop yield impacts can be used as an additional means of assessing the channel through which climate drives suicides. I disaggregate suicide response functions by state to detect a clear geographic pattern in which southern states — which are generally hotter, have higher average suicide rates, and display steeper suicide trends over time — have much stronger responses to growing season temperature (Figure 3.2c). I obtain similarly heterogeneous responses of agricultural yields to growing season temperatures for each of the 13 states included in the crop data. While these estimates have large uncertainty, the correlation between each state's yield sensitivity to growing season temperature and its suicide sensitiv-

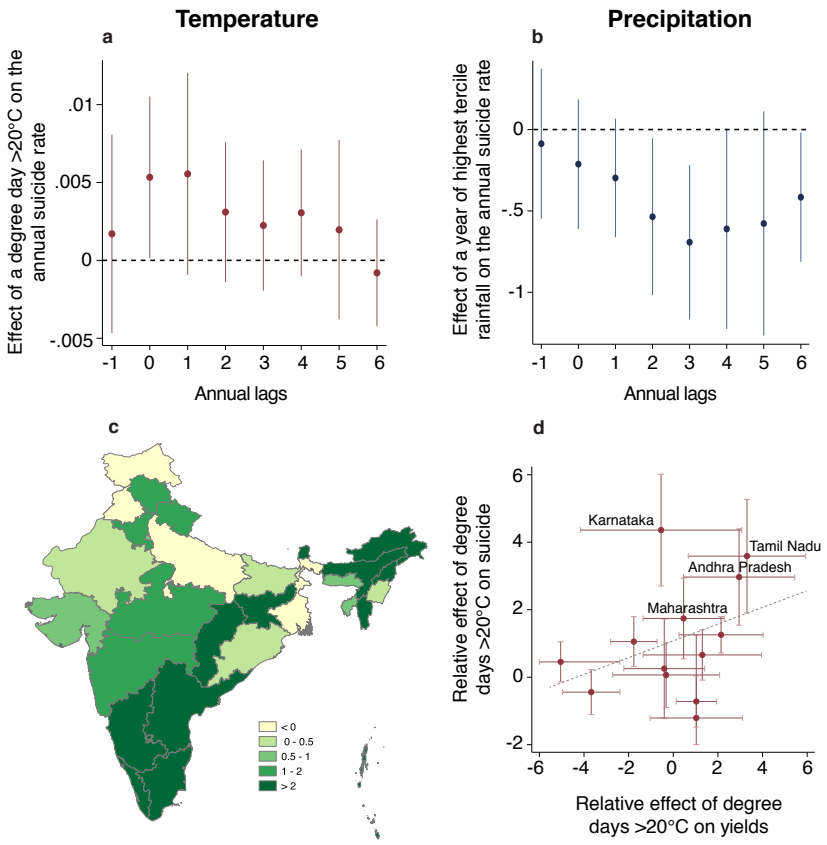

Figure 3.2: Evidence for the agricultural income channel

Lagged effects of growing season (**a**) temperature degree days and (**b**) high precipitation (years in which precipitation falls into the highest tercile of the long-run rainfall distribution) on annual suicide rates per 100,000 people suggest an economic mechanism for climate impacts. (**c**) Geographic heterogeneity in the suicide-temperature response, where states are colored by the state-specific temperature sensitivity as a fraction of the average treatment effect. Darker colors indicate more severe responses of suicide to growing season temperature, relative to the average; yellow indicates a negative effect. (**d**) Correlation between state-level suicide sensitivities and the additive inverse of corresponding state-level crop yield sensitivities. Temperature effects are shown as relative to the average treatment effect. Coefficients in all models were estimated in a degree days model with a cutoff of 20°C. Standard errors are clustered at the state level for suicide and district level for yield, and 95% confidence intervals are shown around each coefficient.

ity is positive, suggesting that states where agricultural yields are more damaged by high temperatures are also the states where these temperatures increase suicide rates substantially (Figure 3.2d). Three states that have been at the center of India's public debates regarding agricultural influences on suicide (Maharashtra, Karnataka and Andhra Pradesh) not only have severe suicide responses to temperature, but also exhibit large negative impacts of temperature on yield.

3.4 Testing for evidence of adaptation

As anthropogenic climate change raises temperatures throughout the world, a central question for global welfare is the extent to which populations adopt adaptive behaviors to prevent climate damages (Carleton and Hsiang, 2016). I conduct four sets of tests to assess the evidence for four distinct hypotheses regarding adaptive behavior in the context of suicide in India: 1) locations that are hotter on average exhibit lower sensitivity to temperature, as populations acclimatize; 2) locations that are wealthier on average exhibit lower sensitivity to temperature, as wealth enables investment in adaptation; 3) temperature sensitivity has declined over time as incomes and access to modern agricultural technologies have risen; and 4) sensitivity to longer-run gradual trends in temperature will be lower than sensitivity to short-run variations in temperature, as populations require time to adapt. My estimation strategies for testing these hypotheses are detailed in Appendix B. Across all four tests, I find no evidence of any type of adaptive behavior. In hotter locations I detect higher than average sensitivity to temperature, contradicting my first hypothesis (Figure 3.3a). Temperature sensitivity is indistinguishable in wealthier locations from that in poor locations, failing to support my second hypothesis (Figure 3.3b). Temperature sensitivity of suicide has remained remarkably stable over time, despite India's robust economic growth and dramatic improvements in agricultural yields over this period (Figure 3.3c). And finally, the impact of gradual changes ("long differences") is, in contrast to my final hypothesis, more severe than that of short-run variations in temperature (Figure 3.3d). Taken together, these tests reveal no evidence of adaptive behavior in the context of temperature damages to suicide rates in India.

3.5 Discussion

As India's suicide rate continues to rise, the causes of these deaths remain heavily debated. In this study, I find that variations in temperature during India's main growing season exert substantial influence over suicide rates. To explore the significance of this effect to total trends in India, I extend my results to calculate the share of this upward trend that is attributable to changes in India's climate over recent decades. In particular, I measure the additional number of deaths attributable to warming temperatures throughout India since 1980 (see Appendix B for details on this approach). I find that by 2013, temperature trends are responsible for over 4,000 additional deaths annually across India, accounting for approximately 3% of annual suicides (Figure 3.4). Across all states and all years since 1980, a cumulative total of 59,300 suicides can be attributed to warming, accounting for 6.8% of the national upward trend in suicides over this time period. As shown in Figure 3.4, there is substantial geographic heterogeneity behind this aggregate value, due both to differential warming trends and to uneven population distributions across the country.

My study has important limitations. Of primary concern is that I do not have a quasi-experiment in which agricultural incomes were randomly allocated across popu-

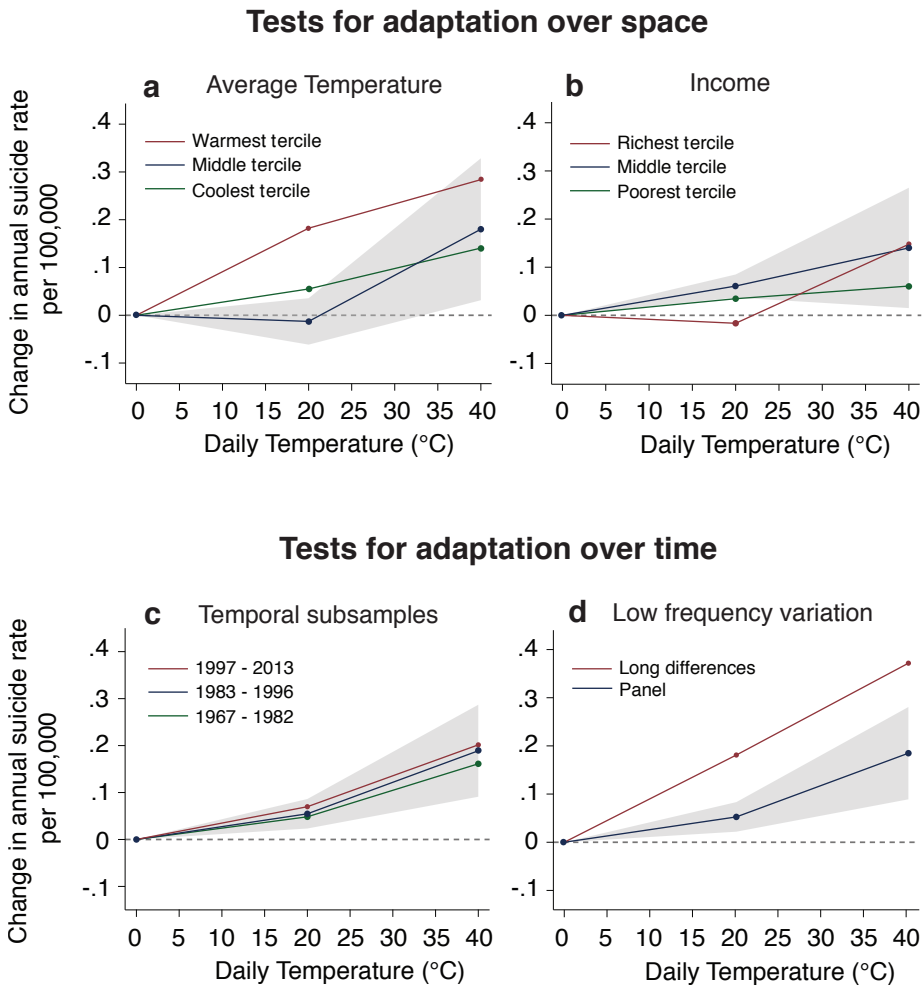

Figure 3.3: Tests of adaptation in the suicide-temperature relationship

Panels show heterogeneity in the suicide response to growing season degree days above 20°C, by (**a**) terciles of long-run average growing season degree days, (**b**) GDP per capita in 2010, (**c**) periods within the sample, and (**d**) across two different estimation strategies ("long differences" estimates the effect of long-run climate trends, while "panel" estimates the effect of year-to-year variation). Shaded areas indicate the 95% confidence interval around (**a**)–(**b**) the middle tercile response function, (**c**) the period 1983 to 1996, and (**d**) the panel method.

lations within India, and suicide rates were monitored in response. Thus, while I employ multiple distinct approaches aimed at pinning down the agricultural mechanism through which climate affects suicide, I do not have a direct test of the common hypothesis that climate-induced economic hardship can lead some individuals to respond with self-harm. Secondly, my empirical strategy relies on estimating the effects of year-to-year variation in temperature and precipitation on suicides within a given state; while this facilitates a causal interpretation of estimated coefficients, it does not guarantee that there are no other factors correlated with both suicide and climate within a state

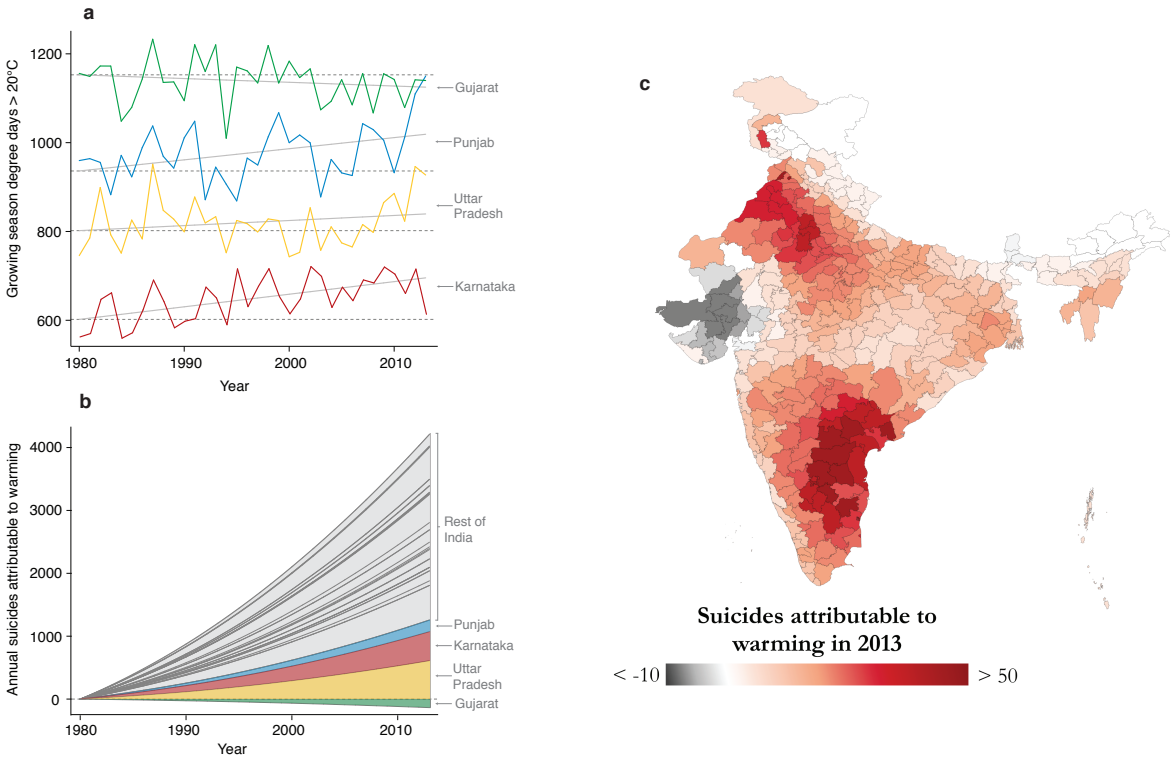

Figure 3.4: Attribution of suicides to warming trends in growing season temperatures since 1980

Panel (**a**) shows trends in degree days above 20°C during India's main growing season for four example states. Panel (**b**) shows the total number of deaths annually that can be attributed to warming trends, using the estimated marginal effect of degree days on suicide rates, the observed warming since 1980, and population estimates from the Indian census. Panel (**c**) shows the geographic distribution of suicides attributed to warming in 2013 at the district level.

that could confound my estimation. However, the robustness of the effect of growing season temperature on suicide rates across many specifications (Tables B.3-B.13) and subsamples (Figure 3.3) makes such confounding factors extremely unlikely.

Despite these necessary shortcomings, my findings convey important lessons for current and future generations. Suicide rates are a salient indicator of serious human hardship. My identification of a substantial effect of climate variation on this measure of human suffering in one fifth of the global population provides the first empirical support for policies which aim to prevent suicides through tools that alleviate the impacts of climate on income, such as crop insurance. These findings are also critical inputs into policy decisions regarding future climate change mitigation and adaptation. As I find no evidence that adaptation has occurred over 46 years in a large and rapidly-developing country, and because suicide prevalence is a valuable measure of well-being, the magnitude of effects I detect has important consequences for assessing the likely impact of future climate change on human welfare globally. India alone is predicted to

experience an average temperature increase of up to 3°C by 2050 (Hijioka, Y. et al., 2014). Without investments in adaptation, my findings suggest that this warming will be accompanied by a rising number of lives lost to self-harm.

3.6 Data and methods summary

Suicide and agricultural data

Annual suicide data are reported by the Indian National Crime Records Bureau (NCRB) at the state level beginning in 1967 for 27 of India's 29 states and 5 of its 7 Union Territories. Suicide records are in the NCRB's "Accidental Deaths and Suicides in India" report and include total number of state suicides per year. I calculate suicide rates as the number of total suicides per 100,000 people, with population data linearly interpolated between Indian censuses. I use agricultural data from Duflo and Pande (2007). These are district-level annual yield records for major crops (rice, wheat, sugar, sorghum, millet and maize) between 1956 and 2000, compiled from Indian Ministry of Agriculture reports and other official sources. These data cover 271 districts in 13 major agricultural states: Andhra Pradesh, Bihar, Gujarat, Haryana, Karnataka, Madhya Pradesh, Maharashtra, Orissa, Punjab, Rajasthan, Tamil Nadu, Uttar Pradesh, and West Bengal. These data provide log annual yield values of a production-weighted index across all crops measured in constant Indian Rupees, where prices are fixed at their 1960-1965 averages. Details on these data sets and summary statistics are in Appendix B.

Climate data

Climate data are generally available at higher spatial and temporal resolution than social outcome data. Although suicides and yields are only measured annually, if the relationship between these outcomes and temperature is nonlinear, daily climate data are required, as annual average temperatures obscure such nonlinearities (Hsiang, 2016). For daily temperature data, I use the National Center for Environmental Protection (NCEP) gridded daily reanalysis product, which provides observations in a $1^\circ \times 1^\circ$ grid (Kalnay et al., 1996). These data include daily mean temperature for each grid over my entire sample period. To convert daily temperature into annual observations without losing intra-annual variability in daily weather, I use the agronomic concept of degree days. Degree days are calculated as follows, where t^* is a selected cutoff temperature value and t is a realized daily temperature value:

$$D^{t^*}(t) = \begin{cases} 0 & if \;\; t \leq t^* \\ t - t^* & if \;\; t > t^* \end{cases} \tag{3.1}$$

Degree days allow temperature to affect an outcome variable only once its value surpasses the threshold t^*. Because there are multiple grid cells per state, I aggregate grid-level degree day values to state-level observations using an area-weighted average

(see Appedix B for robustness checks using weights based on population and area planted to crops). When these state-level degree day values are summed over a time, regressing an outcome on cumulative degree days $\sum_{t=1}^{T} D^{t^*}(t)$ imposes a piecewise linear relationship, in which the outcome response has zero slope for all temperatures less than t^*. While a body of literature identifies biologically-determined cutoffs t^* for yields of a variety of major crops, there is no empirical support to draw on in selecting t^* for suicides. Thus, while I use $t^* = 20°$C throughout the study, I show robustness for a range of plausible cutoffs based on the distribution of my temperature data, and in Figure 3.1 I estimate a flexible piecewise linear function using four different degree day cutoffs to impose minimal structure on the response function.

Because reanalysis models are less reliable for precipitation data, and because non-linearities in precipitation that can't be captured with a polynomial appear to be less consistently important both in the violent crime literature (Burke, Hsiang, and Miguel, 2015a) and in the agriculture literature (Schlenker and Roberts, 2009), I use the University of Delaware monthly cumulative precipitation data to complement daily temperature observations (Willmott and Matsuura, 2014). These data are gridded at a $0.5° \times 0.5°$ resolution, with observations of total monthly rainfall spatially interpolated between weather stations. I again aggregate grids up to states using area-based weights, after calculating polynomial values at the grid-level first. Summary statistics for both rainfall and temperature are shown in Appedix B.

Regression estimation

To identify the impact of temperature and precipitation on suicide rates, I estimate a multivariate panel regression using ordinary least squares, in which the identifying assumption is the exogeneity of within-state, annual variation in degree days and cumulative precipitation. My primary estimation approach employs a flexible piecewise linear specification with respect to temperature and a cubic polynomial function of cumulative precipitation. To capture the distinct impact of economically meaningful climate variation, I separately identify the temperature and precipitation response functions by agricultural seasons (see Appendix B for details). My empirical model takes the general form:

$$\begin{aligned} suicide\_rate_{it} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\sum_{d\in s} DD_{idt}^{k} \\ &+ \sum_{s=1}^{2} g\left(\sum_{m\in s} P_{imt}\right) + \delta_i + \eta_t + \tau_i t + \varepsilon_{it} \end{aligned} \tag{3.2}$$

Where $suicide\_rate_{it}$ is the number of suicides per 100,000 people in state i in year t, $s \in \{1, 2\}$ indicates the season (growing and non-growing), and $k = 1, ..., \kappa$ indicates a set of degree day cutoffs that constrain the piecewise linear response. In my most

flexible model I let $\kappa = 7$ with degree day intervals of 5°C, and in my simplest model I let $\kappa = 2$ and estimate a standard degree day model with just one cutoff. DD^k_{idt} is the degree days in bin k (e.g. degree days between 10°C and 20°C) on day d in state i, and P_{imt} is cumulative precipitation during month m in year t. I estimate $g(\cdot)$ as a cubic polynomial. State fixed effects δ_i account for time-invariant unobservables at the state level, while year fixed effects η_t account for India-wide time-varying unobservables. In most specifications, I include state-specific time trends $\tau_i t$ to control for differential trends in suicide driven by time-varying unobservables. Robustness to different fixed effects specifications is shown in Appendix B.

Equation 3.2 identifies $\hat{\beta}_k$, the season-specific estimated change in the annual suicide rate caused by one day in bin k becoming 1°C warmer. This annual response to a daily forcing variable is similar to that estimated and described in Deryugina and Hsiang (2017). The polynomial response function for precipitation generates marginal effects of one additional millimeter of rainfall, again estimated seasonally. Due to likely correlation between errors within states, I cluster standard errors at the state level. This strategy assumes spatial correlation across states in any time period is zero, but flexibly accounts for within-state, across-time correlation. I estimate a nearly identical specification as shown in Equation 3.2 for agricultural yields. However, in this case I have district level data and therefore include district fixed effects, state-specific time trends, and I cluster standard errors at the district level.

Tests for evidence of adaptation

Exploiting heterogeneity across both time and space, I use four independent empirical tests designed to uncover evidence of adaptation. These regression models each are a variant of Equation 3.2 in which $\kappa = 2$, the degree day cutoff is set to 20°C, and state-specific linear trends are included. Figure 3.3 shows the results from these four tests, the details of which can be found in Appendix B. In panels (a)–(b), I estimate Equation 3.2, but add an interaction term between degree days in the growing season, DD_{idt}, and (a) an indicator for the tercile of average growing season degree days that state i falls into; (b) an indicator for the tercile of average GDP per capita that state i falls into. These distributions are taken over all states and all years in the sample. In panel (c), I split the 46 years in my sample into three temporal subsamples, and estimate the coefficient on an interaction between growing season degree days and an indicator for each of these three temporal subsamples. In panel (d), I estimate a "panel of long differences" empirical model in addition to the standard panel regression in Equation 3.2 (Burke and Emerick, 2016). To do so, I first construct a moving average of the suicide rate and climate variables with a window of 5 years, over the entire sample. I then calculate the 10-year change in this average at four points in my sample: 1970, 1980, 1990 and 2000. That is, my data are collapsed to 4 observations for each state, where each observation measures the 10-year change in suicide rates and climate variables for each decade, and where these changes are "smoothed" by taking 5-year averages at the end points. I then estimate the effect of changes in average degree days

and precipitation on changes in average suicide rates. This model takes the following form:

$$\Delta suicide\_rate_{i\tau} = \sum_{s=1}^{2} \beta_s \Delta DD_{is\tau}^{k} + \sum_{s=1}^{2} \gamma_s \Delta P_{is\tau} + \delta_i + \nu_\tau + \varepsilon_{i\tau} \tag{3.3}$$

Where δ_i are state fixed effects, ν_τ are fixed effects for each of the four decadal starting points in my sample, s indicates the growing and non-growing seasons, and Δ indicates the 10-year change in each variable. In Table B.9, I report results both including and excluding the decadal starting point fixed effect ν_τ.

Attribution of climate trends

To compute estimates of the effect of warming temperature trends since 1980, I follow the approach outlined in Lobell, Schlenker, and Costa-Roberts (2011) and in Carleton and Hsiang (2016). I first estimate a state-specific linear trend in growing season degree days above 20°C for the years 1980-2013. I then generate a de-trended degree days residual that is normalized to temperature in 1980 and predict suicide rates using actual and de-trended growing season degree days. In so doing, I use the coefficient estimates from Column (1) in Table 3.1. The elevated risk of suicide attributable to the trend, relative to the de-trended counterfactual, is the difference between these two predictions. Multiplying by the population in each state and each year recovers the total additional number of suicides. Figure 3.4 panel (b) displays these additional deaths in each year; integrating over states and years gives the cumulative effect of temperature trends for all of India over the entire period since 1980 (see Appendix B for details).

Chapter 4

Global mortality consequences of climate change

4.1 Introduction*

Understanding the likely economic impacts of climate change is a topic that is of considerable interest to policymakers and researchers. It is perhaps not surprising then that a variety of research approaches have been employed in the last several decades, each with their own strengths and weaknesses (e.g. Nordhaus, 1994; Tol, 1997; Stern, 2006; Kopp and Mignone, 2012; Greenstone, Kopits, and Wolverton, 2013; Burke, Hsiang, and Miguel, 2015b; Hsiang et al., 2017). An impediment to making progress across these different approaches has been a failure to articulate the criteria for the "ideal" study. The result has been a growing collection of studies, where scholars rely on their preferences or research style to guide their choice of which study is most informative.

We propose that global estimates of the economic impacts of climate change should meet three criteria. First, they should rely on data, not just researcher intuition about the nature of the relationship. These data ought to be globally representative, not only from wealthy countries. Second, the underlying relationships between human well-being and climate should be based on estimated relationships that are plausibly causal and unlikely to confound climate with unobserved factors. Third, estimated damages should reflect both the costs and benefits associated with adaptations that populations will undertake in response to climate change. The literature, certainly including much of our own prior work, has been plagued by an inability to meet all three of these criteria in a single paper. The result is that the voluminous climate impacts literature can be categorized into papers that have been successful when judged by one or two of these criteria, but not by all three.[2]

*The material from this chapter is from work in progress co-authored with many members of the Climate Impact Lab — Michael Delgado, Michael Greenstone, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert Kopp, Ishan Nath, James Rising, Ashwin Rode, Justin Simcock, and Jiacan Yuan — as well as with Samuel Seo, Arvid Viaene, and Alice Zhang.

[2]A prime example of this is given by two of the major papers in this literature, both dealing with agriculture: Mendelsohn, Nordhaus, and Shaw (1994) and Deschênes and Greenstone (2007). The

This paper develops estimates of the impact of climate change on human mortality while making progress toward meeting each of these criteria simultaneously. Specifically, we develop a plausibly causal multi-country response surface characterizing the effect of temperature on mortality, accounting for costly adaptations that societies undertake to protect themselves. Importantly, we use all available data, describing half the world's population, to construct these estimates, and then use observed within-sample heterogeneity to predict temperature responses in regions lacking data. We then use these results to construct a climate change damage function covering the world's entire population, capturing current and projected future heterogeneity across nearly 25,000 regions of the world. We focus on mortality because previous literature indicates that mortality damages from climate may be among the largest impacts (Diaz, 2014; Burgess et al., 2014; Hsiang et al., 2017), although, importantly, this method can be applied to other economic outcomes in future work. Multiple estimates of this form, spanning all relevant sectors of the economy, can then be integrated to develop a comprehensive damage function that reflects overall economic gains and losses across populations, thereby supporting the eventual calculation of an empirically-grounded estimate of the full social cost of carbon (Interagency Working Group on Socal Cost of Carbon, 2010; Metcalf and Stock, 2017).

Three features of our analysis merit highlighting. First, we estimate the current mortality-temperature relationship using the most exhaustive dataset ever collected on annual, sub-national mortality statistics. These data cover the universe of deaths from 41 countries totaling 56% of the global population; due to their detail, the unit of observation in our regression is the 2$^{\text{nd}}$-administrative level (e.g., in the U.S. this is a county) by year by age category (i.e., <5, 5-64, and >64). This allows us to estimate the mortality-temperature relationship with substantially greater coverage and accuracy for a larger proportion of the human population than previous studies; the most comprehensive previous work has been for a single country or individual cities from several countries.[3] In our global analysis, we uncover a clear U-shape in the mortality-temperature relationship, consistent with previous work. When averaged across all age groups, we find that in our sample an additional 35°C day (-10°C day), relative to a day at 20°C, increases the annual mortality rate by 0.45 (0.32) deaths per 100,000—an effect that is largest for populations younger than 5 years old or older than 64 years old. Throughout our analysis, we maintain distinct models for all age categories, allowing us to better account for different demographic structures in populations outside of our sample and evolving demographics into the future.

former credibly shows adaptation at the cost of a causal interpretation, the latter sacrifices information on adaptation in order to produce a well-identified causal result.

[3]This problem of poor coverage has been especially troublesome for the integrated assessment models (IAMs) that aim to represent global relationships. For example, the mortality damage relationship underlying FUND (Tol, 1997), a major IAM, derives its global empirical estimates from Tol (2002), which in turn draws on studies conducted in the 1990s that are not global and often use limited data or no data at all in their estimations. Among these studies, only Martens (1998) relies on mortality data and this is only via a meta-analysis of results from a select handful of locations (i.e., New York, Los Angeles, Tokyo, the UK, Netherlands, Israel, and Taiwan).

Second, we identify and directly model heterogeneity in the mortality-temperature response function within each age category to account for the benefits of adaptation. Specifically, we allow the effect of temperature to vary as a function of climate (Barreca et al., 2015; Auffhammer, 2018) and income per capita (Hsiang and Narita, 2012; Burgess et al., 2014). These variables were carefully chosen based the intersection of prior evidence from the literature, economic theory, and variables that are included in standard projections of the global economy developed for integration with physical climate models (O'Neill et al., 2014). To condition on these variables globally, including in regions of the world with limited data availability, we construct new high-resolution datasets using satellite imagery. We find that there is substantial heterogeneity in the mortality-temperature relationship: moving from the poorest to richest tercile in our sample saves on average 0.54 deaths per 100,000 at 35°C. Similarly, moving from the coldest to hottest tercile of long-run average temperature saves on average 0.25 deaths per 100,000 at 35°C.[4]

The estimated effects of climate and income on the mortality-temperature relationship enable two exercises that are at the heart of the paper's mission. The first is that we use these covariates to "interpolate" mortality-temperature response functions for populations outside of our sample, thereby "filling in" missing information regarding sensitivity to climate for the 44% of people for whom mortality data are unavailable. This allows for a globally comprehensive, representative, and spatially heterogeneous characterization of the current mortality-temperature relationship. We estimate the *global average treatment effect* today for an additional 35°C day (relative to day at 20°C) is 2.94 deaths per 100,000, which is approximately 6-7 times larger than the average treatment effect from the regions of the world where data are currently available. Part of this difference is explained by the fact that data are disproportionately available in wealthier countries where sensitivities to temperature are lower, so exclusively relying on data from these countries understates the global costs imposed by high temperatures. An additional benefit to our explicit modeling of heterogeneity is that it allows us to predict mortality-temperature response functions individually for each of 25,000 compact geographic units engineered to have roughly equivalent populations and internally homogenous climates.

Our explicit modeling of heterogeneity additionally allows us to project the benefits of future climate adaptation in a realistic fashion. As populations within these small geographic units become richer and warmer in the future, they are likely to respond to temperature in a manner that is similar to the response of populations that are richer and warmer today. To project these patterns of future adaptation, we obtain standard projections[5] of the evolution of income and population for each of the 25,000 geographic units between now and the end of the century. We also compute the distribution of downscaled climates that each population will experience in each simulation of multiple possible greenhouse gas emissions scenarios. These data allow us to estimate how

[4]This finding is qualitatively similar to findings in the previous literature (e.g., Barreca et al., 2016), but is estimated with a far more globally representative sample than in any prior analysis.

[5]See Section 4.3 for details.

the mortality-temperature relationship will likely evolve as climate and income change over the remainder of the century for a range of plausible future pathways. This is a substantial step forward from assuming that response functions are constant over time, as has been the norm in the literature to date. A key finding from this exercise is that the mortality consequences of an additional 35°C day, relative to day at 20°C, are projected to decline by 70% (resulting in 0.89 deaths per 100,000) by the end of the century. This decline is due to both the protective effects of higher incomes, as well as the costly adaptations that individuals are predicted to undertake in response to warmer climates.

Characterizing the benefits of adaptation that people undertake in response to higher temperature is important, as failing to account for these changes may over-estimate future impacts. However, resources expended by future populations on adaptation could alternatively have been devoted to other economic activities, such as consumption or investment. A full accounting of the economic burden of warming must account for the opportunity costs of all resources used to achieve these reduced sensitivities to temperature through adaptive adjustments. Empirically measuring all of these costs directly has not been achieved in the economics literature because the range of potential responses to warming—whether defensive investments (e.g. building cooling centers) or compensatory behaviors (e.g. exercising earlier in the morning)—is enormous, making enumeration of these costs extraordinarily challenging. Indeed, the previous literature has frequently noted that adaptation will involve costs, and occasionally produced partial costs estimates (e.g. Deschênes and Greenstone, 2011; Barreca et al., 2016) or estimates of total impacts net of costs (e.g. Schlenker, Roberts, and Lobell, 2013; Deryugina and Hsiang, 2017), but has made little progress in building a comprehensive empirical measure of direct costs borne in order to adapt to warming.

Within this context, the third feature of our analysis is the development of a general revealed preference method capable of bounding the full adaptation costs that populations will incur to obtain the adaptation benefits that we project. Although adaptation costs cannot be directly observed, they can be inferred by observing the climates under which they are adopted, since adaptations will only be adopted when their costs are less than or equal to their benefits. For intuition, consider Seattle, WA and Houston, TX, which have similar income levels, institutions, and other factors (in a global sense), but have very different climates; on average Seattle has just 0.001 days per year where the average temperature exceeds ≈32°C, while Houston experiences 0.31 of these days annually.[6] Houston has adapted to this hotter climate, evidenced by the fact that a day above 32°C produces $\frac{1}{40}^{\text{th}}$ of the excess mortality in Houston than it does in Seattle (Barreca et al., 2016). If these outcomes are the result of revealed preferences, then it must be the case that the costs required to achieve Houston-like mortality are not worth it for Seattle, since these extreme temperatures only occur $\frac{1}{300}^{\text{th}}$ as often. Anecdotal support for this intuition is provided by the observation that the air conditioning penetration rate is roughly 27% in Washington state and 100% in Texas.[7]

[6]These values are calculated using the data of Barreca et al. (2016) for the years 2000-2004.

[7]2000-2004 average.

We leverage this intuition to build a formal theoretical framework that allows us to derive precise analytical expression for the marginal costs of adaptations undertaken to avoid mortality from climate. Further, we show how these marginal costs can be recovered from data since all populations must be indifferent on the margin between adopting the adaptations of populations that are just slightly warmer or cooler on average. We are able to directly estimate the gradient in these marginal benefits which must equal the gradient of marginal costs in equilibrium.[8] After reconstructing the profile of marginal costs, we integrate them to reconstruct the total cost of any change in adaption levels. Thus, in our projections of the future, we are able to simultaneously estimate how populations will reduce direct mortality from the climate associated with their projected patterns of adaptation, while also tracking the costs incurred in order to achieve these adaptation benefits. Importantly, our approach allows for an arbitrarily large number of unknown adaptive adjustments, and it accounts for the possibility that some adaptations generate consumption value that is independent of their mortality benefits (e.g. the consumption value of air conditioning).

Together, these three features of our analysis allow us to develop measures of the full mortality-related costs of climate change (and the associated damages from a marginal ton of CO_2 emissions) for the entire world, reflecting both the direct mortality costs (accounting for adaptation) and all adaptation costs. As a starting point, we note that allowing for adaptation has a substantial impact on the estimated direct mortality impacts of climate change. Specifically, we estimate that, without taking account of adaptation, as is the norm in the literature (Deschênes and Greenstone, 2011; Hsiang et al., 2017), the mortality cost of climate change will be 125 deaths per 100,000 in 2100. However, after accounting for the effects of rising future income and projected climate adaptation, this value falls to 28 deaths per 100,000. Of course, these are just the direct costs; implementation of our revealed preference approach indicates that societies will incur adaptation costs equal in value to 16 deaths per 100,000 in excess mortality risk. Thus, the total mortality burden of climate change is projected to be worth 44 deaths per 100,000 at the end of the century.[9] Because it is not standard, and somewhat unappealing, to use mortality risk as the numeraire good in valuations, we convert these costs into dollars using various conversions, such as the VSL or the value of life years lost. In our preferred estimates under a low mitigation climate scenario (RCP8.5), we estimate that the total mortality related costs of climate change amount to 6.8% of projected global GDP during 2080-2100.[10] Projected adaptation costs make up one third of this economic burden. Our full set of projections include over 10 billion

[8]We note that Schlenker, Roberts, and Lobell (2013), Guo and Costello (2013) and Deryugina and Hsiang (2017) exploit similar arguments regarding the equality of marginal adaptation costs and marginal adaptation beneifts, and we describe in detail in Section 4.2 how our approach relates to these earlier contributions.

[9]This is a median estimate across a range of values from 28 different climate models.

[10]We explore a range of valuation assumptions in the results section, including high and low VSLs, age-adjustment, and assumptions about varying the VSL over space. The values presented here are for the U.S. EPA VSL, and adjust for age at death. We assign each local region a VSL based on an income elasticity of one relative to the U.S. EPA value, and then apply the global median VSL value to all impact regions in each year. Details on this procedure are in Section 4.5.

region-by-day-by-climate model probability distributions of excess mortality for each scenario. While rich in information, these results are difficult to incorporate into future studies without some simplification. Thus, we use our projections to construct a simple "mortality damage function" following the approach in Hsiang et al. (2017), which can be easily inserted into integrated assessment models of the global economy used by other researchers.

The rest of this paper is organized as follows: Section 4.2 outlines a conceptual framework for the problem of projecting climate damages into the future, accounting for adaptation and its cost; Section 4.3 describes the data used in the estimation of impacts and in the climate change projected impacts; Section 4.4 provides details of each of the steps of the overall methodology; Section 4.5 describes the main results of the analysis; and Section 4.6 concludes.

4.2 The costs and benefits of endogenous adaptation to climate

Our aim is to calculate the full mortality costs of climate change, accounting for both the costs and benefits of adaptive behavior. Adaptation benefits, the more heavily studied of the two, remain rarely accounted for when empirically-derived climate response functions are used to generate climate change impact predictions. While there is increasing evidence that populations adapt to warming temperatures through investing in compensatory behaviors or technologies (e.g. Barreca et al., 2016, 2015; Graff Zivin and Neidell, 2014), previous literature has only recently begun to document and account for these benefits when making climate change projections (e.g. Auffhammer, 2018; Heutel, Miller, and Molitor, 2017). To the best of our knowledge, there exist no estimates of observable adaptation benefits at the global scale, either in the context of mortality risk, or in any other category of impact under climate change. Moreover, it is even more rarely acknowledged that achieving these adaptive benefits is costly; the costs necessary to realize adaptation have been left largely unaddressed in the climate change literature. Here, we develop a framework that allows us to theoretically ground and empirically estimate both the benefits and costs of adaptation.

Extremely few empirical studies explicitly address the theoretical importance of weighing adaptation benefits against associated costs when considering climate damages. Hsiang and Narita (2012) derive theoretically how an unobserved adaptation cost function could explain estimated heterogenous damages from hurricanes. Schlenker, Roberts, and Lobell (2013) derive a closed-form solution for the structure of these costs under specific log-linear functional form assumptions that are suitable for the context of crop yields and extreme heat, and Lobell et al. (2014) demonstrate that cross-sectional patterns of average yields in the US are consistent with the notion that farmers sacrifice average yields (a cost) in order to adapt to extreme heat. Studying the partial equilibrium of the California forestry sector, Guo and Costello (2013) point out that, on the margin, the value of marginal adaptations are zero when the outcome of in-

terest is a maximized quantity, since the marginal costs of adaptation must equal the marginal benefits. Recently, Deryugina and Hsiang (2017) generalized this result to a general equilibrium setting with an arbitrary number of unknown margins of adaptation, exploiting the equality of marginal benefits and costs of adaptation to identify the marginal product of climate using random weather variation.

Our theoretical contribution advances this literature in multiple ways. First, our result is derived from microfoundations and is generalizable in the sense that it does not rely on functional form, as was the case in Schlenker, Roberts, and Lobell (2013). Second, our framework extends beyond contexts where the outcome of interest is itself a maximized quantity, as was the case in Guo and Costello (2013) and Deryugina and Hsiang (2017), who rely heavily on the Envelope Theorem. This second point is crucial, because much of what we wish to learn about the economic costs of climate change examines outcomes that are clearly not maximized quantities (e.g. crime and conflict, agricultural yields, migration). Furthermore, we are the first to address tradeoffs in adaptation costs and benefits for a non-market outcome (mortality risk), whereas the method developed in Deryugina and Hsiang (2017) is restricted to market outputs with nonzero prices. Finally, we are the first to explicitly recover the structure of an adaptation cost function using data. Previous studies invoke the existence of adaptation costs in order to reconcile empirical observations, but these adaptation costs cannot be observed. Our approach does not illuminate unobserved adaptation costs directly, but we develop the following procedure to recover their exact structure implicitly.

Accounting for the full costs of climate change

We define the climate of a location as the joint probability distribution over a large vector of possible conditions that can be expected to occur at a given location over a specific interval of time. Following the notation of Deryugina and Hsiang (2017), let $\boldsymbol{C}$ be a vector of parameters describing the entire joint probability distribution over all relevant climatic variables. For example, $\boldsymbol{C}$ might contain the mean and variance of daily average temperature and rainfall, among other parameters. Weather realizations are a random vector $\boldsymbol{c}$ drawn from this distribution. Mortality risk is a function of both $\boldsymbol{c}$ and a vector of K endogenous social and economic conditions $\boldsymbol{b} = \{b_1, ..., b_K\}$. The vector $\boldsymbol{b}$ captures all choice variables available to individuals (except consumption of a numeraire good), including possible adaptive behaviors and investments that could interact with individuals' exposure to a warming climate, such as installation of air conditioning and the allocation of time spent indoors.

Climate change has the potential to influence mortality risk through two pathways. First, a change in $\boldsymbol{C}$ will directly alter realized weather draws, changing $\boldsymbol{c}$. Hsiang (2016) called this the "direct effect" because it captures the fact that any change in the probability distribution of weather events must generate a change in the distribution of events that individuals actually experience. Second, a change in $\boldsymbol{C}$ can alter individuals' beliefs about their likely weather realizations, shifting how they act, and ultimately changing their endogenous choice variables $\boldsymbol{b}$. This "belief effect" captures the fact that individuals adapt by preparing for the distribution of weather events and

corresponding direct effects that they expect, thus changing $\boldsymbol{b}$. Adjustments to $\boldsymbol{b}$ should be thought of as a generalization of what prior work refers to as any type of long-run adaptation (e.g. Mendelsohn, Nordhaus, and Shaw, 1994; Kelly, Kolstad, and Mitchell, 2005; Barreca et al., 2016), since these adaptations are necessarily factor reallocations based on beliefs about the climate (Deryugina and Hsiang, 2017). Therefore, since the climate $\boldsymbol{C}$ determines both $\boldsymbol{c}$ and $\boldsymbol{b}$, for notational simplicity we write

$$\boldsymbol{c} = \boldsymbol{c}(\boldsymbol{C}), \quad \boldsymbol{b} = \boldsymbol{b}(\boldsymbol{C})$$

noting that $\boldsymbol{c}$ is a random vector drawn from a distribution characterized by $\boldsymbol{C}$. At an initial time $t = 0$ with initial climate $\boldsymbol{C}_0$, we then can write the probability of death as

$$\Pr(\text{death at } t = 0 \mid \boldsymbol{C}_0) = f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_0)). \tag{4.1}$$

To date, most empirical estimates of climate impacts on social or economic conditions calculate these effects assuming no adaptation takes place (e.g. Deschênes and Greenstone, 2007; Houser et al., 2015). Specifically, these approaches calculate changes in an outcome variable imposed by changing the distribution of $\boldsymbol{c}$, assuming the choice vector $\boldsymbol{b}$ is evaluated at its initial values before the climate changes. This accounts only for the direct effect of the climate. Following this approach, the change in mortality risk incurred due to the direct effect of changes in the climate between period $t = 1$ and $t = 0$, but ignoring the fact that individuals' beliefs about their climate may also shift, is

$$f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_1)) - f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_0)). \tag{4.2}$$

One might also imagine isolating the belief effect of climate change on mortality risk. For example, we can envision the change in mortality risk that would arise from households in Seattle receiving the same amount of air conditioning that residents of Houston currently enjoy, without changing the actual temperatures experienced in Seattle. Because air conditioning is captured by elements in the choice vector $\boldsymbol{b}$, the benefits to mortality risk from this change would be expressed as

$$f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_0)) - f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_0)) \tag{4.3}$$

where $\boldsymbol{C}_0$ denotes the climate of Seattle and $\boldsymbol{C}_1$ denotes the climate of Houston.

In reality, we do not observe Houston-level air conditioning in Seattle because rational endogenous allocations $\boldsymbol{b}$ will match population's actual beliefs over the $\boldsymbol{c}(\boldsymbol{C})$ values they are likely to experience. Thus, under climate change, we would expect populations to update their behaviors and technologies $\boldsymbol{b}$ as their beliefs about $\boldsymbol{C}$ evolve. To the extent that populations may attenuate mortality damages by adjusting $\boldsymbol{b}$ in response to climate change, these adaptations will generate benefits by partially countering the direct effect of climate on mortality. Accounting for both the belief effect and the direct effect simultaneously would produce a more realistic estimate for the change in mortality risk due to climate change

$$f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) - f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_0)). \tag{4.4}$$

Note that if the climate is changing such that the direct effect of $\boldsymbol{C}_1$ is more damaging than $\boldsymbol{C}_0$, then the adjustment of $\boldsymbol{b}$ will generate *benefits of adaptation* since these damages will be partially mitigated. Stated another way, the change in Equation (4.2) will be larger than the change in Equation (4.4).[11]

Several analyses have estimated reduced-form versions of Equation (4.4), confirming that accounting for endogenous changes to technology, behavior, and investment may mitigate the direct effects of climate in a variety of contexts (Schlenker and Roberts, 2009; Burgess et al., 2014; Hsiang and Narita, 2012; Hsiang and Jina, 2014; Barreca et al., 2015; Heutel, Miller, and Molitor, 2017; Auffhammer, 2018). Importantly, however, while this approach accounts for the *benefits* of adaptation, it does not account for its *costs*. If adjustments to $\boldsymbol{b}$ were costless and provided protection against the climate, then we would expect universal uptake of highly adapted values for $\boldsymbol{b}$ so that all populations would be inoculated against the climate. But we do not observed this to be true: Seattle does not have air conditioning levels that mirror levels in Houston. The previous studies that demonstrate benefits of adaptation do so by demonstrating reduced sensitivity to marginal environmental changes ($\frac{\partial f}{\partial \mathbf{c}}$) in more adverse climates (Houston) and larger sensitivity in less adverse climates (Seattle). Carleton and Hsiang (2016) document that such wedges in observed sensitivities to climate—which they call "adaptation gaps"—are a pervasive and unexplained feature of the broader climate damages literature. We follow Schlenker, Roberts, and Lobell (2013), Houser et al. (2015) and Deryugina and Hsiang (2017) by arguing that the *unobserved costs of adaptation* are what maintain this cross-sectional structure of sensitivities to climate. In order for the lower adoption of air conditioning in Seattle to be rational when high levels of air conditioning provide large benefits in Houston, it must be that there are costs associated with high levels of air conditioning that are not justifiable in the climate of Seattle.

The compensatory investments that individuals incur in order to adapt take many forms, from investing in air conditioning to spending less time outdoors. We develop a generalizable revealed preference approach that allows us to empirically recover the sum of all unobserved costs incurred from adjustment of $\boldsymbol{b}$ as the climate warms. Importantly, our approach does not require that we observe these adjustments, making us agnostic to which specific forms of adaptation are utilized, and we need not attempt to enumerate all potential pathways of adaptation, which we think is likely impossible. Rather, we derive sufficient statistics for capturing the full nonlinear cost of adaptation $A(\boldsymbol{b})$ by micro-founding observed patterns of sensitivity to climate. Once we are able to account for these costs, we can add them to our estimated direct and belief effects of the climate in order to capture the full value of excess mortality risk imposed by

[11] Deryugina and Hsiang (2017) note that, as estimated in the empirical literature, this inequality might not appear *prima facie* always true. However, as they demonstrate, this is a result of inappropriate pooling of locations across different $\boldsymbol{C}$ in the estimation of Equation (4.2).

climate change, net of the costs and benefits of adaptive adjustments

$$climate\_change\_damages = VSL \times \underbrace{[f(\boldsymbol{b}(\boldsymbol{C}_1), \boldsymbol{c}(\boldsymbol{C}_1)) - f(\boldsymbol{b}(\boldsymbol{C}_0), \boldsymbol{c}(\boldsymbol{C}_0))]}_{\text{direct + belief effects}} + \underbrace{A(\boldsymbol{b}(\boldsymbol{C}_1)) - A(\boldsymbol{b}(\boldsymbol{C}_0))}_{\text{adaptation costs}} \tag{4.5}$$

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

In the next subsection, we build on the theoretical framework developed in Deryugina and Hsiang (2017), using revealed preference intuition from Schlenker, Roberts, and Lobell (2013), to derive adaptation costs as a function of climate. While related to these previous analysis, the following approach is novel because it applies in a context where the outcome (mortality risk) is not a directly maximized quantity (as in Deryugina and Hsiang, 2017) nor a market good (as in Schlenker, Roberts, and Lobell, 2013); furthermore, it allows us to trace out the otherwise unobservable cost function $A(\boldsymbol{b}(\boldsymbol{C}))$, which to our knowledge has never previously been accomplished.

A revealed preference approach to estimate adaptation costs

Let the probability of death $f(\boldsymbol{b}, \boldsymbol{c})$ depends both on weather realizations $\boldsymbol{c}$ and on choice variables $\boldsymbol{b}$, just as above. Taking account of this temperature-dependent mortality risk, agents will maximize expected utility based on the climatological distribution of $\boldsymbol{c}$ that they expect. We assume agents have rational expectations and can integrate over the distribution of realizations of $\boldsymbol{c}$ for each climate to compute the expected probability of death conditional on the climate and their actions

$$\mathbb{E}_{\boldsymbol{c}}[f(\boldsymbol{b}(\boldsymbol{C}), \boldsymbol{c}(\boldsymbol{C})) \mid \boldsymbol{C}] = \tilde{f}(\boldsymbol{b}(\boldsymbol{C}), \boldsymbol{C}). \tag{4.6}$$

Let agents derive utility $u(.)$ both from consumption of a numeraire good 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). Agents then solve for the optimal $\boldsymbol{b}$, subject to a budget constraint

$$\max_{x,\boldsymbol{b}} \; [1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})] \, u(x, \boldsymbol{b}) \quad \text{s.t.} \quad h(\boldsymbol{b}) + x = Y \tag{4.7}$$

where $h(\boldsymbol{b})$ is the pecuniary cost of adaptation and Y is income. Summing across the $K+1$ first order conditions for Equation (4.7) and exploiting the definition of the value

of the statistical life (VSL), we know that the optimal choice vector $\boldsymbol{b}^*$ must satisfy:[12]

$$-VSL \sum_k \frac{\partial}{\partial b_k} \underbrace{\tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}_{\substack{\text{expected} \\ \text{mortality} \\ \text{given } \boldsymbol{b}}} = \sum_k \frac{\partial}{\partial b_k} \underbrace{\left[\underbrace{h(\boldsymbol{b}^*)}_{\substack{\text{pecuniary} \\ \text{costs of } \boldsymbol{b}}} - \underbrace{u(x^*, \boldsymbol{b}^*) \frac{1}{\partial u(x^*, \boldsymbol{b}^*)/\partial x}}_{\text{non-mortality benefits of } \boldsymbol{b}} \right]}_{\text{net cost of } \boldsymbol{b} = A(\boldsymbol{b}^*)} \tag{4.8}$$

The lefthand side of Equation 4.8 is the sum of marginal changes to expected mortality risk for adjustments to the choice vector, times minus the VSL. If there exist potential benefits of additional adaptation on the margin, in the form of further reducing mortality risk by adjusting $\boldsymbol{b}$, then the summation will have an overall negative value. Multiplied by minus the VSL, this represents the total marginal benefit of adjusting $\boldsymbol{b}$ in terms of reducing mortality risk. The righthand side of Equation 4.8 has two parts. The first term inside the square brackets, $h(\boldsymbol{b}^*)$, represents all pecuniary expenditures required to achieve the optimal set of adaptive actions $\boldsymbol{b}^*$, such as spending on air conditioning. The second term in the brackets represents minus the dollar value of all direct utility benefits derived from each adaptive action $\boldsymbol{b}^*$, such as the utility of enjoying air conditioning expressed in dollars (based on willingness-to-pay). The sum of these two terms can thus be interpreted as the *net cost* of achieving $\boldsymbol{b}^*$, from the perspective of adapting to climate in order to reduce climate-caused mortality. For simplicity, we denote the total net cost as $A(\boldsymbol{b})$, matching the earlier notation in Equation 4.5 (for our purposes, it is neither important or empirically feasible to separately identify these two components of net cost, and throughout our analysis adaptation costs should be interpreted as net of any direct non-pecuniary benefits). Substituting $A(\boldsymbol{b})$ into Equation (4.8) indicates clearly that the optimum choice vector $\boldsymbol{b}^*$ equalizes total marginal

[12]The first order conditions take the form:

$$[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})] \frac{\partial u(x, \boldsymbol{b})}{\partial x} = \lambda, \qquad [1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})] \frac{\partial u(x, \boldsymbol{b})}{\partial \boldsymbol{b}} - \frac{\partial \tilde{f}(\boldsymbol{b}, \boldsymbol{C})}{\partial \boldsymbol{b}} u(x, \boldsymbol{b}) = \lambda \frac{\partial h(\boldsymbol{b})}{\partial \boldsymbol{b}}$$

Note that because $\boldsymbol{b}$ is a vector, the second equation above represents K separate equations. Substitution gives:

$$\underbrace{-\frac{\partial \tilde{f}(\boldsymbol{b}, \boldsymbol{C})}{\partial \boldsymbol{b}} u(x, \boldsymbol{b})}_{\substack{\text{marginal benefit of a} \\ \text{change in } \boldsymbol{b} \text{ (in utils)}}} = \underbrace{[1 - \tilde{f}(\boldsymbol{b}, \boldsymbol{C})] \left[\frac{\partial u(x, \boldsymbol{b})}{\partial x} \frac{\partial h(\boldsymbol{b})}{\partial \boldsymbol{b}} - \frac{\partial u(x, \boldsymbol{b})}{\partial \boldsymbol{b}} \right]}_{\text{net marginal cost of a change in } \boldsymbol{b} \text{ (in utils)}}$$

The righthand side of the above expression represents the marginal cost of changing b, valued in utils, net of any change in utility realized from direct utility effects of adaptation. We call this object "net marginal costs". The lefthand side is simply the marginal mortality benefit realized from a change in $\boldsymbol{b}$. Multiplying by $\frac{1}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\frac{\partial u(x,\boldsymbol{b})}{\partial x}}$ translates this expression into dollars, and summing across all choice variables k in the vector $\boldsymbol{b}$ gives the expression in Equation 4.8. This translation arises from the standard Value of a Statistical Life (VSL) definition: $VSL = \frac{u(x,\boldsymbol{b})}{[1-\tilde{f}(\boldsymbol{b},\boldsymbol{C})]\frac{\partial u(x,\boldsymbol{b})}{\partial x}}$. Note that the VSL, which may depend on both income and the climate, is held constant throughout our conceptual framework. In our empirical derivation, we allow the VSL to evolve with income, under a range of different choices for the income elasticity. For simplicity, the VSL does not vary with climate in our projection of climate change impacts.

benefits of adaptation, in terms of mortality risk, with its total marginal net costs:

$$\underbrace{-VSL \sum_k \frac{\partial \tilde{f}(\boldsymbol{b}^*, \boldsymbol{C})}{\partial b_k}}_{\substack{\text{marginal benefit} \\ \text{of } \boldsymbol{b} \text{ (in dollars)}}} = \underbrace{\sum_k \frac{\partial A(\boldsymbol{b}^*)}{\partial b_k}}_{\substack{\text{net marginal costs} \\ \text{of } \boldsymbol{b} \text{ (in dollars)}}} \tag{4.9}$$

Equation (4.9) explains many patterns of adaptation characterized in the previous studies mentioned above, and can reconcile the "adaptation gaps" documented in Carleton and Hsiang (2016). Populations will adjust every dimension of $\boldsymbol{b}^*$ in response to their climate $\boldsymbol{C}$ such that their total marginal costs equal their marginal benefit as adaptive strategies. Unless the marginal cost of adaptation on the right hand side is zero for every dimension of $\boldsymbol{b}^*$ (i.e. adaptation is free), then we should expect to observe heterogeneity in the marginal benefits of adaptation across different climates.

Because Equation (4.9) implicitly defines an optimal vector of choice variables as a function of climate and income (which sets the marginal utility of consumption through the budget constraint), it is useful to write the optimum as an implict function $\boldsymbol{b}^* = \boldsymbol{b}^*(\boldsymbol{C}, Y)$.

To calculate the full mortality costs of climate change in Equation (4.5), we must recover the change in total unobserved adaptation costs as the climate warms from $\boldsymbol{C}_0$ to $\boldsymbol{C}_1$:

$$A(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) - A(\boldsymbol{b}^*(\boldsymbol{C}_0, Y)) \tag{4.10}$$

where Y is written as if it is "held fixed" because it is exogenously determined, although it can be varied continuously in actual calculations.[13] This is the missing term in prior empirical damage estimates that account for adaptation, because neither $\boldsymbol{b}$ nor A are directly observable. However, we can compute the magnitude of this term by exploiting Equation (4.9), since by revealed preference we know how large total unobserved costs must be on the margin in order to support the adaptation benefits we observe.

To compute the total cost of adaptations that are endogenously adopted by populations as they experience a transition between two climates that differ by a non-marginal amount, we integrate marginal costs. Similar to the intuition developed in Schlenker, Roberts, and Lobell (2013) and Deryugina and Hsiang (2017), we integrate a sequence of marginal costs that are individually implied by the observed gradient in marginal benefits of adaptation. Utilizing the Gradient Theorem, the Chain Rule, and substitution from Equation 4.9, we can rewrite the non-marginal cost of adapting to climate

[13]Note one key difference between this expression and the cost term in Equation 4.5. Here, the cost function $A(\cdot)$ depends both on the climate and on income, while in Equation 4.5, income doesn't enter. This difference arises because in our formal derivation we impose a budget constraint, which influences the ability of agents to reallocate resources among choices $\boldsymbol{b}$, as well as the numeraire good x. For example, as individuals get richer, they may invest in air conditioning even under a constant climate. Under this derivation, we can account for changes in income that unfold in coming centuries.

changes as:

$$\boldsymbol{A}(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) - \boldsymbol{A}(\boldsymbol{b}^*(\boldsymbol{C}_0, Y)) = \int_{\boldsymbol{C}_0}^{\boldsymbol{C}_1} \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}, Y))}{\partial \boldsymbol{C}} d\boldsymbol{C}$$

$$= \int_{\boldsymbol{C}_0}^{\boldsymbol{C}_1} \sum_k \frac{\partial A(\boldsymbol{b}^*(\boldsymbol{C}, Y))}{\partial \boldsymbol{b}_k^*(\boldsymbol{C}, Y)} \frac{\partial \boldsymbol{b}_k^*(\boldsymbol{C}, Y)}{\partial \boldsymbol{C}} d\boldsymbol{C}$$

(substitution from Equation 4.9)
$$= -\int_{\boldsymbol{C}_0}^{\boldsymbol{C}_1} VSL \sum_k \frac{\partial \tilde{f}(\boldsymbol{b}, \boldsymbol{C})}{\partial \boldsymbol{b}_k^*(\boldsymbol{C}, Y)} \frac{\partial \boldsymbol{b}_k^*(\boldsymbol{C}, Y)}{\partial \boldsymbol{C}} d\boldsymbol{C}$$

$$= -\int_{\boldsymbol{C}_0}^{\boldsymbol{C}_1} VSL \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}, Y))}{\partial \boldsymbol{C}} d\boldsymbol{C} \tag{4.11}$$

where the the final form represents an object that can be empirically recovered from the data, without every directly observing adaptive actions or their costs. The relationship between mortality risk and the climate—the gradient $\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}$—is a sufficient statistic for the total change in mortality risk arising from all unobservable adaptive behaviors—the Jacobian $\frac{\partial \boldsymbol{b}^*}{\partial \boldsymbol{C}}$. Because we can empirically derive the relationship between mortality risk and the climate, we can numerically compute the value of the integral in Equation 4.11, despite our inability to measure any of the actions in $\boldsymbol{b}$. The details of the implementing Equation 4.11 are in Appendix C.1.

Figure 4.1 shows a cartoon to help communicate the intuition of our model with a single endogenous choice variable b and a single parameterization of the climate, e.g. average temperature. On the vertical axis are costs, while the horizontal axis is quantity of adaptation. In panel **A**, two adaptation benefit curves are shown for two distinct average temperatures, C_0 and C_1, where $C_1 > C_0$. Total gain to the individual is increasing downward (negative costs) and to the right (higher adaptation). For each climate, the shape of the benefit curve is determined by the shape of the marginal mortality benefit of b, the lefthand side of Equation 4.9, which in turn depends on the climate through $\tilde{f}(b, C)$. In the example shown here, the marginal mortality benefit of any additional adaptation is smaller in the colder location at all levels of b than it is in the hotter location. This is consistent with our empirical results shown below, and is intuitive, as more lives are saved through adaptation when exposure to health-threatening temperatures is on average higher. In equilibrium, each location chooses an optimal level of adaptation $b^*(C)$ at which the marginal mortality benefits exactly equal the marginal net costs, i.e. where the location-specific benefits curve is tangent to the cost curve. If the climate warms from C_0 to C_1, adaptation would increase to b_1^*, the point at which the marginal costs of additional adaptation exactly offset the marginal benefits in the hotter climate. One can imagine many such transitions, for infinitesimally small changes in climate, where the sequence of equilibria trace out the entire net cost function $A(b)$.

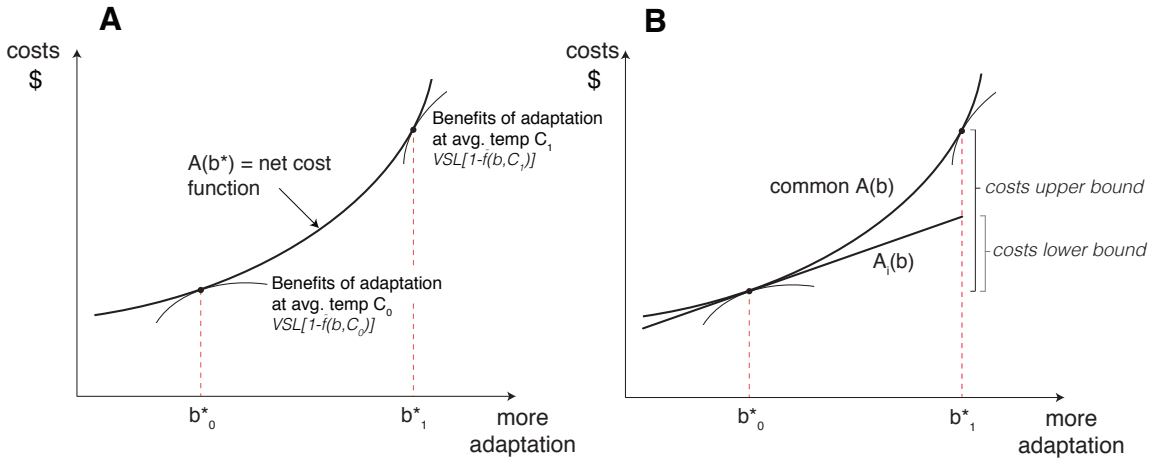

Figure 4.1: Revealed preference approach used to estimate costs of adaptation

Panel **A** shows two mortality benefit curves for climates with temperatures $C_0 < C_1$, where welfare is increasing down and to the right (lower costs and more adaptation). The function $A(b^*)$ traces out all points of tangency between the unobservable net cost function, and the mortality benefit curve, across all locations with heterogeneous climates. Panel **B** shows the two adaptation cost estimates we empirically derive, for a location moving from climate C_0 to climate C_1 under climate change. The upper bound is the result of integrating marginal changes along a universally shared net cost function, assuming agents reoptimize continually as their climate warms and that marginal costs of adaptation, conditional on climate, are common across all populations. The lower bound is the result of integrating marginal changes along the tangency line at the initial period's benefit curve, assuming that agents optimally choose adaptation behaviors and technologies at the initial climate. This lower bound allows for the possibility that each location faces distinct marginal costs of adaptation, even conditional on climate.

In our empirical implementation, we compute the cost of adaption by integrating the cost function under two different sets of assumptions, shown in in panel **B** of Figure 4.1. In the first case, we assume that *all populations share a common, possibly nonlinear, cost function* $A(b)$ and continuously adapt as the climate evolves, such that all populations are always at a tangency point with this common function $A(b)$. As the climate warms, populations adopt the adaptive strategies of their nearby neighbors and experience the adaptation costs experienced by those neighbors. If those neighbors experienced marginal costs that were higher, we assume that the warming populations become subject to that higher marginal cost. Integration of Equation 4.11 under this assumption produces an upper bound on costs, as in this case all benefits of adaptation are completely subsumed by compensatory costs. A second alternative assumption is that populations *do not share a common adaptation cost function*, rather each population achieves equilibrium based on their own unique cost function $A_i(b)$. This assumption would be appropriate if there are fundamental reasons (perhaps geographic, cultural or political) why the same adaptation technology has different costs when utilized by different populations. In such a scenario, the marginal costs exhibited by neighboring

populations are not indicative of the costs a warming population would experience if employing the same adaptations, thus it would be inappropriate to allow marginal costs to evolve with warming along a curve defined by warmer neighboring populations. In this case, we have much less information about the adaptation costs a population would experience with warming that is out of their historical experience, since costs of neighbors are not a valid proxy. Thus, in this scenario, we make the simplest assumption that the marginal costs of adaptation for each population are fixed across all b, such that total cost $A_i(b)$ increases linearly based on the marginal cost currently observed for that population. This approach produces lower estimates for the cost of adaptation, as we empirically uncover lower marginal costs of adaptation today in locations which are cooler – fixing these marginal costs at lower levels before the climate warms gives rise to a lower bound on adaptation costs. Notably, one implication of this assumption is that the costs of adapting from b_0 to b_1 for population i need not be the same as the savings achieved by population j that is reducing adaptation from b_1 to b_0. In order for such a condition to hold in equilibrium, it must be the case that populations cannot freely trade factors in b, otherwise these different marginal prices would necessarily equalize.

We use these two different assumptions about the evolution of marginal costs, combined with our result in Equation 4.11, to bound adaptation costs experienced by populations as their climate warms. All components required to empirically derive both of these cost bounds are estimated or directly observable, as detailed in Appendix C.1.

4.3 Data

Mortality data

Our mortality data represent 41 countries.[14] In some cases our data represent the universe of reported deaths in those countries, while in others (namely China and India), data are representative samples, as no vital statistics registry system exists. Combined, our dataset covers mortality outcomes for 56% of the global population. Spatial coverage, resolution, and temporal coverage are shown in Figure 4.2, and each dataset is described in Table 4.1. Data are drawn from multiple, often restricted, national and international sources, with details on each given in Appendix C.2. All mortality datasets contain information on deaths per 100,000 population from all causes at a monthly or annual frequency. All datasets contain age-specific mortality rates at various resolutions with the exception of India, which only provides all-age mortality. We harmonize these diverse sources into a single multi-country panel dataset of age-specific annual mortality rates, using three age categories: 0-4, 5-64, and over 64.

[14]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, owing to the fact that the India data, also used in Burgess et al. (2014), do not have information on age-specific mortality rates.

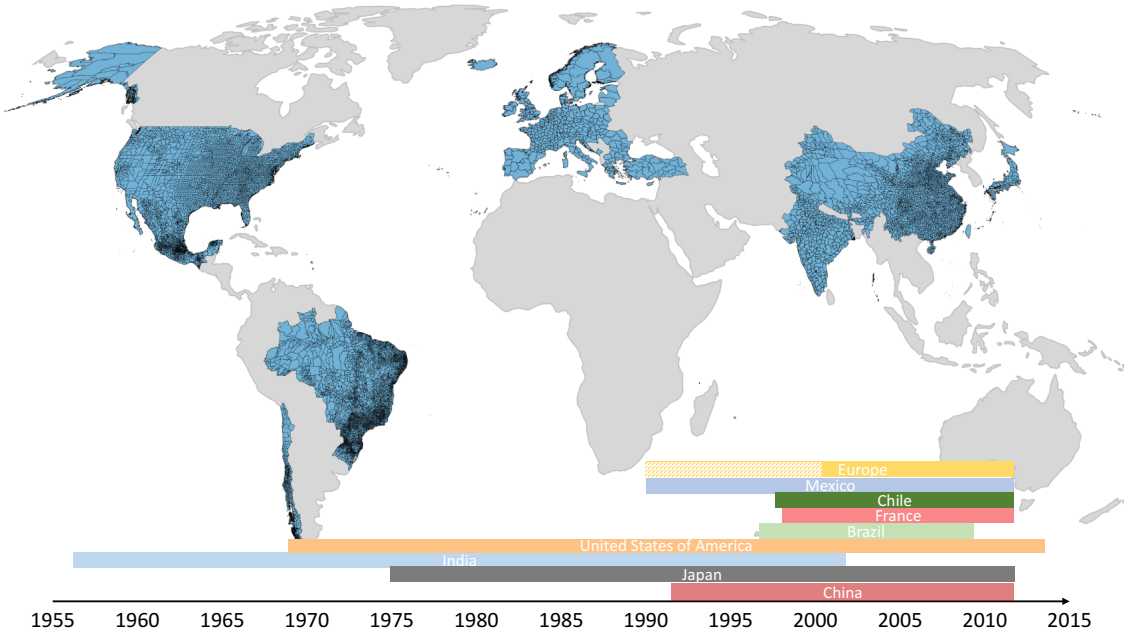

Figure 4.2: Spatial and temporal coverage of mortality statistics

Regression estimates of the temperature-mortality relationship rely on mortality statistics from all countries shown in blue, at the spatial resolution indicated by the black boundary lines. Temporal coverage for each country is shown under the map.

Country	N	Spatial scale	Years	Age categories	Avg. mortality rate[a,b]		Global pop. share	Covariate values		
					All-age	Over 65		GDP p.c.	$\bar{T}$	# Days $> 28°C$
Brazil	208,560	ADM2[c]	1997-2009	0-4, 5-65, 65+	488	4,107	0.027	6,433	24.1	44.5
Chile	13,314	ADM2	1997-2010	0-4, 5-65, 65+	554	4,186	0.002	8,943	13.2	0
China	7,488	ADM2	1991-2010	0-4, 5-65, 65+	635	7,507	0.183	3,793	14.1	24.1
EU[2]	11,733	NUTS2[d]	1990[e]-2010	0-4, 5-65, 65+	806	5,233	0.091	26,273	10.5	1.2
France	2,444	ADM2	1998-2010	0-20, 20-65, 65+	531	1,968	0.009	28,181	11.9	0.3
India	13,246	ADM2	1957-2001	ALL	540	N/A	0.169	1,092	25.2	118.6
Japan	4,968	ADM1	1975-2010	0-4, 5-65, 65+	787	4,361	0.017	23,216	14.2	8.2
Mexico	139,615	ADM2	1990-2010	0-4, 5-65, 65+	458	3,723	0.017	7,707	20.9	27.6
USA	399,590	ADM2	1968-2010	0-4, 5-65, 65+	1,013	5,256	0.043	28,836	13.5	11.6
Global	800,958	-	-	-	748	4,663	0.559	18,222	17.7	-

Table 4.1: Mortality data summary statistics

[a] In units of deaths per 100,000 population.

[b] Mortality rate is winsorized at the 1 % level at the high end of the distribution only.

[c] ADM2 refers to the second level of administrative region in each country, while ADM1 refers to the first level. For example, in the USA, ADM2 refers to the county, and ADM1 to state.

[d] European Union data are grouped here as it was obtained from a single source. Mortality data from 33 countries are included in this sample. Detailed description of the countries with data within this region is presented in Appendix C.2.

[e] 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 C.2 for details.

Historical climate data

We perform analyses with two separate groups of data on precipitation and temperature from independent sources. First, we use a reanalysis product, the Global Meteorological Forcing Dataset (GMFD), which relies on a climate model in combination with observational data to create globally-comprehensive gridded data on daily temperature and precipitation. Second, we test the robustness of our results to the use of products which generate gridded output from spatial interpolation of observational data. Here, we use the Berkeley Earth Surface Temperature dataset (BEST) in combination with the University of Delaware precipitation dataset (UDEL). A brief description of the data is provided in Table 4.2, with full data descriptions in the Appendix. The final column in Table 4.2 signifies the way the data are grouped for analysis, i.e., "group 1" data sets are used as the sources of temperature and precipitation in a specific regression. Results using the GMFD data are treated as the primary results throughout the paper, due to their higher spatial and temporal resolution for both rainfall and temperature, and because the climate projections we use for future climate scenarios are bias-corrected to the GMFD dataset. Regression results using the "group 2" data are shown as robustness.

Gridded daily temperature data are aggregated to the same administrative level as the mortality data by taking any nonlinear transformations at the pixel level before spatially averaging, either using population or area weights.[15] Precipitation quadratic polynomials are similarly calculated and weighted averages are taken over administrative units.

Covariate data

Our analysis exploits heterogeneity in the mortality-temperature relationship as a function of two long-run covariates: average temperature and income. Data sources for these two covariates are described below.

Temperature Data on long-run average temperature are used as one of our covariates to explain spatial heterogeneity in the mortality-temperature relationship. Our average temperature data are calculated from the same temperature datasets described above, and used in the main analysis to estimate a nonlinear mortality-temperature relationship (e.g. GMFD). Long-run averages are taken over every year in the mortality sample for a particular location. Note that this time scale will therefore vary across countries, as described in Table 4.1. These average values are calculated at the ADM1-

[15]Note that we take nonlinear transformations of the climate data at the pixel-day level before spatially aggregating to preserve the tails of the distribution within the administrative region. Failing to preserve this ordering in calculation biases results (Hsiang, 2016). We show robustness of the mortality-temperature relationship to four different nonlinear functional forms, all of which undergo the pixel-level transformation before being spatially averaged to administrative regions. In our main results, spatial averages across pixel values within an administrative region are conducted using population weights.

Dataset	Citation	Variables	Method	Resolution	Source	Group
GMFD	Sheffield, Goteti, and Wood (2006)	T, P	Reanalysis & interpolation	0.25°	Terrestrial Hydrology Group, Princeton	1
BEST	Rohde et al. (2013)	T	Interpolation	1°	Berkeley Earth	2
UDEL	Willmott and Matsuura (2014)	P	Interpolation	0.5°	University of Delaware	2

Table 4.2: Sources of climate data

We show robustness across multiple climate datasets, each of which uses a distinct methodology to generate globally-comprehensive gridded products. For the main analysis, the Global Meteorological Forcing Dataset (GMFD) is used.

level (e.g., province or state) in order to match the level at which the administrative income data are available.

Sub-national Incomes Sub-national incomes for all non-EU countries are obtained from Gennaioli et al. (2014), who collate sub-national incomes from administrative data around the world. These values are generally at the ADM1 (e.g. state) level. We match the ADM2-year observations in our sample to cross-sectional income values by assigning each region to the average ADM1 GDP per capita over the sample period. The data in Gennaioli et al. (2014) are typically not annual, as many are drawn from census data. The panel in Gennaioli et al. (2014) is also unbalanced, so we interpolate between years and take the average for a region at the ADM1 level over those interpolated data. Details of the coverage of both samples is given in Table C.1. Sub-national incomes for locations within the EU are obtained from Eurostat (Eurostat, 2013), which provides incomes at the NUTS2 level for each of the European countries in our dataset. All sub-national income data are in constant 2005 dollars PPP.

Projection data

Spatial data for projections We create a set of boundaries that define the spatial units onto which we extrapolate temperature-mortality sensitivities derived from our estimation, and for which we create location-specific projected damages of climate change. To do so, we utilize politically defined regions, as they form a better scale for analysis than regular grids due to their use in collecting socioeconomic data. Moreover, these regions are generally more 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 the Global Administrative Region dataset (Global Administrative Areas, 2012) to delineate boundaries, but require fewer than the approximately 295,000 spatial units present in that dataset. We thus create a set of 24,378 agglomerated regions that allow for greater comparability and computational feasibility than unagglomerated regions. We establish a set of criteria to create these regions that makes them approximately comparable with respect to population, and internally consistent with respect to mean temperature, diurnal temperature range, and mean precipitation. More detail of this process is given in Appendix C.4, along with a map of the regions.

Climate projections To comprehensively capture spatial heterogeneity in climate damages, we use a high-resolution (0.25° X 0.25°) set of global, bias-corrected climate projections produced by NASA Earth Exchange (NEX) Global Daily Downscaled Projections (GDDP) (Thrasher et al., 2012).[16] The NEX-GDDP dataset constitutes 21 climate projections, which are downscaled from the output of General Circulation Model (GCM) runs in the Coupled Model Intercomparison Project Phase 5 (CMIP5) archive (Taylor, Stouffer, and Meehl, 2012).[17] For each model's climate projection, we employ the results of a historical experiment, which simulates the response of the climate to historical forcing from 1850 to 2005 (data after 1981 are used in our analysis), and both the Representative Concentration Pathways 4.5 and 8.5 (RCP4.5 and RCP8.5) experiments. RCP4.5 represents a "stabilization" scenario in which total radiative forcing is stabilized around 2100 (Riahi et al., 2011). RCP8.5 simulates climate change under intensive growth in fossil fuel emissions from 2006 to the end of 21^{st} century. The gridded data from NEX-GDDP are aggregated to impact regions for further analysis on climate change impacts.

The ensembles of CMIP5 GCMs do not produce probability distributions of climate projections and they systematically underestimate the tail risks in future climate (Tebaldi and Knutti, 2007; Rasmussen, Meinshausen, and Kopp, 2016). To provide a probabilistic ensemble of climate projections, 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 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 (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. More detail on this method can be found in Appendix C.5. At end of century (i.e. 2080-2100), the median warming of global mean surface temperature across all GCM and surrogate models, relative to pre-industrial, is 2.7°C for RCP4.5 and 4.3°C for RCP8.5.

[16]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).

[17]Details of the models used are provided in Appendix C.5.

Population projections Projections of national populations are derived from the Organization for Economic Co-operation and Development (OECD) (Dellink et al., 2015) and International Institute for Applied Systems Analysis (IIASA) (Samir and Lutz, 2014) population projections 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. The population data are accessed from the SSP database (IIASA Energy Program, 2016). 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 generate local-level projections of climate change impacts, and to account for substantial heterogeneity in adaptation and exposure within countries, we need to assemble high-resolution information on on population distributions. To do so, we downscale the country-level projections from the SSPs using current LandScan estimates of populations that fall within our impact regions within each country. Populations for impact regions in countries or areas not given in the SSP Database are held constant at their LandScan estimated values. Thus, for any given impact region i in year t, population for scenario s ($P_{i,t,s}$) is given by

$$P_{i,t,s} = \begin{cases} \sum_{j \in J} d_{i,j} * \left(\hat{P}^{SSP}_{j,t,s} \frac{P_i^{LandScan}}{\sum_k d_{k,j} P_k^{LandScan}} \right), & \text{if } i \in J \\ P_i^{Landscan}, & \text{if } i \notin J \end{cases} \tag{4.12}$$

where $\hat{P}^{SSP}_{j,t,s}$ is the SSP population given for country j, $P_i^{LandScan}$ is the LandScan estimate for impact region i, $d_{i,j}$ is a dummy variable which is set to 1 when i is in country j and to 0 otherwise, and J is the set of all countries in the SSP Database. Note that while this approaches 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.

Income projections Projections of national per-capita income are also taken from the socioeconomic conditions making up the SSPs. Multiple models are used to estimate SSP trajectories. These models predict different subsets of countries around the globe, and describe different pathways within each shared scenario. Across the models used to estimate SSP trajectories, populations are fairly similar, so we merge all models by averaging estimates across all models that predict each country.

Only the IIASA GDP model and OECD Env-Growth provide GDP per capita 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 these two models to capture uncertainty within socioeconomic scenario. 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 in the given year for the baseline period, and allow this income to grow at the globally averaged rate. We smoothly interpolate between the time series data provided in 5-year increments in the SSP Database. For each 5-year segment, we calculate the average annual growth rate, and apply to produce each year's estimate.

As described above for population, our high-resolution analysis requires estimates of location-specific GDP within country borders. To generate such estimates, we allocate national GDP per capita across the impact regions within a country through a downscaling procedure that relies on night lights imagery from the NOAA Defense Meteorological Satellite Program (DSMP). Using available subnational income data from Gennaioli et al. (2014) in combination with higher-resolution income data from the United States, China, Brazil, and India, we empirically estimate the relationship between GDP per capita and nightlight intensity. We use this estimated relationship to allocate GDP data provided at the national level heterogeneously across impact regions within each country, based on relative intensity of night lights in the most recent DSMP images available (2013). While this approach models heterogeneity in income levels across impact regions, each region grows in the future at the same rate as the national country projection. This results in subnational, impact region level, GDP per capita projections that, like the population projections, fix the relative within-country distribution of income to the distribution observed today.

4.4 Methods

Dose-response function estimation

We begin by estimating a mortality temperature dose-response function across all countries in our sample by pooling subnational observations across the 41 nations for which we have mortality records.[18] We use daily variation in weather to identify a high-frequency response of mortality to temperature, following, for example, Deschênes and Greenstone (2011), Barreca et al. (2016), and Deryugina and Hsiang (2017). Our model is:

$$M_{aict} = f(T_{it}) + g^c(P_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{4.13}$$

where c denotes *country*, i denotes 2^{nd}*-level* administrative division (ADM2),[19] a indicates age category with $a \in A = \{\text{0-4}, \text{5-64}, 65+\}$ and t indicates years. P_{it} is cumulative annual precipitation, and it is modeled via $g^c(\cdot)$ as a country-specific quadratic function. We estimate the nonlinear function $f(T_{it})$ using four distinct functional

[18] India data can be included as a single, distinct "age-group" in the pooled analysis. In later, age-specific regressions, we are forced to omit India due to its lack of age data.

[19] This is usually the case. However, the EU data is reported at Nomenclature of Territorial Units for Statistics 2^{nd} (NUTS2) level. This falls between first- (e.g., state or province) and second- (e.g., district or county) levels. Japan reports mortality at 1^{st}-level administrative division, but they are of comparable surface area to US ADM2 division.

forms: *(i)* 4^{th}-order and 5^{th}-order polynomials of cumulative daily average temperatures; *(ii)* 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, where the bin edges of k in degrees Celsius are given by the following set: $K = \{-\infty, -17, -12, -7, -2, 3, 8, 13, 18, 23, 28, 33, +\infty\}$; *(iii)* restricted cubic spline where the knots in degrees Celsius are given by $N = \{-12, -7, 0, 10, 18, 23, 28, 33\}$; and finally *(iv)* a 2-part linear spline measuring the effect of heating degree days below 0°C and cooling degree days above 25°C. In this pooled model, we weight observations by age-specific population.[20]

Informed by a specification test described in Appendix C.11, we use the 4^{th}-order polynomial, with $age \times ADM2$ and $country \times age \times year$ fixed effects as our preferred model. 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 plausibly randomly assigned (Deschênes and Greenstone, 2011; Carleton and Hsiang, 2016). The $country \times age \times year$ fixed effects δ_{act} account for any time-varying trends or shocks to age-specific mortality rates which are unrelated to the climate. Standard errors are clustered at the 1^{st} administrative level. Robustness of this model to alternative fixed effects and error structures is shown in Section 4.5, and to alternative functional forms and climate datasets in Appendix C.9. While the binned model is the most flexible functional form, this method is demanding of the data, a constraint that binds particularly in our interaction model designed to capture heterogeneity in the response function (see Section 4.4). In contrast, the polynomial model provides sufficient flexibility while limiting the demands on the data. However, we present results from multiple specifications, and our results are remarkably similar across these functional forms.

In Equation 4.13, the estimated response function $\hat{f}(T_{it})$ is an average treatment effect across each of our three age categories. However, there is substantial evidence from previous research and within our data that individuals of different ages respond heterogeneously to temperature variation (e.g., Deschênes and Greenstone, 2011). Because shifts in demographics in the future may be large, such heterogeneity may have important consequences for projected mortality rates under climate change. Therefore, the model we focus on throughout this paper is a slightly modified version of Equation 4.13, in which a separate response function is identified for each age category:

$$M_{aict} = f^a(T_{it}) + g^c(P_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \tag{4.14}$$

All variables are defined as above, with the only difference being that $f(\cdot)$ is age-dependent. This approach necessitates dropping India from the analysis. However, we perform cross-validation exercises using the India data as a benchmark; this exercise is detailed in Section 4.5.

[20] We constrain population weights to sum to one for each year in the sample. That is, our weight for an observation in region i in year t for age group a is $w^a_{it} = pop^a_{it} / \sum_i \sum_a pop^a_{it}$. 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.

Extrapolation of climate sensitivity over space

A key challenge of conducting a comprehensive analysis of global mortality-temperature relationships is that data availability renders direct estimation of effects impossible in data-sparse locations. In order to overcome this limitation, we develop a simple two-factor interaction model using average temperature and average incomes to explain cross-sectional variation in our estimated mortality-temperature relationship. This approach enables us to use two observable characteristics – average temperature and income – to predict the sensitivity of mortality to temperature in any location around the world. However, it additionally provides estimates of adaptation, as observed in the historical record. That is, the extent to which acclimatization (through higher average temperatures) and income (through loosening the budget constraint influencing investment in adaptive technologies or behaviors) ameliorate temperature's impact on mortality indicate the extent of observable adaptation to climate.

Our first step is to estimate heterogeneity in the estimated response function $f^a(T_{it})$ in Equation 4.14 due to each of these factors. We then use the coefficients from this interaction model to predict dose-response values in each of our 24,378 impact regions, where we have developed downscaled data to identify the values of average income and average temperature within each of these high resolution regions (see Section 4.3). We refer to this two-factor global space of predicted climate sensitivities as the "interpolation surface", because our two covariates create a multidimensional response surface, in which each level of income and climate produces a unique vector of temperature sensitivities for each level of daily temperature exposure.

The factors defining the interpolation surface come directly from the theoretical framework in Section 4.2. First, a higher average temperature incentivizes investment in adaptive behaviors, as the return to any given adaptive mechanism is higher the more frequently the population experiences damaging high-heat days. In Section 4.2, this can be seen directly because the marginal benefit of increasing adaptation (the value of lives saved from lowering temperature sensitivity) is a function of $\boldsymbol{C}$, the climate. We use a parsimonious parameterization of the climate, interacting our nonlinear temperature response function with the location-specific long-run average temperature. Second, higher incomes relax agents' budget constraints, facilitating adaptive behavior. In Section 4.2, this is seen when we define optimal adaptation $\boldsymbol{b}^*$ as a function of exogenous income Y.

Empirically, both factors have been shown to enable adaptation. For example, the sensitivity of mortality to hot days is significantly lower in warmer average climates within the United States, as opposed to cooler regions of the country (Barreca et al., 2015). Similarly, countries with higher average cyclone exposure display substantially depressed sensitivities of income to cyclone wind exposure (Hsiang and Jina, 2014). With respect to income, Hsiang and Narita (2012) show that the marginal damages imposed by cyclones on both human health and economic output fall with income per capita, and Carleton, Hsiang, and Burke (2016) document that the strength of the relationship between temperature and violent conflict declines with GDP per capita.

To create this interpolation surface, we interact these two covariates with our tem-

perature variables, where all covariates are observed at the 1^{st} administrative level (e.g. state):

1. $TMEAN$ is the sample-period average of daily temperature
2. $GDPpc$ is the sample-period average of annual GDP per capita

To see how we estimate this interaction model, note that any of our four nonlinear parameterizations of $f_a(T_{it})$ in Equation 4.13 can be written as:

$$f_a(T_{it}) = \sum_{k \in K} \beta_a^k T_{it}^k$$

where k subscripts indicate a transformation of daily average temperature T_{it}. For example, in the polynomial case, T_{it}^k indicates daily average temperature raised to the power k, while in the binned case, T_{it}^k indicates the total number of days in year t that fall into bin k. In any of our specifications, the coefficients β_a^k define the age-specific response function. Given this, we estimate the following equation for each of our nonlinear specifications:

$$\begin{aligned} M_{aict} = & \sum_{k \in K} \gamma_{0,a}^k T_{it}^k \\ & + \sum_{k \in K} \gamma_{1,a}^k T_{it}^k \times TMEAN_s \\ & + \sum_{k \in K} \gamma_{2,a}^k T_{it}^k \times \log(GDPpc)_s \\ & + g^c(P_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \end{aligned} \tag{4.15}$$

with s referring to ADM1-level (e.g., state or province), and $k \in K$ indicates the term in the nonlinear function of temperature. All other variables are defined as in Equation 4.13. In contrast to the uninteracted models in Equations 4.13 and 4.14, we estimate Equation 4.15 without any regression weights, following the recommendation in Solon, Haider, and Wooldridge (2015) to omit weights when explicitly modeling heterogeneity in treatment effects.

This set of estimates, $\hat{\gamma}$ creates the interpolation surface that we use to predict temperature sensitivities for all impact regions r around the world. Call these predicted temperature sensitivities $\hat{\hat{\beta}}_{r,a}^k$, which are specific to each temperature term k in the regression and each age group a in each impact region r. To identify these $\hat{\hat{\beta}}$s, however, we first must assemble data for each factor in Equation 4.15 for every region r. Because such data are not readily available for GDP per capita (to our knowledge, Gennaioli et al. (2014) is the most comprehensive income dataset available for subnational administrative units, and it provides neither global coverage nor sufficiently high resolution to match our impact regions r), we use the downscaling procedure described in Section 4.3 to predict impact region level income using national income statistics in combination

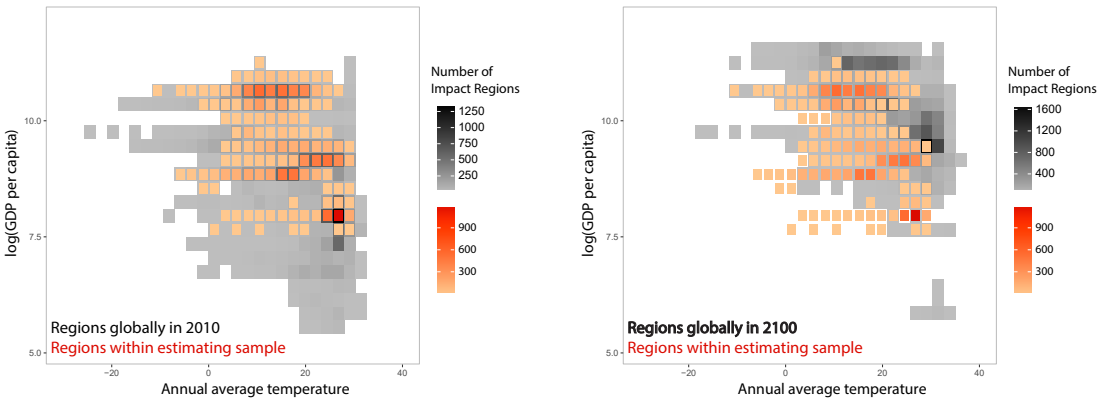

Figure 4.3: Joint coverage of income and long-run average temperature

Coverage in our sample evaluated over impact regions (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 2010. 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 (SSP 3). In both panels, the in-sample frequency in red-orange indicates coverage for impact regions within our data sample in 2010.

with night lights imagery from the NOAA's Defense Meteorological Satellite Program (DSMP). To generate $TMEAN$, we use the same temperature data as was compiled for the regressions in Equations 4.14 and 4.15, calculating average daily temperature in each impact region. Figure 4.3 shows the coverage of these covariate data for income and long run average temperature in 2010 for the impact regions contained within the countries covered by our mortality data (and thus used to estimate Equations 4.14 and 4.15). Coverage in-sample (in red-orange) is contrasted in Figure 4.3 with the global set of impact regions (in grey-black). We note that temperatures in the global sample are generally well-covered by our data, but we fail to capture many of the poorer regions in the world, due to the scarcity of mortality data in many poorer countries (panel **A**). However, panel **B** in Figure 4.3 shows that our sample is fairly representative of global temperatures and incomes in 2100, after locations have warmed and become more wealthy.

With these impact region level data in hand, we predict the response curve for each age group for each region, with values $\hat{\hat{\beta}}^k_{r,a}$ determined as:

$$\hat{\hat{\beta}}^k_{r,a} = \gamma^k_{0,a} + \gamma^k_{1,a} TMEAN_r + \gamma^k_{2,a} \log(GDPpc)_r \tag{4.16}$$

where values of the covariates are estimated for a baseline period of 2001 to 2015. This results in a spatially heterogeneous yet globally comprehensive set of response functions, each of which is specific to an impact region and age group. Note that this approach implies that every impact region is assigned a *predicted* response function generated from the interpolation surface, even if it is within a country with an estimated response (i.e., where data exist). Once this interpolation step is complete, we have a unique

set of age-specific response functions for each impact region, completely covering the earth's surface.

Projected climate impacts

Our approach to generating projected climate impacts on mortality satisfies two essential problems that have been acknowledged in the literature. First, we obtain globally representative coverage of dose-response relationships (Section 4.4), and second, we allow these relationships to change through time as a function of projected climate and income. As described in Section 4.4, in the first step we extrapolate over *space*. In this second step, we extrapolate over *time*, such that regions' dose-response functions exhibit adaptation via changes that unfold over the coming century in the two factors modeled in Equation 4.15.

To extrapolate over time, we allow response functions to change smoothly from the baseline interpolated response shown in Section 4.4, as the average climate and as GDP per capita evolve into the future. We assign a specific time period for averaging each of the covariates: in the case of temperature we specify a 15-year lagged kernel, while in the case of income we empirically derive a 13-year lagged kernel from our data.[21] Details of the income adaptation kernel derivation are given in Appendix C.6. Relying on the parameters $\boldsymbol{\gamma}$ from Equation 4.15, an update process occurs that modifies the response curve in region r and in year t as follows:

1. A 15-year moving average of temperatures for region r is updated using the temperature from year t to generate a new level of $TMEAN_{rt}$.

2. A 13-year moving average of income per capita in region r is calculated from the OECD Env-Growth Shared Socioeconomics Pathways (SSP) model using our national-to-regional income downscaling procedure to generate a new value of $\log(GDPpc)_{rt}$.[22]

3. The response curve is calculated according to Equation 4.16 with the updated values of $TMEAN_{rt}$ and $\log(GDPpc)_{rt}$ shown in the steps above.

We apply these spatially and temporally heterogeneous response functions to projected changes in the climate in order to generate predicted impacts of climate change on mortality through the 21$^{\text{st}}$ century. To do so, we first calculate the nonlinear transformations of daily average temperature that are used in the function $f(T_{it})$ under both the RCP4.5 and RCP8.5 emissions scenarios for the 21 GCM projections and 7

[21] We can derive a duration over which updating occurs in the case of incomes due to the larger time series variation of incomes in the data. For temperature, the historical trends have so far been small and this makes derivation of a comparable duration difficult.

[22] GDP per capita is provided in the SSP projections only at 5-year increments. The values are held constant between SSP 5-yearly updates of values. The temporal pattern of income growth from the SSPs is maintained, but we scale the value by $\frac{avgincome_r[t]}{avgincome_r[0]}$, the ratio of a future income per capita to the baseline value in 2010.

model surrogates in our data (as described in Section 4.3). We sample from individual GCMs in order to account for uncertainties in the climate system as projected into the 21st century. A second source of uncertainty in our projected impacts arises from our econometric estimates of response functions. There are two contributions to this – uncertainty in the estimates of β^k and uncertainty in the estimate of each of the γs. In order to account for many of these sources of uncertainty, we draw multiple values from across the full distribution of these terms, as well as base impact projections where we hold all estimated terms at their median values. The resulting Monte Carlo calculation is highly computationally intensive and fully incorporates uncertainty from climate and econometric sources.

In generating these projected impacts, we impose two simple assumptions, guided by basic theory as well as physiological literature.

Assumption #1: Full adaptation is defined as a flat line. We define the fully adapted state as one in which variation in temperature has no effect on mortality. To implement this assumption in projections of future impacts, we first identify a region of physiologically optimal temperatures at which the fully adapted state may, in theory, occur. Drawing on extensive research across epidemiology and medicine (e.g., Curriero et al., 2002; Guo et al., 2014), and ergonomics (e.g., Seppanen, Fisk, and Lei, 2006b; Hancock, Ross, and Szalma, 2007), we let this range of possible minimum mortality risk cover the temperatures 10°C to 25°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 (e.g. Guo et al. (2014) demonstrate that mortality risk is smallest around the 75th percentile of local temperatures in 12 different countries), it is important to note that we allow these minimum mortality risk temperatures to be spatially heterogeneous (although fixed over time). In calculating projected impacts, we then allow adaptation to occur (i.e. allow the response function to flatten) until the minimum mortality risk level, as defined in 2010, is reached. We note that while this approach allows for adaptive behaviors captured by future changes in our two factors, income and average climate, our imposition of this assumption precludes any changes in technology or adaptation investments which may alter the minimum mortality risk level itself.

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 when estimating Equation 4.15, we do not allow any income gains through time to raise the marginal effect of temperature on mortality. We implement this restriction by bounding the marginal effect of income on mortality; that is, when $\sum_{k \in K} \gamma_2^k T_{rt}^k > 0$ for some region r in period t, we set $\sum_{k \in K} \gamma_2^k T_{rt}^k = 0$. Note that this condition will only be binding if the marginal effect of income estimated in Equation 4.15 is positive for some nonempty set of temperatures. Further note that we impose this assumption first, before ensuring that adaptation does not lead to impacts that

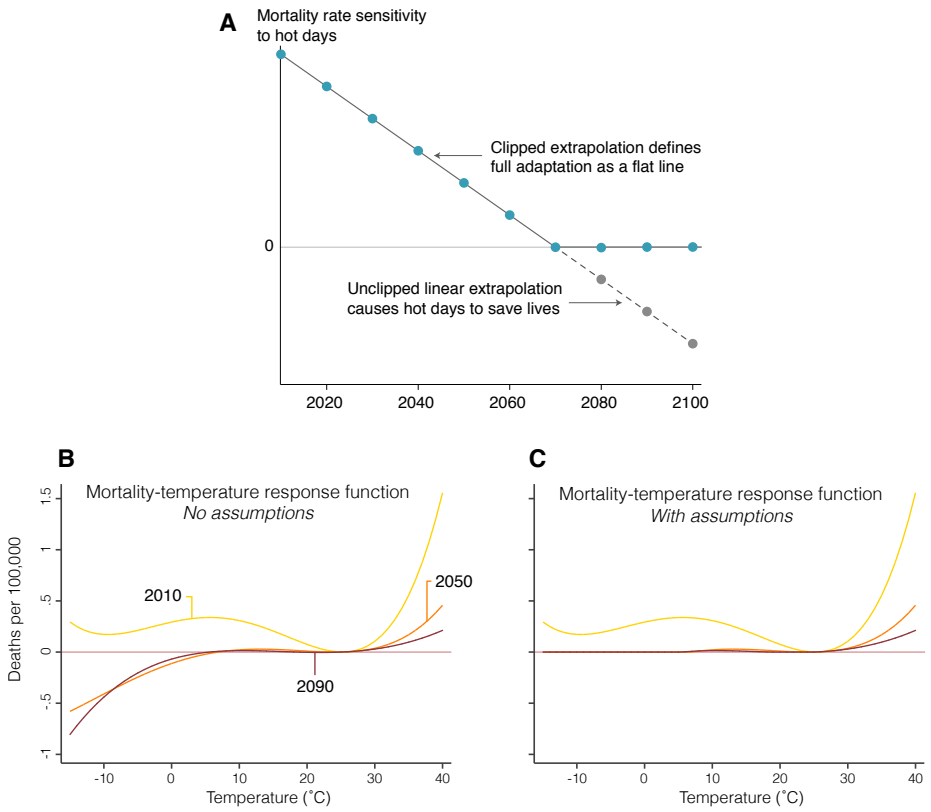

Figure 4.4: Full adaptation defined as a flat line

Two assumptions are imposed in climate projections to ensure that full adaptation is defined as a flat-line response function. 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 the that full adaptation is realized when hot days impose zero additional mortality risk. Panels **B** and **C** represent an empirical example of how the imposition of these restrictions can change the shape of the response function, for a given impact region.

fall below the minimum mortality risk level described under assumption #1. A visual example of the influence of these two assumptions can be seen for one example impact region in Figure 4.4.

Under these two assumptions, we estimate projected impacts separately for each impact region, and then aggregate these high resolution effects to state, country, and global levels. This aggregation is done using population weighting, as follows:

$$W_{ct} = \frac{\sum_{i \in R(c)} W_{it} P_{it}}{\sum P_{it}} \tag{4.17}$$

Since all regions in a country grow at the same rate, the weighting stays constant across

time. However, different countries grow at different rates, and the global estimate is computed the same way, but using all regions, r, as

$$I_t = \sum_k \hat{\hat{\beta}}_r^k[t] T_t^k \tag{4.18}$$

where T_t^k is the number of days in bin k in year t. After computing impacts for each high-resolution region, we aggregate them by computing $\sum_{r \in R(j)} \frac{P_r[t] y_r[t]}{P_r[t]}$, the weighted average. For $t = 2015$, $\hat{\hat{\beta}}_r^k$ with baseline covariate data values is used for each impact region.

4.5 Results

Our results proceed in a series of sequential steps. We first establish that mortality has a U-shaped and significant relationship with temperatures in our multi-country pooled sample, consistent with previous studies with a narrower geographic scope. We then demonstrate that there is substantial age-heterogeneity in this relationship, which motivates separating the data into three broad age groups for the remainder of the analysis. We show that this result is robust to a variety of specification choices, population-weighting schemes, estimation strategies that limit the potential bias introduced by variability in data collection capacity, and varying lag lengths of exposure. Heterogeneity due to differences in incomes and average climates is then demonstrated. These results are used to extrapolate responses to the parts of the world where we have now data, allowing us to create the first global average treatment effect of temperature on mortality. Finally, we project these responses into the future to calculate the full mortality cost of climate change, accounting for adaptation benefits costs.

Multi-country mortality dose-response function

Pooling subnational mortality records across 41 countries, we estimate Equation 4.13, showing results for the dose-response function obtained with a 4^{th}-order polynomial in daily average temperature. Table 4.3 displays our main result, showing the marginal effects at various temperatures. These estimates can be interpreted as the change in the number of deaths per 100,000 for a day at each temperature, compared to a moderate day at 20°C. 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 and state-level linear trends in column (3). The U-shaped response common in prior literature is evident across all specifications. Choosing column (2) as our preferred specification, we find, for example, that a day at 35°C leads to an increase in the mortality rate of approximately 0.5 extra deaths per 100,000, relative to a day at 20°C. A day at -10°C similarly increases the mortality rate, relative to a moderate day, by 0.3 deaths per 100,000.

	Mortality rate				
	(1)	(2)	(3)	(4)	(5)
35°	0.434***	0.454***	0.156	0.647**	0.471***
	(0.123)	(0.151)	(0.108)	(0.302)	(0.147)
30°	0.313***	0.308***	0.113*	0.300**	0.319***
	(0.066)	(0.072)	(0.063)	(0.121)	(0.069)
25°	0.150***	0.142***	0.053	0.097***	0.148***
	(0.036)	(0.033)	(0.038)	(0.036)	(0.031)
0°	0.048	0.042	0.066	0.143**	-0.012
	(0.109)	(0.112)	(0.090)	(0.069)	(0.099)
-5°	0.214	0.184	0.146	0.233**	0.106
	(0.134)	(0.139)	(0.091)	(0.101)	(0.121)
-10°	0.382**	0.321*	0.225**	0.335**	0.218
	(0.164)	(0.170)	(0.103)	(0.138)	(0.145)
Adj R-squared	.987	.994	.994	.999	.994
Observations	787712	787329	787329	787329	787329
Adm2-Age FE	YES	YES	YES	YES	YES
Cntry-Yr FE	YES	-	-	-	-
Cntry-Yr-Age FE	-	YES	YES	YES	YES
Adm1-Yr-Age Lin TR	-	-	YES	-	-
Pseudo Precision-Weighted	-	-	-	YES	-
13-month exposure	-	-	-	-	YES

Standard errors clustered at the ADM1 level. All regressions are population-weighted
*** p<0.01, ** p<0.05, * p<0.1

Table 4.3: Multi-country temperature-mortality response function

This table shows coefficient estimates for a temperature-mortality response function estimated using pooled subnational data across 41 countries and 39% of the global population. 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. Point estimates indicate the marginal effect of increasing daily average temperature by 1°C, evaluated at each temperature value shown.

Temperature functional form selection To ensure robustness of our result across specifications, we estimate Equation 4.13 for each of the functional forms of temperature described in the previous section. Results for these are displayed in Figure C.8. However, some of the functional forms have better properties for projection or extrapolation, which are essential components of our exercise of projecting responses into the future at a global scale. In particular, we prefer a parametric functional form over a non-parametric binned form as is it less demanding of the data. We note, however, that the binned functional form allows us to estimate the full nonlinearity of the response with no restrictions. Therefore, we perform a test of model selection that takes the binned response as the "true" response within the data and evaluate each other specification against that benchmark with a set of model-fit criteria. Details of these results

are given in Appendix C.11. We find that the 4^{th}-order polynomial performs well, and so for clarity of exposition we display only those results for the remainder of the paper unless stated otherwise.

Age group heterogeneity Because of substantial physiological differences, it is likely that age cohorts response differently to a given temperature (Deschênes and Moretti, 2009). This heterogeneity is important to account for, as there exist large demographic disparities across countries today, and because there are likely to be considerable demographic transitions in the future. We test for heterogeneity across age groups by estimating Equation 4.14, which provides separate coefficients for each of our three main age categories. Results are shown in Table 4.4. We see that there is substantial heterogeneity across age groups within our multi-country sample. Again, selecting column (2) as our preferred specification, we see that people over the age of 65 experience approximately 4.4 extra deaths per 100,000 for a day at 35°C compared to a day at 20°C. The results are plotted in Figure 4.5, with the all-age 95% confidence interval (i.e., the results plotted from the estimates shown in column (2) of Table 4.3) shown for reference. We see that, while all age groups have positive mortality responses to hot days, the magnitude of the effect for the older population (ages > 65) is substantially larger than for the younger cohorts. This group does, however, represent a relatively small proportion of the global population today, limiting its similarity to the average treatment effect in Table 4.3. This demographic is also more severely affected by cold days than are the younger age cohorts; we find that people over the age of 65 experience 2.8 extra deaths per 100,000 for a day at -5°C compared to a day at 20°C, while other age categories display a flat response, such that much of the cold-related mortality in the overall response is driven by the older age group. This implies that many of the benefits of decreased mortality due to fewer cold days under climate change will accrue to the elderly and not to children or working age individuals, and is consistent with prior evidence on age heterogeneity in the mortality-temperature relationship in the United States (Deschênes and Moretti, 2009).

Alternate climate data The main effect for our all-age response function is plotted in Figure C.8 for the GMFD climate data (top) and the BEST climate data (bottom). These data are drawn from independent sources, as described in Section 4.3, and we see broad similarity for the responses across all functional forms and both sets of data. The figure illustrates the importance of a nonlinear treatment of temperature impacts, while also showing broad agreement across various functional forms for both climate datasets.

Pseudo precision-weighting Due to the large variety of sources for our mortality data, with some sources being drawn from countries which may have less capacity for data collection than others in the sample, we perform a two-stage estimation that attempts to down-weight observations that may be of lower quality. To do this, after estimating the model once, we use the ADM1 average value of residuals to inverse-

	Mortality rate				
	(1)	(2)	(3)	(4)	(5)
Panel A: <5 years of age					
35°	2.415***	0.084	0.134	0.292	0.040
	(0.522)	(0.262)	(0.184)	(0.258)	(0.256)
30°	1.510***	-0.037	0.068	0.099	-0.040
	(0.235)	(0.107)	(0.071)	(0.114)	(0.106)
0°	-2.543***	-0.075	-0.113	-0.032	-0.112
	(0.304)	(0.114)	(0.102)	(0.056)	(0.111)
-5°	-2.683***	-0.198*	-0.142	0.004	-0.221*
	(0.377)	(0.116)	(0.100)	(0.088)	(0.115)
Panel B: 5 - 64 years of age					
35°	4.970***	0.056	-0.003	0.087	0.073
	(0.671)	(0.106)	(0.073)	(0.214)	(0.106)
30°	2.760***	0.072	0.028	0.023	0.081
	(0.261)	(0.061)	(0.037)	(0.094)	(0.061)
0°	-4.150***	-0.134**	-0.096*	-0.005	-0.138**
	(0.295)	(0.063)	(0.050)	(0.040)	(0.053)
-5°	-4.732***	-0.130	-0.092*	-0.015	-0.133**
	(0.369)	(0.079)	(0.050)	(0.060)	(0.066)
Panel C: ≥65 years of age					
35°	-4.417**	4.361***	1.619	4.889**	4.439***
	(1.787)	(1.635)	(1.229)	(2.181)	(1.584)
30°	-2.348***	2.552***	0.838	2.414***	2.605***
	(0.798)	(0.700)	(0.569)	(0.833)	(0.670)
0°	8.120***	1.515***	1.445***	1.242***	1.072**
	(0.694)	(0.578)	(0.437)	(0.422)	(0.543)
-5°	10.145***	2.807***	2.133***	2.083***	2.155***
	(0.823)	(0.770)	(0.456)	(0.577)	(0.699)
R-squared	.988	.993	.993	.999	.993
Observations	787712	787329	787329	787329	787302
Adm2-Age FE	YES	YES	YES	YES	YES
Cntry-Yr FE	YES	-	-	-	-
Cntry-Yr-Age FE	-	YES	YES	YES	YES
Adm1-Yr-Age Lin TR	-	-	YES	-	-
Pseudo Precision-weighting	-	-	-	YES	-
13-month exposure	-	-	-	-	YES

Standard errors clustered at the ADM1 level. All regressions are population-weighted
*** p<0.01, ** p<0.05, * p<0.1

Table 4.4: Temperature-mortality response function with demographic heterogeneity

Regression estimates are from a fourth-order polynomial in daily average temperature and are estimated using GMFD weather data with a multi-country pooled sample that was winsorized at the 1% level. Point estimates indicate the marginal effect of increasing daily average temperature by 1°C, evaluated at each temperature value shown.

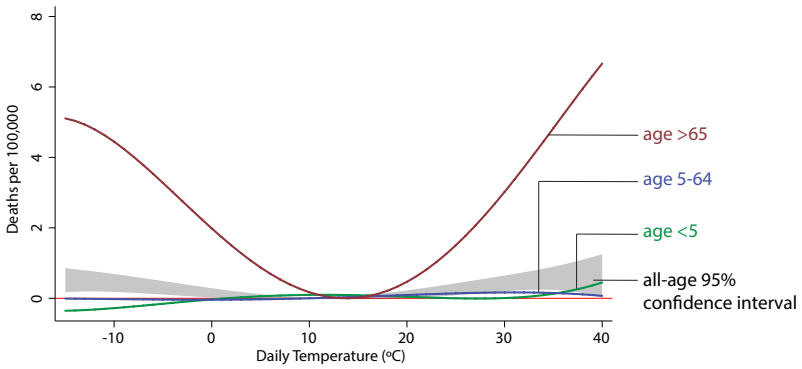

Figure 4.5: Temperature-mortality response function with demographic heterogeneity

Temperature-mortality response function estimated for population below 5 years of age (green), between 5 and 64 years of age (blue), and 65+ 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.

weight the regression in a second stage. All ADM2-age-year observations within a given ADM1-age category are given the same weight in the second stage, where ADM1 locations with lower residual variance are given higher weight. This has the effect of down-weighting those observations that show larger discrepancies from the main estimated effect. The results for this procedure are displayed in column (4) of Tables 4.3 and 4.4. The differences between columns (2) and (4) are small, from which we conclude that the variability in data quality, after winsorizing, is not a major concern for our analysis.

Lagged exposure Numerous previous papers have demonstrated that temperatures have a lagged effect on health and mortality (e.g., Deschênes and Moretti, 2009; Barreca et al., 2016; Guo et al., 2014). In particular, cold temperatures can influence mortality rates for up to 30 days (Deschênes and Moretti, 2009). Our main analysis is conducted at the annual level, so that lagged effects within individual months are accounted for in net annual mortality totals. However, to allow for the fact that January mortality may result from up to a four-week lag of temperature exposure, we define a "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 $t-1$. These results are shown in column (5) of Tables 4.3 and 4.4. Changes due to this specification are minimal, but are mostly evident in the old-age response to cold temperatures, implying that there are longer lagged effects of cold temperatures on winter mortality for the elderly.

Sub-national heterogeneity of dose-response function

We estimate the interaction model in Equation 4.15 to achieve two goals. First, we seek to uncover empirical estimates of the benefits of adaptation by comparing sensitivity of mortality to temperature across heterogeneous locations within our multi-country sample. Second, we aim to create an interpolation surface based on this subnational heterogeneity in sensitivity that will allow us to extrapolate response functions to new regions of the globe, and to future time periods. In this model, we interact our weather variables with 1st-administrative-level covariates of average climate and average income.[23] The results for the γ-terms for this model are shown in Table 4.5 for each of the three age groups of interest. Each coefficient represents the marginal effect of increasing the relevant covariate (e.g. $\log(GDPpc)$) by one unit on the temperature-sensitivity of mortality, evaluated at the daily temperature shown. For example, we see that higher incomes lower the sensitivity of infant mortality to both cold temperatures (coefficient of -0.5 on a -5°C day), as well as to hot temperatures (coefficient of -0.2 on a 35°C day).

Figure 4.6 helps to visualize the heterogeneity we uncover within our sample, taking the oldest age category as an example.[24] Each subpanel in Figure 4.6 shows a dose-response function for the age 65+ mortality rate, evaluated at a particular level of income and average climate. To show heterogeneity across all subsets of income-climate space within our sample, we plot predicted response functions for each tercile of income and average temperature within our data. Average incomes are increasing in the y-direction and average temperatures are increasing in the x-direction. Results in Figure 4.6 demonstrate consistency with 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. Indeed, across all age categories, moving from the coldest to the hottest tercile in our sample saves on average 0.25 deaths per 100,000 at 35°C. Similarly, a loosening of the budget constraint, as proxied by increasing GDP per capita, should enable individuals to invest further in adaptation. We find that the effect of moving from the poorest to the richest tercile in our sample saves on average 0.54 deaths per 100,000 at 35°C across all age categories.

By exploiting the full nonlinearity of the temperature distribution, we are able to see that this "adaptation" to marginal damages occurs primarily in the extremes of the temperature distribution, with little or no adaptation necessary in the central part of the distribution. Income is protective on the hottest days for the middle and older age category, and is additionally beneficial on the coldest days for the youngest age category. As evidenced in Figures 4.6 and C.6, the richest regions exhibit substantial flattening of the response curve under warmer average climates. This is particularly

[23] As mentioned in Section 4.4, we define "average climate" as the average temperature, and average income as the average of GDP per capita, both calculated over the entirety of the sample period for each ADM1 unit.

[24] An all-age version of Figure 4.6 is shown in Appendix C.7.

	Age < 5		Age 5-65		Age 65+	
	lgdppc	Tmean	lgdppc	Tmean	lgdppc	Tmean
35°	-0.204	0.005	-0.179	-0.018	-3.556	-0.282
	(0.297)	(0.052)	(0.142)	(0.017)	(2.408)	(0.322)
30°	0.071	0.010	-0.014	-0.010	0.211	-0.147
	(0.207)	(0.030)	(0.057)	(0.008)	(0.855)	(0.135)
0°	-0.510	0.045*	0.078	-0.031	2.929	-0.751***
	(0.394)	(0.026)	(0.146)	(0.019)	(1.959)	(0.145)
-5°	-0.542	0.051*	0.217	-0.045**	6.835**	-0.974***
	(0.467)	(0.029)	(0.217)	(0.022)	(2.947)	(0.180)
Observations	787,329					
R-squared	0.936					
Admin2-Age FE	YES					
Country-Year-Age FE	YES					

Standard errors clustered at the state level
*** p<0.01, ** p<0.05, * p<0.1

Table 4.5: Marginal effect of covariates on temperature sensitivity of mortality

Coefficients represent the marginal effect of increasing each covariate by one unit on the temperature sensitivity of mortality rates, evaluated at each of the shown daily average temperatures. Regression is a fourth-order polynomial in daily average temperature, 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. Each temperature variable is interacted with each covariate.

the case on hot days, and is consistent with evidence from the United States – Barreca et al. (2016) show a steep decline in mortality sensitivity to heat when moving from the hottest to the coolest average climates within the United States.

Global extrapolation

To derive globally representative climate change damages, we must uncover the relationship between temperature and mortality for the entire globe, not just those locations where data are available. Figure 4.7 demonstrates our extrapolation of mortality-temperature response functions to the entire globe. To execute this extrapolation, we use the estimates of $\boldsymbol{\gamma}$ from Equation 4.15, combined with impact region level incomes and average temperatures, to predict dose response functions at each impact region, r; we denote these predictions $\hat{\beta_{r,a}}$. In panel **A**, these $\hat{\beta_{r,a}}$ are plotted for each impact region for 2010 values of income and climate for the oldest age category and for the impact regions that fall within the countries in our mortality dataset ("in-sample"). Despite a shared overall shape, panel **A** shows substantial heterogeneity across regions

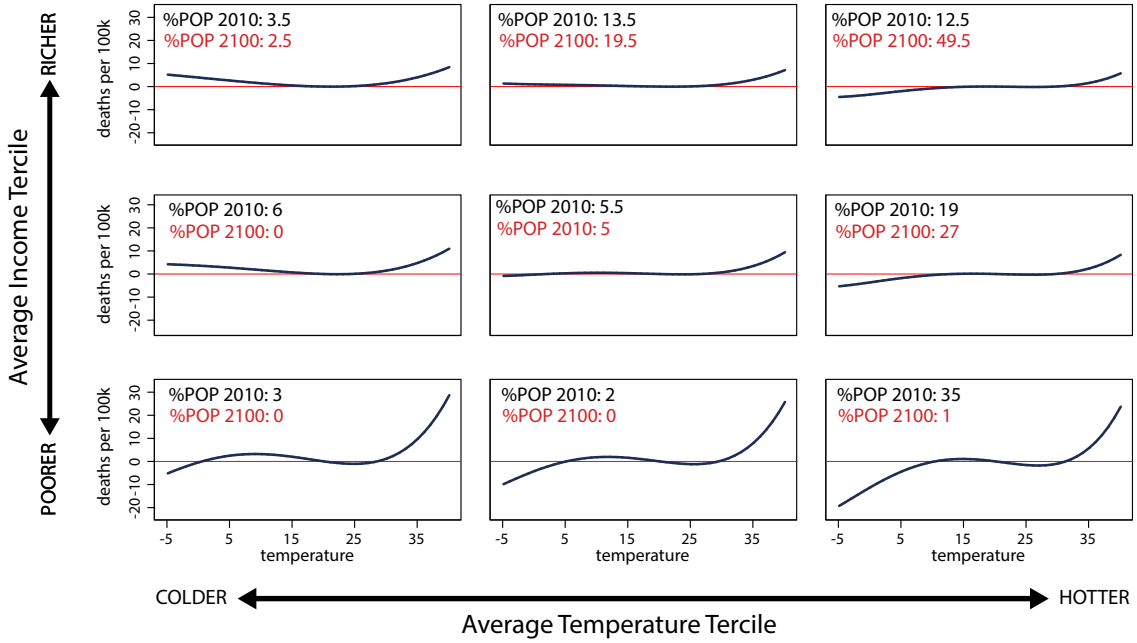

Figure 4.6: Heterogeneity in the mortality-temperature relationship

Each panel represents a predicted response function for the 65+ age category 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 poor, cold regions of our sample, while those in the upper right apply to the wealthy, 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. 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. The all-age response functions shown here are population-weighted averages of estimated coefficients in each of three age categories.

in this temperature response. On average, the response function across these impact regions (shown in black) lines up roughly with the estimated in-sample average treatment effect estimated as described in 4.5 (shown in red). In panel **B**, we show an analogous figure for the youngest age category; this result demonstrates our model's ability to predict the substantial within-location differences in mortality risk across age groups. Geographic heterogeneity within our sample is shown for hot days in the maps in panels **C** and **D**. On each map, the colors indicate the marginal effect of a day at 35°C day, for ages 65+ (panel **C**) and for ages < 5 (panel **D**).

In panels **E**–**H** of Figure 4.7, we show analogous figures, but now extrapolated to the entire globe (in red in panels **E** and **F**, we show the in-sample response function for reference). The predicted responses at the global scale imply that a 35°C day increases the annual global mortality rate by 2.94 deaths per 100,000 relative to day at 20°C, which is considerably larger than the estimated effect within our sample. This

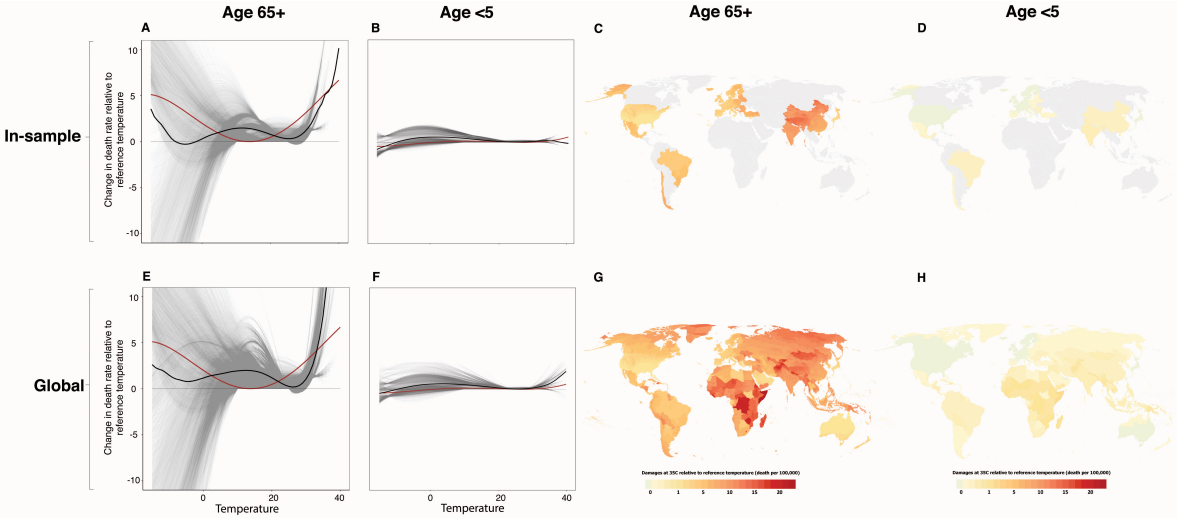

Figure 4.7: Using a model of adaptation to predict response functions out of sample

In panels **A, B, E** and **F**, grey lines are predicted response functions for each impact region, and solid black lines are the unweighted average of the grey lines. The red line in these panels, for reference, is the "true" average response function – the estimated average treatment effect from a regression model without heterogeneity in temperature sensitivity (column (2) in Table 4.4). Panels **C, D, G** and **H** show maps of 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 our data sample, and the bottom row shows extrapolation to all impact regions globally.

difference in sensitivity is largely due to the large protective benefits of income and the lack of low income regions within our estimating sample. In panels **G** and **H**, the estimated geographic heterogeneity in mortality sensitivity to heat gobally reveals the influence of both income and average climate – while hot average climates are generally better adapted, poor locations that are also relatively warm remain highly susceptible to mortality effects of heat, due to the importance of income in facilitating adaptation.

Figure 4.7 makes clear the substantial degree of out of sample extrapolation that is required to take our multi-country model of heterogeneity and apply it to a global sample. To ensure that we are accurately representing response functions in new locations, we perform cross-validation tests (presented in Appendix C.8) that involve sequentially leaving out one country in estimation and predicting a response function in the omitted country. We then compare our predicted response function for the omitted country to the response estimated on that country's data alone. In Figure 4.8, we show the result of this cross-validation exercise for India. India represents the poorest and hottest region for which we have mortality records, and provides an important check on our extrapolation performance, as the global sample is substantially poorer and hotter than our data (see Figure 4.3). Figure 4.8 compares our predicted responses for all impact regions in India (in grey) to the mortality-temperature response estimated using India's

data (in red).[25] 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. Similarly reassuring results arise for cross-validation tests in other countries, as shown in Appendix C.8.

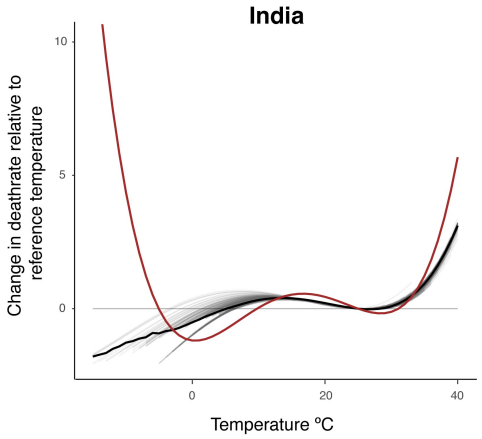

Figure 4.8: Cross-validation of adaptation model: India

Cross-validation in India of the interaction model of temperature sensitivity demonstrates that the model performs well in the poorest and hottest locations of our sample. Grey lines are predicted response functions for each impact region in India, estimated from the interaction model using a 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 data from India. The congruence between red and black lines shows that our interaction model generates accurate predicted response functions in the poorest and hottest regions of the world, the subset of covariate space for which we have the least representation.

Projection of damages and adaptation into the future

In Section 4.5, we demonstrated that our model of heterogeneity allows us to extrapolate temperature-mortality relationships to regions of the world without mortality data today. However, to calculate the total mortality costs of climate change, recall from Section 4.2 that we must allow these response functions to change through time, reflecting the benefits of adaptation. Here, we use our model of heterogeneity to extrapolate impact region response functions over time, completing a key step towards a valuation of the total mortality costs of climate change. To do so, we use downscaled projections of income and climate to predict impact region level response functions for each age group and year, yielding $\hat{\hat{\beta}}_{rt,a}$. We plot the spatial distribution of marginal damages of a 35°C day over the 21^{st} century in Figure 4.9. The maps in Figure 4.9 reveal that in most

[25] We generate an all-age response function from our age-specific interaction model by population-weighting the coefficients from each age category with age-specific population values.

regions of the world, we see a clear downward trend in the sensitivity of mortality rates to high temperatures, as locations get both richer and hotter as the century unfolds. Calculating a global dose-response function for 2100 results in an additional 35°C day relative to day at 20°C increasing the annual mortality rate by 0.89 deaths per 100,000. When compared to the 2.94 deaths per 100,000 calculated in 2010, this represents large benefits of adaptation, i.e. reoptimization of the vector $\boldsymbol{b}$ from Section 4.2. We find that increasing incomes account for 88% of the decline in marginal damages.

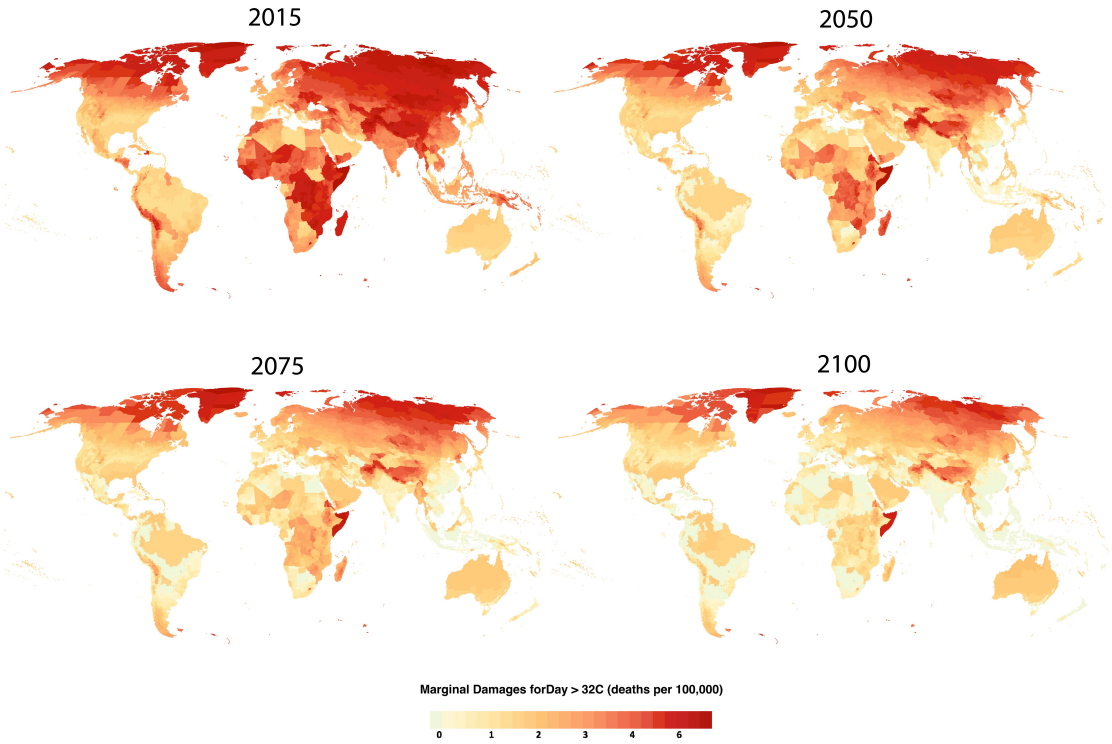

Figure 4.9: Mortality sensitivity to high temperatures decays through time

Rising incomes and gradually warming climates lower mortality sensitivity to high temperatures through the21st century. Colors indicate the increase in the mortality rate for the oldest age category (65+) due to a hot day (32°C), relative to a moderate day (location-specific minimum mortality temperature). Both income and average climate affect temperature sensitivity, such that by end of century, substantial adaptation is predicted.

The response functions illustrated in Figure 4.9 reflect the protective effect of projected economic growth, as well as adaptation to the warming average temperature distributions. To calculate the costs of climate change, we have to be careful to isolate the role of income as separate from the role of climate change. In particular, our interpolation surface allows us to calculate three types of response functions for each impact region at each point in time: *(i)* a constant β evaluated at 2010 income and average climate levels; *(ii)* a time-varying β^{inc}, which accounts for future economic growth; and *(iii)* a time-varying $\beta^{inc,clim}$, which accounts for both economic growth

and future shifts in average climate. We isolate the impact of climate change on future temperature-driven mortality as follows:

$$\hat{M}^{CC} - \hat{M}^{NoCC} = \beta^{Inc,Clim} T_1 - \beta^{Inc} T_0 \tag{4.19}$$

The first term in Equation 4.19, $\beta^{Inc,Clim} T_1$, represents the predicted mortality rate under a future climate in a population that adapts both in response to income and a warming climate. The second term, $\beta^{Inc} T_0$, represents a counterfactual predicted mortality rate that would be realized under current temperature exposure in a population that adapts in response to rising incomes over the coming century. This counterfactual includes the prediction, for example, that air conditioning will become much more prevalent in a country like India as the economy grows, regardless of whether climate change unfolds or not. By subtracting off this counterfactual, our predicted mortality levels isolate the role of climate change from the role of economic growth in determining future mortality rates, although mortality sensitivity to temperature is still moderated by the protective effects of rising incomes.

We generate impact-region specific predictions of mortality damages of a warming climate for all days between 2015 and 2100 under a range of climate and socioeconomic scenarios. Following the approach outlined in Section 4.2 and Appendix C.1, we also generate associated measures of adaptation costs for each location at each point in time. To visualize the magnitude of these effects at the global scale, we aggregate total impacts up to the global level to calculate the full mortality damages of climate change. Figure B.2 shows our median predictions across all 28 climate models of the global increase in the mortality rate (deaths per 100,000) due to climate change, over time. This time-series is decomposed into different drivers of adaptation benefits and costs, where adaptation is reflected in the decrease in the sensitivity to temperature, and is governed by our changing $\hat{\hat{\beta}}_{rt,a}$ values.

The "No Adaptation" scenario in Figure B.2 uses each impact region's estimated response at their 2010 level of income and average climate, not allowing temperature sensitivity to change with future income or temperature. As discussed in Section 4.2, this is a benchmark model often employed in previous work, and accounts only for the direct effects of climate change. The "Future Income" scenario allows the response function to change with future incomes. Thus, the difference between these two scenarios captures the temperature-driven mortality benefits of future economic growth. The "Future Income and Climate Adaptation" scenario additionally allows the response to adjust to future temperature distributions, thereby accounting for additional protective effects of people adjusting to shifts in average climate. Since this anticipated adaptation will require agents to incur a cost, which we bound using our revealed preference methodology detailed in Section 4.2, the total mortality cost of climate change is given by the black line. This estimate adds back in the costs $A(\boldsymbol{b}(\boldsymbol{C}_1)) - A(\boldsymbol{b}(\boldsymbol{C}_0))$, measured in lives lost, therefore fully accounting for adaptation as well as its costs.

Without taking account of adaptation, we estimate that the mortality cost of climate change will be 125 deaths per 100,000 by 2100 under an RCP8.5 emissions scenario. Increased income reduces this further to approximately 48 deaths per 100,000, and

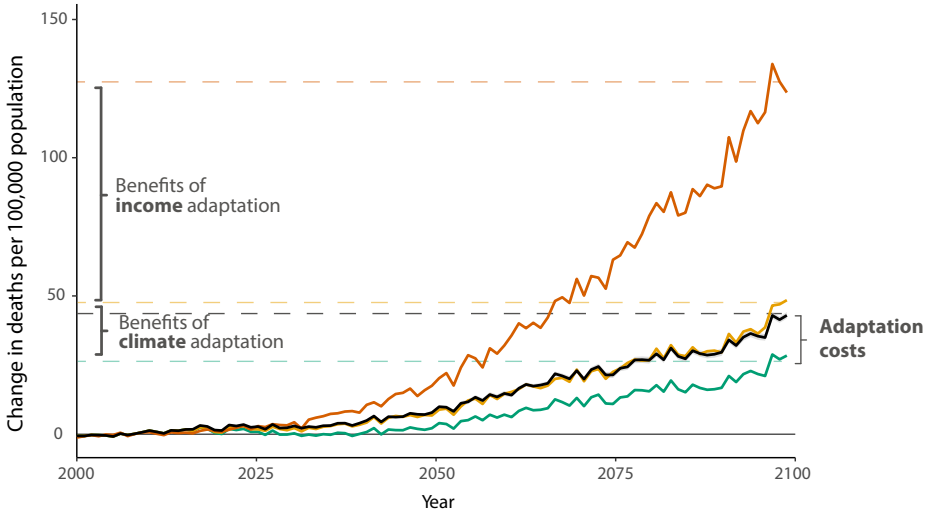

Figure 4.10: Projected mortality costs of climate change

Time series of the total mortality costs of climate change (black line), shown 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, represented by a median estimate across 28 climate models. Each colored line represents a partial mortality cost, while the black line shows the total mortality costs of climate change, accounting for both adaptation costs and benefits. Orange: no adaptation. Yellow: adaptation via income growth only. Green: adaptation via income growth and in response to a warming climate. Black: adaptation via income growth and in response to a warming climate, plus the average estimate of costs incurred to achieve adaptation.

climate adaptation reduces this to approximately 28 deaths per 100,000. While these latter two values indicate substantial adaptation benefits, it is important to note that even these smaller values are economically meaningful: the current mortality rate from automobile accidents in the United States is approximately 12 per 100,000. Moreover, we find that the cost of climate adaptation, using the approach outlined in Section 4.2, is equivalent to an approximate increase of 16 deaths per 100,000, resulting in an estimate of the total mortality costs of climate change of 44 deaths per 100,000 by the end of the century under an RCP8.5 emissions scenario.

These projected damages depend heavily on the mitigation scenario. Figure 4.11 shows a comparison between our projections for the total costs of climate change under RCP8.5 (relatively high emissions) and RCP4.5 (relatively low emissions). By end of century, the costs fall from 44 deaths per 100,000 under RCP8.5 to approximately 10 per 100,000 under RCP4.5, demonstrating remarkable gains from mitigation.

Finally, the aggregate global time series masks considerable heterogeneity in where mortality impacts will be experienced. A key benefit of our approach is that we can derive location-specific projections at high resolution globally. Figure 4.12 shows the

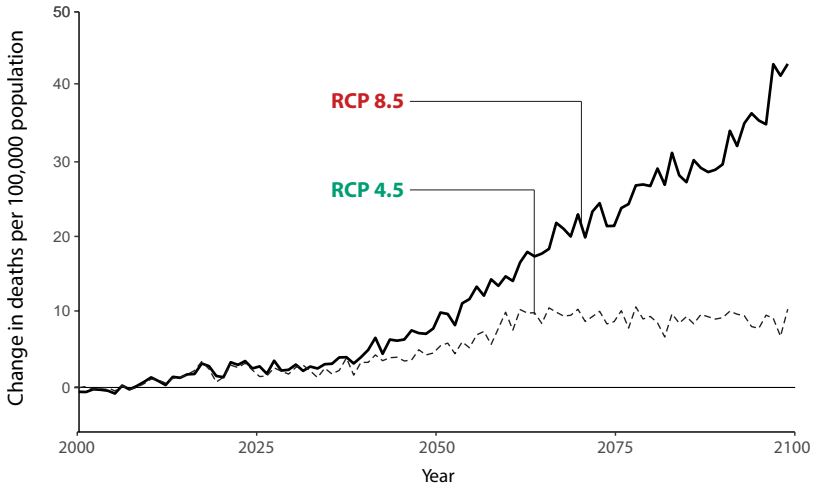

Figure 4.11: Dependence of projected mortality costs of climate change on emissions

Time series of the total mortality costs of climate change shown across distinct emissions trajectories. Both lines indicate total predicted mortality costs due to climate change, accounting for both adaptation benefits and costs, and indicate the median estimate from across 28 climate models. RCP8.5 is a high-emissions scenario, while RCP4.5 is a scenario with aggressive emissions reductions.

evolution of mortality impacts through time, displaying the total number of deaths per impact region. We see that by the end of the century, much of the mortality impact is concentrated in the global south, despite the gains from adaptation shown in Figure 4.9.

Valuation and Damage function estimation

Constructing an empirical damage function relating temperature change to aggregate global welfare losses proceeds in three steps. First, we estimate the excess number of deaths attributable to climate change in each impact region for every future year. These estimates account for projected adaptation as well as its costs. Second, we assign a value of statistical life to climate change driven deaths to translate the losses into dollar values. Third, we repeat this exercise for a range of GCMs under different emissions scnearios to map out the relationship between a given increase in mean global temperature and the estimated mortality cost of climate change.

We value the impact of climate change on mortality using estimates from the literature on the value of a statistical life (VSL). We produce damage estimates using both the VSL used by the EPA in regulatory cost-benefit analysis, $7.6 million in 2006 dollars, and that estimated by Ashenfelter and Greenstone (2004), $2.1 million in 1997

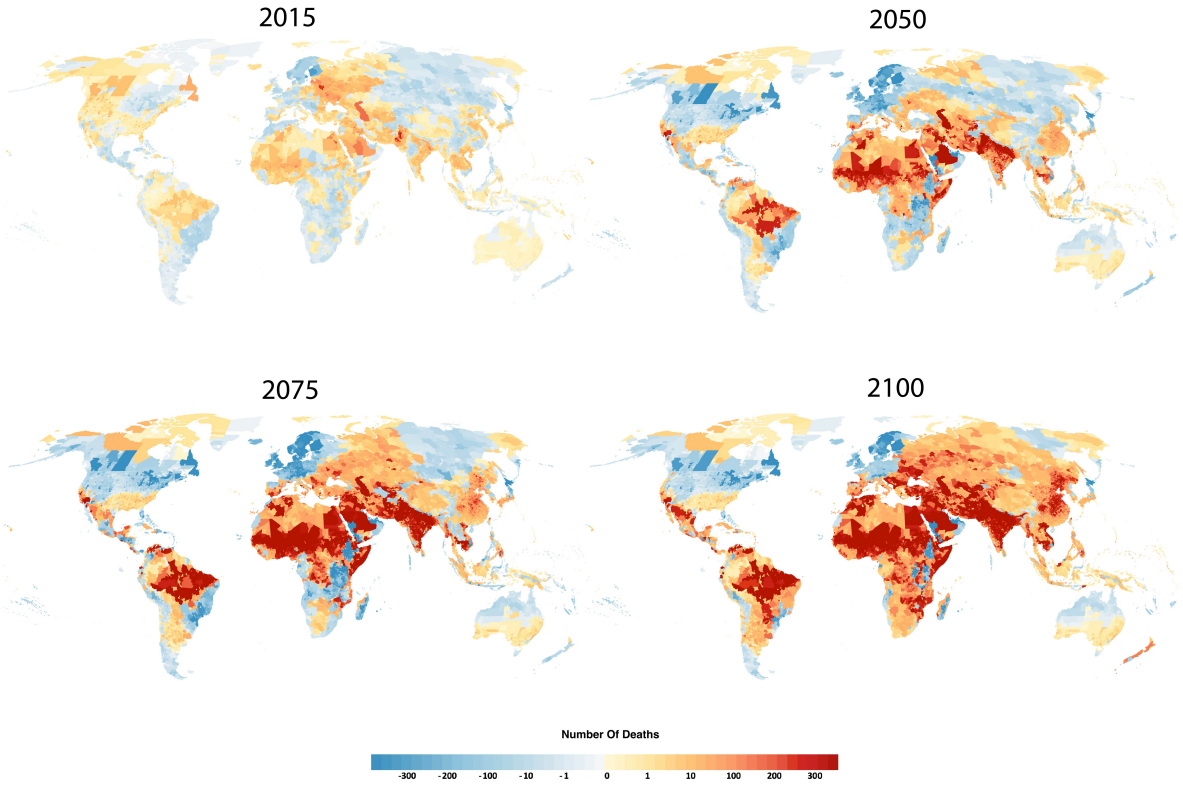

Figure 4.12: Global inequality in mortality costs of climate change

Accounting for adaptation and modeling impacts at high resolution uncovers substantial global inequality in mortality impacts of climate change. Colors show the total number of deaths due to climate change. Estimates come from a model accounting for both the costs and the benefits of adaptation, and indicate a median across 28 climate models. Estimates shown refer to the RCP8.5 emissions scenario.

dollars. To account for the possibility of VSL heterogeneity by age of death, we show the results from two methods of valuing deaths in different age categories. First, we value all age groups at the same standard VSL. Second, we downscale the VSL into a uniform value per life-year and calculate the total value of expected life-years lost.[26]

Assigning a total dollar value to global deaths also requires accounting for differences in income levels across the places where deaths occur. In general, the VSL might vary with income because the level of consumption affects the relative marginal utilities of a small increment of consumption and a small reduction in the probability of death. Consistent with the literature, we use an income-elasticity of 1 to adjust the US estimates of the VSL to different income levels across the world.[27] However, our preferred

[26] We calculate expected life-years lost by applying deaths uniformly within each broad age category and utilizing the future projected life expectancies at each age given by the SSPs.

[27] The EPA uses a VSL income elasticity of 0.7 and a review by Viscusi (2015) estimates 1.1 for the income-elasticity of the VSL.

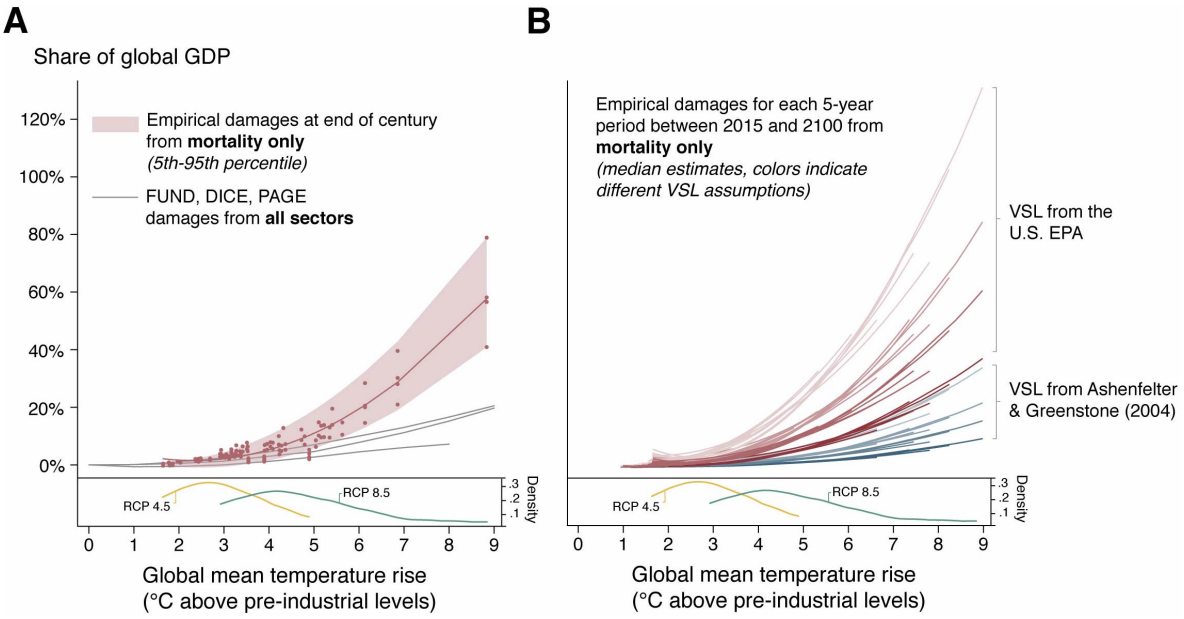

Figure 4.13: Mortality-only climate damage functions

This figure compares empirically-derived mortality-only damage functions to all-sector damages from leading Integrated Assessment Models (IAMs). Panel **A** shows damages at end-of-century, as a function of end-of-century warming above pre-industrial. Each red dot indicates the value of total mortality impacts of climate change between 2080 and 2100, accounting for both costs and benefits of adaptation, for a single climate model. Estimates are pooled across all climate models, emissions scenarios, and growth scenarios, use the U.S. VSL, an income elasticity of one, a globally-averaged VSL applied to all impact regions, and a life-years adjustment. The shaded area represents a quantile regression at the 5^{th} and 95^{th} percentiles, and the light grey lines are estimates of all-sector damages under the IAM models DICE, FUND, and PAGE. Panel **B** shows distinct damage functions estimated independently for each 5-year period between 2015 and 2100 for each set of VSL assumptions. Blue curves use the Ashenfelter and Greenstone (2004) VSL, while red curves use the U.S. EPA VSL. Darker shades indicate more conservative assumptions, as described in the text. In the bottom of both panels are the density of temperature anomalies under two emissions scenarios.

estimate for the value of global damages assigns a homogeneous VSL to different income levels. This preference arises because the social planner does not optimize welfare by income-weighting deaths when consumption sacrifices and averted death come from different agents. Intuitively, the idea underlying heterogeneous VSLs is that lower-income agents value an increment of expected life less than an increment of consumption due to their high marginal utility of consumption, meaning they might prefer to invest less in life-saving technology such as airbags. This logic does not follow in the context of climate change, however, since the tradeoff is between an increment of local mortality risk and the marginal unit of consumption from some global agent that would be sacrificed to mitigate climate change. We capture this idea by valuing all deaths at the VSL of the median global agent, though we also show results for applying the VSL heterogeneously to deaths at different income levels. Results for each valuation assumption and

emissions scenario are given in Appendix Table C.5, showing the damage as a percent of global GDP at three different points in the 21st century.

Finally, to construct a damage function, we aggregate and value the total mortality costs of climate change under a range of future climate scenarios. Plotting these against the global mean surface temperature change for a given model allows construction of an empirical damage function relating costs to the level of temperature increase. Figure 4.13, panel **A**, shows a scatterplot of global mean temperature increase and mortality costs at end-of-century as a proportion of global GDP under all RCP4.5 and RCP8.5 GCM scenarios. Panel **A** contains estimates only for our preferred method of valuation – the EPA VSL, with an age-adjustment for expected life-years lost, and with a common global average VSL across all impact regions. The result in panel **A** shows the core relationship that would be expected to exist between mortality costs and global mean surface temperature change and is an essential component of IAMs for assessing climate change costs. Our central estimate under our preferred method of valuation for the total cost of an approximate expected 4.3°C rise of global mean surface temperature under RCP8.5[28] is 2.7-12.8% of GDP by the end of the century, which is roughly comparable to the estimates from leading IAMs that account for all categories of climate damages (e.g. energy, agricultural output, etc.). Panel **B** of Figure 4.13 shows how this estimated damage function varies both across time, and across sets of valuation assumptions. Each curve represents an independent damage function estimation for a particular 5-year period between 2015 and 2100 and for a particular set of valuation assumptions. Shorter lines indicate damage functions estimated earlier in the century, while curves spanning the full support of temperature anomalies are estimated in later years. Table C.5 in the Appendix reports the value of these curves at three different points through the century for the full range of valuation assumptions and the median of the climate model projected temperatures.

An empirical cost function

While our revealed preference approach enables us to include adaptation costs in an estimate of total mortality damages at end of century, it also facilitates the construction of an empirically-derived cost function. Recall from Section 4.2 that the true net cost function $A(\boldsymbol{b}^*)$ is unobservable, as we cannot measure the costs incurred as populations undertake a wide range of adaptive actions. However, our framework shows that we can use observable effects of a changing long-run climate on the mortality sensitivity to daily temperature to back out both a lower and an upper bound on this unobservable cost function. This innovation allows us to plot bounds on $A(\boldsymbol{b}^*(\boldsymbol{C}_1)) - A(\boldsymbol{b}^*(\boldsymbol{C}_0))$, the change in costs incurred to adapt to a climate that warms from $\boldsymbol{C}_0$ to $\boldsymbol{C}_1$. Figure 4.14 represents this empirically-derived adaptation cost function at the global scale. We predict that costs amounting to approximately 2.3% of global GDP are likely to be

[28]The median global mean surface temperature anomaly relative to pre-industrial across all our models is 4.3°C.

incurred in order to adapt to the median predicted warming at end of century under RCP8.5 (4.3°C).

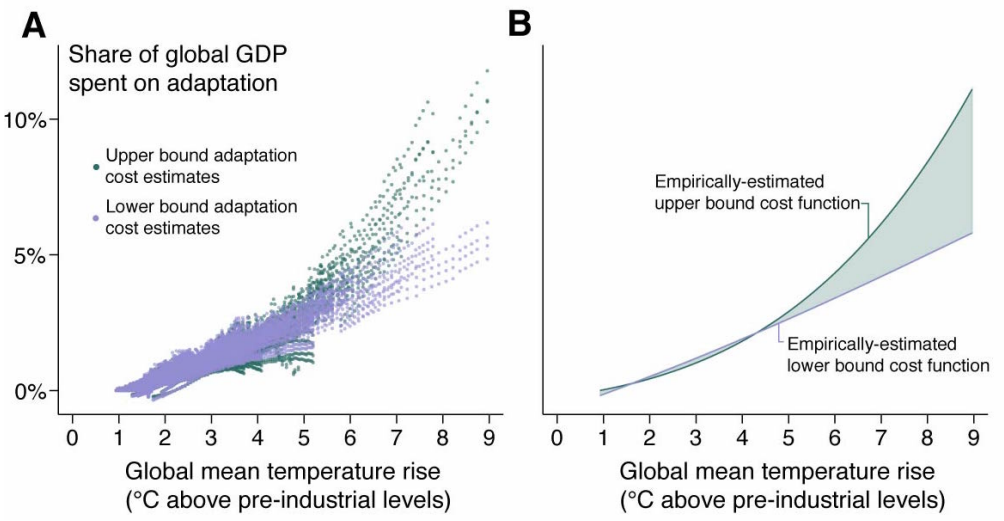

Figure 4.14: Empirically-derived adaptation cost function

Panel **A** shows upper bound (green) and lower bound (purple) estimates of global adaptation costs incurred as a function of global mean temperature rise. Each point represents one climate model in one year, and all years from 2015 to 2100 are included. Costs are shown as a function of current-year global GDP. In Panel **B**, the points in the first panel are used to trace out bounds on an empirical cost function. Regressions in Panel **B** are cubic in global mean temperature rise.

4.6 Discussion

This paper develops a data-driven approach to projecting climate damages into the future, accounting for both the costs and benefits of adaptation. Ultimately, we develop a damage function for the relationship between global health and global temperature that solves many of the problems that have plagued the climate impact literature to date. The damage function is a central component of any analysis of the climate-economy relationship, and is a critical input to IAMs. We define reliable damage functions as those that are based on plausibly causal empirical relationships and that reflect both direct impacts as well as the benefits and costs of adaptation. We argue that the voluminous empirical climate impacts literature can largely be classified into studies which pursue one of these goals at the expense of the other. Our global mortality damage function, derived from the most comprehensive mortality records ever assembled, aims to meet both of these criteria simultaneously.

Projecting climate damages into the future requires us to estimate the extent, benefits, and costs of adaptation. This work builds on a number of distinct advances that have been made in the climate economics and climate impacts literature over the past

few decades. One of the first major steps was to conceptually identify the components of analysis for future research on climate change and the economy, for which the foundation was laid in early work on integrated assessment (e.g., Nordhaus, 1994). A major step in how we understand adaptation was achieved with the development of the "Ricardian" methods in Mendelsohn, Nordhaus, and Shaw (1994), who examine the relationship between farmland values and long-term climate. The Ricardian approach relies on cross-sectional variation in climate; because populations invest in adaptive technologies appropriate to their local climates, this method has the appeal of reflecting many components of feasible adaptation. However, its major weakness is that climate is not randomly assigned across space, and unobserved determinants of human well-being (e.g., soil quality in the context of agricultural productivity) covary with climate. This challenge implies that unobservables confound estimates of the marginal effect of the climate, such that Ricardian methods do not lend themselves to a causal interpretation. The next significant step forward in the literature was the application of panel data methods to climate impacts research (e.g., Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009). Designed to isolate plausibly random variation in climatic conditions, panel models rely on inter-annual variation in weather at a given location to quantify causal effects on human well-being. However, the weakness of this approach is that the range of adaptations available to populations in response to short-run weather variation is likely more limited than, and is definitely distinct from, the adaptive strategies that may be available as agents respond to long-run gradual changes in the climate.

There has been a dramatic increase in the volume of climate impacts research relying on the panel data method in recent years, and it has begun to transform our understanding of the relationship between society and the climate (Dell, Jones, and Olken, 2014; Carleton and Hsiang, 2016). While these findings have guided the analysis in this paper, we have also built upon a trend that has returned to cross-sectional variation to demonstrate heterogeneity in the response to high-frequency variation or shocks (i.e., weather) by examining areas with different distributions of those shocks (i.e., climate) (e.g., Hsiang and Narita, 2012; Hsiang and Jina, 2014; Barreca et al., 2015; Auffhammer, 2018; Heutel, Miller, and Molitor, 2017). While convincingly showing that populations adapt, these previous papers concentrated on the benefits of adaptation, without capturing what the costs of such adaptations will be. Our revealed preference approach to bounding adaptation costs allows us to provide the first empirical estimate of climate damages that explicitly measures both the benefits and costs of present and future adaptation.

To make this contribution, we assemble the largest collection of mortality microdata used in any analysis of the mortality-temperature relationship. Consistent with prior literature, we show in this multi-country, high-resolution sample that extreme temperatures are most damaging to health, particularly for older populations. Direct modeling of heterogeneity reveals important variation in adaptation across the world, with higher average temperatures leading to decreased marginal health damages from heat, and higher average incomes being protective against the hottest days. In addition to showing the benefits of adaptation, we estimate the costs—our revealed preference

approach allows us to quantify the costs of an arbitrarily large set of unobservable adaptive adjustments, relying on sufficient statistics that are measurable in historical data. We find that accounting for both the benefits and costs of adaptation dramatically influences predictions of mortality damages from climate change.

We estimate large damages to human health under climate change, both in terms of number of deaths and monetary costs. In addition to these global aggregate costs, we estimate substantial heterogeneity in projected impacts across the world, showing that climate change will not affect populations equally, and implying that distributional concerns are of first-order importance for climate policy. Mortality costs due to climate change are expected to be one of the largest contributors to the total costs of climate change, and yet they will not be the only costs. A large and growing literature has demonstrated impacts on outcomes as diverse as migration, crime, labor, and education (Carleton and Hsiang, 2016), and the methodology we outline and implement here is generalizable to these other sectors. In particular, the adaptation cost methodology provides a framework for understanding one of the largest missing pieces of this literature to date. The approach used in this paper, when applied across sectors, can be used to produce an overall damage function that reflects expected damages across a wide variety of outcomes. Ultimately, we believe these advances will lead to the development of an empirically-grounded estimate of the Social Cost of Carbon that is transparent in its construction, can keep pace with new data and new research developments, and may form the foundation of climate change cost-benefit analyses and policies around the world.

Chapter 5

The implicit global water market

5.1 Freshwater management as a global problem

Human history is punctuated by water-abundant eras of flourishing, and by water-scarce periods of strife, hunger, and even societal collapse (Diamond, 2005; Hsiang, Burke, and Miguel, 2013; Gleick, 2016). For centuries, we have devised means of collecting, storing, and distributing freshwater resources to ensure resilience against the vagaries of natural supply. Over the last 100 years, massive infrastructural investments and remarkable advances in hydrology, civil engineering and agronomy have ensured that water supplies are clean and reliable in most developed countries, that food production has kept pace with steady population and economic growth, and that hydropower is one of the world's fastest growing sources of energy. However, the challenges of securing stable and safe supply are also rapidly changing: diets are becoming more water-intensive, economic growth has continued to raise water demand per capita, and the environmental damages of large-scale hydrologic interventions, such as wetland losses, fish population declines, and salinization, are becoming more severe. Climate change looms as a driver of increased demand and a threat to both coastal aquifers and alpine storage of freshwater in ice and snow. Superimposed upon these challenges is the fact that throughout the globe, most ground and surface water resources are accessed through some form of common property regime, leading to extraction rates that far exceed socially optimal levels (Hardin, 1968; Gleick, 1993; Culp, Glennon, and Libecap, 2014).

While a large and influential literature in resource economics studies the design of institutions that can ameliorate over-extraction of freshwater (e.g. Gisser, 1983; Schoengold and Zilberman, 2007; Olmstead, 2010; Leonard and Libecap, 2016), empirical validation of theoretical findings has been hampered by poor data. Even in the U.S., basic water withdrawal, storage, and use data are not collected or distributed, despite public outcries over groundwater overdraft in places like California's Central Valley or the High Plains' Ogallala Aquifer (Gleick, 2016). Moreover, this literature has been disproportionately local or regional in nature, despite the growing influence of global-scale forces like climate change and food trade on local water use (Dalin et al., 2017). In this essay, I utilize a novel satellite data source measuring changes in total water

mass at all points in space to make a small contribution toward building empirical estimates of global drivers of local water depletion. In particular, I ask: How do changes to agricultural trade policies impact freshwater resources? I demonstrate that political interventions which increase incentives to domestically produce staple crops have substantial consequences for water availability, providing the first global-scale assessment of the impact of agricultural policy on water depletion rates.

This essay makes an important step forward in empirical estimation of non-local drivers of water depletion. Recent studies of water extraction have begun to provide measurements of management successes and failures (Grafton et al., 2012; Donna and Espín-Sánchez, 2016; Ayres and Meng, 2017; Hagerty, 2017; Debaere and Li, 2016), but these analyses apply to particular basins or aquifers, and are unlikely to generalize. While these case studies have been incredibly important for designing local water policy, the restriction of the existing literature to local-level analyses is potentially troubling. This is because of the growing importance of two global-scale phenomena. First, climate change is altering the spatial distribution of both the demand and supply of freshwater resources, linking global greenhouse gas emissions to local-level hydrology (Piao et al., 2010; Milly et al., 2008). Second, trade in water-intensive goods is expanding, implying that production and consumption decisions in one location have substantial implications for basins and aquifers across the globe. This explicit international market for water-intensive goods generates a global, implicit market for water in which water "embedded" in tradable goods regularly crosses country lines (Chapagain and Hoekstra, 2008; Allan, 2011; Dalin et al., 2017). Regional analyses can generate suboptimal policy recommendations when these global features are ignored; for example, defining a water trading scheme based on historical stream flow may cause overuse if surface water supplies fall under climate change.

Until recently, it has been virtually impossible to study water empirically at a scale that allows for measurement across regions of the world that are linked through the climate or through trading patterns. To overcome this challenge, I employ a novel, globally comprehensive satellite-based dataset that has, to my knowledge, yet to be used in social science. My data come from the Gravity Recovery and Climate Experiment (GRACE), a satellite mission that uses the time-varying distance between two satellites to uncover changes in Earth's gravitational field, which in turn allow for derivation of estimates of changes in water mass at all points on Earth (Tapley et al., 2004). These water mass anomalies have been used extensively in hydrology, geology, and earth systems science to document water loss over time in particular regions (e.g. Rodell et al. (2006); Rodell, Velicogna, and Famiglietti (2009); Famiglietti et al. (2011)). While many studies uncover trends in agricultural regions that are consistent with irrigation-driven water depletion (Rodell et al., 2006; Rodell, Velicogna, and Famiglietti, 2009; Famiglietti et al., 2011; Richey et al., 2015), no prior analyses pair GRACE with social science data to causally attribute water loss to agricultural production or policy.

I address my research question in two steps. First, noting that most political interventions in agricultural markets manifest as wedges between domestic and global food commodity prices, I set up a simple conceptual framework to generate a set of testable hypotheses regarding the influence of these wedges on water use. Second, I combine

satellite-based measures of water – "total water storage" – with a variety of climate and socioeconomic data to estimate a differences-in-differences empirical model that tests for the influence of these policy levers on water availability at the grid cell level throughout the globe. I find, consistent with my conceptual framework, that policies which subsidize domestic production of agricultural output lead to losses in domestic total water resources. These effects are largest when water-intensive crops are subsidized, and are particularly severe in locations suitable to these crops. To put these impacts in perspective, I compare the magnitudes of these effects with jointly-estimated relationships between total water storage and variation in both climate and in night lights, a universally-available proxy for electrification, and hence access to irrigation. I find that on average, a 100 percentage point increase in the wedge between domestic agricultural producer prices and global prices has roughly the same impact on total water storage as changing monthly average temperature from 30°C to 40°C, or as increasing night lights intensity from zero (e.g. rural Ghana) to 14 (e.g. suburban Spain).

This essay proceeds as follows. Section 5.2 outlines a conceptual framework that formalizes the notion of an implicit global water market and its dependence on policy interventions in the explicit market for agricultural goods. In Section 5.3, I describe the water, climate, socioeconomic, and agricultural policy datasets I use, as well as the processing required to combine them. I detail my empirical methodology in Section 5.4. I provide the first evidence that climatic and agricultural footprints can be uncovered statistically in the GRACE dataset in Section 5.5. In Sections 5.6 and 5.7 I show estimation results testing three hypotheses generated in the conceptual framework. I discuss the implications of this work and future avenues of research in Section 5.8.

5.2 Conceptual framework

Explicit water markets are rare within countries, and virtually nonexistent internationally (Schoengold and Zilberman, 2007). Lack of water trade can theoretically lead to misallocation in water use across diverse users, with important welfare consequences (Culp, Glennon, and Libecap, 2014), and much research in water economics focuses on understanding how to build well-functioning local-level water markets (Chong and Sunding, 2006). However, lessons from trade theory suggest we need not necessarily be concerned about water as an immobile factor of production: early trade theory demonstrated (Samuelson, 1953) and very recent theory reaffirmed (Adao, Costinot, and Donaldson, 2017) that there is an equivalence between competitive equilibria in which factors are directly traded and where outputs are traded, while factors are not. This implies that free trade in water-intensive goods could lead to welfare-maximizing allocations of this non-traded resource. As agricultural production accounts for 70-90% of total consumptive freshwater use (Chapagain and Hoekstra, 2008; O'Bannon et al., 2014), freely trading agricultural outputs could enable such an optimal allocation.

Two key impediments limit the possibility that efficient allocation of water resources arises through trade in water-intensive outputs. First, water is grossly underpriced, such that negative externalities permeate market interactions. Second, agricultural markets

are notoriously distorted by political interventions such as subsidies, import tariffs, and export bans (Anderson, Rausser, and Swinnen, 2013). The presence of both distortions generates a classic example of the theory of second-best, in which output market liberalization may not generate welfare gains, nor increase water use efficiency (Chichilnisky, 1994). In this section, I adapt a classic, industry-level small open economy model from the trade and environment literature to formalize the notion of the global implicit market for water, in which both these distortions are captured. I use this model to generate two intuitive testable hypotheses regarding the influence of changes in policy distortions in the explicit market for agricultural goods on domestic usage of water resources. After gaining insight at the aggregate level from this simple model, I add slightly more detail to the production technology to account for within-country and across-crop heterogeneity. This addition provides a final testable hypothesis, and highlights potential interaction effects between policy and exogenous location-specific characteristics that are theoretically ambiguous but empirically identifiable.

Industry level aggregate model

Until recently, most of the literature studying the role of trade costs and trade liberalization on the environment employed models and empirical tests derived at the industry level, implicitly abstracting from any relevant heterogeneity within industries (e.g. Ang and Zhang (2000); Levinson (2009)). The canonical model in this literature, discussed in detail in Copeland and Taylor (2005), is an extension of the classic Heckscher-Ohlin trade model to include an environmental externality. It has not, to my knowledge, been applied to study the role of trade policy on water extraction; the vast majority of the trade and environment literature applies to industrial (i.e. manufacturing) air pollution (Cherniwchan, Copeland, and Taylor, 2017). In this section, I slightly modify this model to provide basic intuition and testable hypotheses that relate policy shocks in the explicit market – i.e. agricultural trade policies – to changes in the quantity of water extracted in the implicit market. This section follows Copeland and Taylor (2005) closely.

Consider a small open economy with two industries, a non-agricultural sector producing a good Y, and the agricultural sector, producing agricultural output X. Both goods are produced with constant returns to scale technology using two primary factors of production, labor L and land K. Let the price of X be p, while Y serves as the numeraire. Production of agriculture is land-intensive, such that

$$\frac{K_x}{L_x} > \frac{K_y}{L_y}$$

where K_x indicates land used in sector X, K_y indicates land used in sector Y, and likewise for labor subscripts. The production function for good Y takes the following general form:

$$Y = H(K_y, L_y) \tag{5.1}$$

While production in both sectors requires land and labor, production in agriculture is generated as a joint output with irrigation water.[1] In particular, suppose that irrigation water is extracted from a common pool resource, at some cost, which includes physical costs as well as government-imposed penalties per unit extracted. Call these extraction costs δ. We should think of the policy-induced penalties as a government's attempt to match private and social costs of water use; in reality, water in agriculture is almost ubiquitously underpriced, often with a price of zero. Thus, while not all water use is socially suboptimal from a welfare standpoint, private costs nearly always fall below social costs, at least in agriculture, due to externalities such as overuse from imperfect property rights, water pollution, land subsidence, and salinization, among others. For now, I put the nuances of these externalities in the background, assuming that δ includes the government's (possibly suboptimal) regulatory means of managing these externalities.

To avoid paying higher water extraction costs, the firm can invest in lowering water requirements per unit of agricultural output by reallocating some share $0 \leq \theta \leq 1$ of its inputs. For example, land and labor could be spent on increasing irrigation efficiency, changing varieties of crops grown, or shifting planting times within the year. Thus, while "potential output" is denoted X (agricultural output that could be realized were all irrigation water available used), we write the joint production of "net output" x and irrigation water z as:

$$x = (1 - \theta)F(K_x, L_x) \tag{5.2}$$
$$z = \phi(\theta)F(K_x, L_x) \tag{5.3}$$

Where, like $H(\cdot)$, $F(\cdot)$ is constant returns to scale, and increasing and concave in its arguments, and where $\phi(0) = 1$, $\phi(1) = 0$ and $\frac{\partial \phi}{\partial \theta} < 0$.[2] The scalar θ represents the share of factor resources expended in an effort to lower water use intensity – a lower θ implies that fewer factors of production are sacrificed to minimize irrigation water, and can instead be allocated to direct production of X. Higher levels of θ generate low water-intensity of output, but come at the cost of reallocating a substantial share of inputs K_x and L_x into water-saving technologies instead of direct generation of agricultural output. Note that "potential output" is realized when $\theta = 0$, and by choice of units, we have in this case that one unit of water is used in the production of one unit of agricultural output:

$$x = F(K_x, L_x) \tag{5.4}$$
$$z = x \tag{5.5}$$

[1]The model can easily be generalized to allow non-agricultural production to also require water. However, as approximately 70-90% of freshwater used by humans globally is estimated to be allocated to agriculture (O'Bannon et al., 2014), the simplification used here closely follows actual production patterns.

[2]Note that this structure implies that the technologies which lower irrigation water per unit of output use the same factor intensity as the production of the final agricultural good X.

Substitution of Equation 5.3 into Equation 5.2 shows that this joint production function can be rewritten to show irrigation water as a third input into production of net output x:

$$x = \left(1 - \phi^{-1}\left(\frac{z}{F(K_x, L_x)}\right)\right) F(K_x, L_x) \tag{5.6}$$

Let there be standard iceberg trade frictions in this economy such that the domestic relative price p^d may exceed or fall below the world price p; that is, an exporter faces prices $p^d = \frac{p}{1+\tau}$ while an importer faces prices $p^d = p(1+\tau)$, and τ represents the magnitude of trade costs. Government policy can influence these relative prices: domestic policy distortions can magnify the wedge between the domestic price under iceberg trade costs and the realized price. These distortions include, for example, the output subsidies and export and import tariffs that make agricultural policy so controversial globally (Anderson, Rausser, and Swinnen, 2013). Non-agricultural output is freely traded without frictions.

Given this setup, we can draw a three-dimensional production possibilities frontier, as in Copeland and Taylor (2005). In the top of Figure 5.1, we have a standard 2-good production possibilities frontier between non-agriculture Y and net agricultural output x. Note that this frontier in $x - Y$ space depends on the level of θ – for higher levels of θ, more factors are allocated to water use efficiency, therefore shrinking the possibilities set. The use of water resources is shown in the third dimension in the bottom panel of Figure 5.1: as net agricultural production increases, water use rises; the slope of this increase is governed by θ, with lower θs generating steeper slopes.

In Figure 5.1, production occurs where the frontier is tangent to the relative domestic producer price p^d.[3] This point is indicated by A.

Changes in technology

Suppose a technology shock lowers the cost of extracting water resources. That is, consider a fall in δ. This decrease will lower firm's optimal allocation of factors K_x and L_x dedicated to water-saving technologies in agriculture, as liberal use of irrigation is now less costly. This is derived formally in Copeland and Taylor (2005), but is intuitive: facing lower extraction costs per unit of water used, firms will reallocate factors of production away from water-saving technologies and toward production of agricultural goods, thus lowering θ. This lower θ shifts out the production possibilities frontier, as shown in Figure 5.2, and for a given price, both agricultural output and total water use rise.

Thus, we have our first testable hypothesis:

1. **Hypothesis #1:** *All else equal, a decrease in the cost of extracting irrigation water will increase agricultural output and increase total water use.*

[3]See Copeland and Taylor (2005) for derivation of this tangency. Note that the producer price for net output falls below the consumer price for agricultural output because some share of potential output is lost to water efficiency investments. This distinction does not affect the production-side comparative statics generated in this section, and is therefore omitted for simplicity.

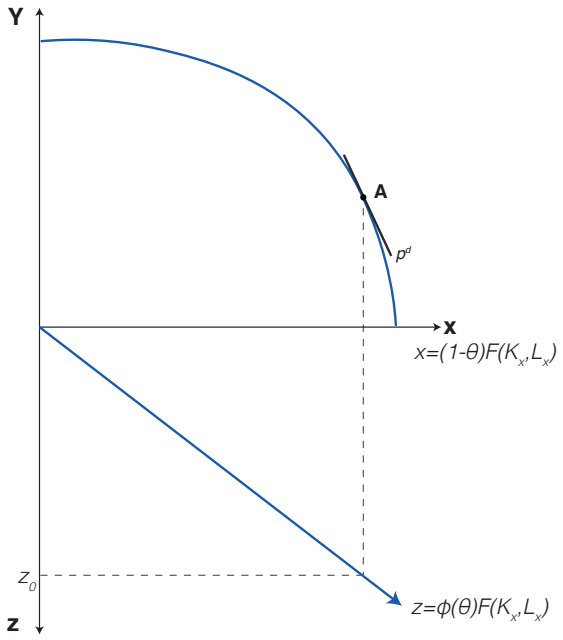

Figure 5.1: Production possibilities frontier with a water-intensive agricultural sector

The top panel illustrates a classic production possibilities frontier (PPF) in $x - Y$ space, representing possible output bundles across non-agriculture Y and net agricultural output x. Agricultural output x is "net" of factors allocated to the reduction of water needs: an endogenous fraction θ of factors used in agriculture are invested in water-saving technologies that lower total water needs in agriculture. Higher government penalties of water use, or increased cost of extraction, will increase θ. In the bottom panel, the use of water z is shown as a function of net output x, where the slope is governed both by θ and the exogenous water intensity function $\phi(\cdot)$. Production occurs at point **A**, the tangency between the PPF and the domestic producer price. Figure adapted from Copeland and Taylor (2005).

Changes in policy

Second, we consider how policies which govern the trade patterns in the explicit market for agricultural goods influence the implicit market for irrigation water. Consider a change in agricultural policies, which govern τ, the wedge between the domestic and global relative price of agriculture. In particular, consider a liberalization which shrinks τ toward zero. Suppose for simplicity that water intensity of agricultural production stays fixed under this change (i.e. θ is unchanged). Under these assumptions, Figure 5.3 demonstrates how such a trade liberalization will affect water use differentially, depending on the comparative advantage of the given country under study.

Figure 5.3 shows liberalization for an importer of agricultural goods. For an importing country, the relative price of X domestically, $p^d = p(1+\tau)$, exceeds the world relative price p. In this case, trade liberalization lowers the relative price of agricultural output, shifting production along the frontier up and to the left. Total water use falls by Δz,

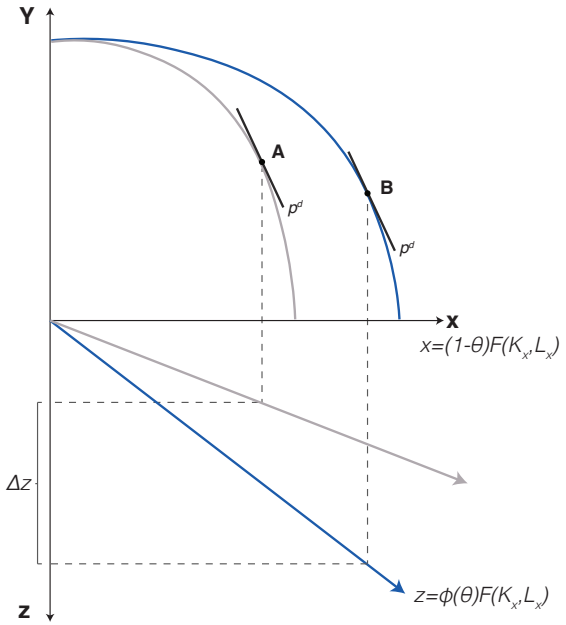

Figure 5.2: Impact on water use of a decrease in cost of using irrigation water

This figure illustrates the effect of a decrease in irrigation costs on water use. The top panel illustrates a classic production possibilities frontier (PPF) in $x - Y$ space, representing possible output bundles across non-agriculture Y and net agricultural output x. Agricultural output x is "net" of factors allocated to the reduction of water needs: an endogenous fraction θ of factors used in agriculture are invested in water-saving technologies that lower total water needs in agriculture. A decrease in costs of extracting and applying irrigation water lowers θ to the blue line, and more water is used for every unit of agricultural production. The PPF shifts asymmetrically out in $x - Y$ space, as more resources are available for agricultural production. The total increase in water use is indicated by Δz, and production moves from point **A** to point **B**. Figure adapted from Copeland and Taylor (2005).

as the economy's mix of goods shifts toward Y and away from X. Thus, liberalization is water-saving in this case. In contrast, an exporting country would have a domestic price below world prices, $p^d < p$. Liberalization brings p^d closer to the world price, shifting the economy toward agricultural production due to the country's comparative advantage in agriculture, and raising water use z. Therefore, the sign of the water use effect of liberalization depends critically on the initial comparative advantage of the country in question. However, if we consider any policy which increases the domestic relative price, be that a liberalization policy or export subsidy for an exporting country or an import tariff for an importing country, factors will reallocate toward agricultural production and water use will increase.

This generates a second testable hypothesis:

3. **Hypothesis #2:** *All else equal, trade liberalization will increase water use for*

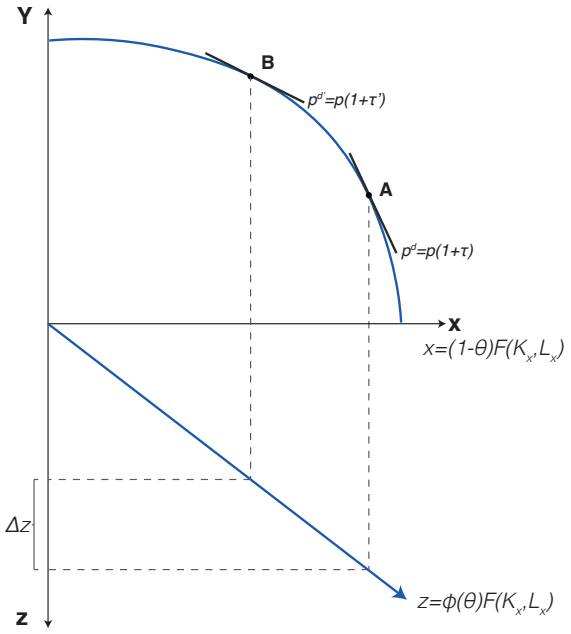

Figure 5.3: The effect of trade liberalization on water use

This figure illustrates the effect of trade liberalization on water use in a country importing agricultural goods. The top panel illustrates a classic production possibilities frontier (PPF) in $x - Y$ space, representing possible output bundles across non-agriculture Y and net agricultural output x. Agricultural output x is "net" of factors allocated to the reduction of water needs: an endogenous fraction θ of factors used in agriculture are invested in water-saving technologies that lower total water needs in agriculture. Trade liberalization in agriculture implies a reduction in τ such that relative producer prices fall from $p^d = p(1+\tau)$ to $p^{d'} = p(1+\tau')$, where $\tau' < \tau$, and production moves from point **A** to point **B**. Figure adapted from Copeland and Taylor (2005).

exporters of agricultural goods, and decrease water use for importers of agricultural goods. For all countries, a policy which increases the domestic relative price of agriculture will increase water use.

Sub-industry, sub-country heterogeneity

The aggregate model above abstracts away from any heterogeneity in production technologies across locations within a country, as well as any heterogeneous production techniques used to produce different crops. In reality, there is wide dispersion in crop-specific productivity within countries (Fischer et al., 2002) and crops have heterogeneous water use requirements (Mekonnen and Hoekstra, 2011). Because the agricultural policies that influence trade patterns are often crop-specific (Anderson, Rausser, and Swinnen, 2013), this high level of aggregation is likely to ignore potentially fundamental relationships between the explicit and implicit markets. Because my empirical setting allows me to observe crop- and location-specific policy shocks and productivities, here

I add detail to the industry-level model described above, to generate two sets of additional hypotheses to take to the data. This addition follows closely the production side of the model developed by Costinot, Donaldson, and Smith (2016). This model was designed to estimate the general equilibrium effects of climate change on agricultural output; it is helpful to frame the study of policy shocks on water use due to its parsimonious means of capturing rich heterogeneity across space and crops. I make small changes to the underlying structure of this model, and then explore basic comparative statics with respect to changes in parameters that affect agricultural production.

As above, our small open economy is endowed exogenously with labor (L) and land (K), and we continue to have just two sectors in the economy, non-agriculture Y and agriculture X. In agriculture, we now allow for multiple crops $k \in \mathcal{K} \equiv \{1, ..., K\}$ to be grown. While labor is freely mobile within a country, land is composed of many fields $f \in \mathcal{F} \equiv \{1, ..., F\}$. Each field f is made up of a continuum of parcels $\omega \in [0, 1]$. In my empirical setting, each field f corresponds to a $1^\circ \times 1^\circ$ grid cell. For simplicity, I assume here that each grid cell is of equal area indicated by a, although in reality the curvature of the Earth implies that areas change with latitude. Under this simplification, however, we can express total land endowment as $L = aF$. Because the consumption side of this model is not necessary for me to generate the partial equilibrium comparative statics I take to the data, I omit its exposition here. However, one should keep in mind that there is a representative agent with utility defined over consumption of non-agriculture and an aggregate of all crops, with a constant elasticity of substitution across crops, as well as a constant elasticity of substitution across varieties of the same crop originating in different countries. Details can be found in Costinot, Donaldson, and Smith (2016).

For non-agriculture Y, let us put some simple structure on the production function in Equation 5.1 and assume this numeraire good is produced using only labor with constant returns to scale technology characterized by an exogenous productivity parameter, A_y:

$$Y = H(L_y) = A_y L_y$$

Recall from Equation 5.6 that the joint production function specified above can be rewritten treating water as an input into agricultural production. Here, we put structure on Equation 5.6 and assume that parcels of land, units of labor, and units of water are perfect complements, such that a firm in location (f, ω) can produce:

$$x^k(\omega) = A_x^{fk}(\omega) \min\{K_x^{fk}(\omega), L_x^{fk}(\omega), z^{fk}(\omega)/v^{fk}(\omega)\} \tag{5.7}$$

In this expression, $A_x^{fk}(\omega) \geq 0$ is location-by-crop specific productivity, which is exogenously determined, and $v^{fk}(\omega) > 0$ is location-by-crop specific water intensity of production, supplanting the industry-wide $\phi(\theta)$ term used in the previous section. The higher is $v^{fk}(\omega)$, the more water is required per unit of labor and land to produce a given level of output.[4] $A_x^{fk}(\omega)$ is randomly assigned, and I denote the average field level productivity for crop k as $A_x^{fk} = \mathbb{E}_\omega[A_x^{fk}(\omega)]$. For simplicity, let us assume that

[4]In the aggregate industry model, more water intensity corresponds to a larger $\phi(\theta)$ value, conditional on θ.

characteristics of each crop fix water intensity such that $v^{fk}(\omega) = v^k \ \forall(i, f, \omega)$. While this is unlikely to be true in reality, in my empirical setting I am only able to exploit heterogeneity in water intensity across crops, but not across space.

As above, let the outside good be freely traded, while crops k have standard iceberg trade costs. In this new setting with preferences defined over varieties of crops originating in different countries, we no longer need to classify countries into exporters or importers of agricultural production. Instead, we let there be iceberg trade costs that are specific both to crops and to trading partners: ν_{ij}^k is the cost associated with shipping a unit of crop k from country i to country j (trade costs within countries are assumed to be zero). Thus, the domestic consumer price in country j for purchasing crop k from country k is $p_{ij}^k = \nu_{ij}^k p_i^k$, where p_i^k is the local price in origin country i.

Profit maximization in the outside good ensures that the wage $w = A_y$, as long as there is sufficient labor to produce the outside good (I assume so). Profit maximization in the agriculture sector requires that each parcel ω is: *(i)* allocated to the crop which maximizes the value of the marginal product, less the costs of water, if this value is less than the wage bill to farm that parcel, and *(ii)* left uncropped if the maximum value of the marginal product less the cost of water falls below the wage bill. Note that the wage bill for any parcel, given the Leontief production technology in Equation 5.7, is simply $w = A_y$.

With δ^f the exogenous marginal cost of accessing water resources for field f, this cost includes, as above, the energy costs of groundwater extraction or the maintenance costs of a canal, in addition to any government penalties for water extraction. Given the conditions *(i)* and *(ii)* above, the probability that a parcel ω in field f is allocated to crop k can be written as:

$$\pi^{fk} = Pr\left\{p^k A_x^{fk}(\omega) - \delta^f v^k = max\{A_y, p^1 A_x^{f1}(\omega) - \delta^f v^1, ..., p^K A_x^{fK}(\omega) - \delta^f v^K\}\right\} \quad (5.8)$$

That is, the probability of allocating ω to k is the probability that the returns to growing k on parcel ω exceed the returns to growing all other crops, as well as exceed the wage bill. Note that because each f is composed of a continuum of parcels, π^{fk} is also the share of f planted to k. If, for simplicity, we let the area of each field be $a = 1$, we can then write total water use at field f as:

$$z^f = \sum_k \pi^{fk} v^k \quad (5.9)$$

Changes in crop-specific policies

Governments can influence domestic producer prices for locally produced varieties of individual crops, by direct subsidies or taxes or indirectly by influencing ν_{ij}^k, which changes demand abroad and ultimately changes domestic producer returns (see Section VI of Costinot, Donaldson, and Smith (2016)). This intervention analogous to the policy-driven changes in producer prices for aggregate agricultural output in the

industry-level model above. Let the effects of these policies on domestic prices be denoted by τ_k. The probability that a parcel ω in field f is allocated to crop k becomes:

$$\pi^{fk} = Pr\Bigg\{ p^k(1+\tau_k)A_x^{fk}(\omega) - \delta^f v^k \\ = max\{A_y, p^1(1+\nu_1)A_x^{f1}(\omega) - \delta^f v^1, ..., p^K(1+\tau_K)A_i^{fK}(\omega) - \delta^f v^K\}\Bigg\} \quad (5.10)$$

All else equal, an increase in τ_k makes production of k more attractive relative to all other crops and the outside option of leaving ω uncropped. This implies that π^{fk} rises, while π^{fm} weakly falls for all $m \neq k$.

$$\frac{\partial \pi^{fk}}{\partial \tau_k} > 0$$

And the effect on water use is:

$$\frac{\partial z^f}{\partial \tau_k} = \sum_{m \in \mathcal{K}} \frac{\partial \pi^{fm}}{\partial \tau_k} v^m = \underbrace{\frac{\partial \pi^{fk}}{\partial \tau_k}}_{>0} v^k + \sum_{m \neq k} \underbrace{\frac{\partial \pi^{fm}}{\partial \tau_k}}_{\leq 0} v^m$$

Note that the first term of this expression is positive, as increasing the returns to k raises the share of f allocated to crop k. The second term is weakly negative, as the share allocated to other crops m may fall. Therefore, the total effect on water use will depend on the relative magnitudes of $\frac{\partial \pi^{fk}}{\partial \tau_k}$ and $\frac{\partial \pi^{fm}}{\partial \tau_k}$, as well as on the relative water intensity of k. The higher is v^k – the more water-intensive is the crop for which the price exogenously rises – the more likely is z^f to rise. Thus, this generates a final hypothesis to take to data:

4. **Hypothesis #3:** *An exogenous increase in the domestic price of crop k will increase water use in field f if crop k is relatively water-intensive.*

Heterogeneous agricultural productivity interacts with policy and technological change

The higher resolution decomposition shown in Equation 5.10 highlights that the influence of crop-specific policies τ_k on water use in a given field f will depend on agricultural productivity A_x^{fk}. Before a price shock, the share of land in field f that is planted to crop k is increasing in exogenous productivity A_x^{fk}. When a policy change of $\Delta\tau_k > 0$ further increases the return to planting crop k, the effect on land allocation and hence water use will depend on initial productivity A_k^{fk}. It is possible that higher productivity locations will generate larger changes in area planted, and thus in water use. Imagine two fields above a shared aquifer resource, where one field is highly productive and the other has poor soil. Initially, wages are sufficiently high that no crops are grown on either field. An increase in τ_k raises the returns to planting crop k for both fields. For

the highly productive field, this cost shock pushes returns for crop k above the wage bill, increasing the share of that field that is cropped. For the low productivity field, this shock is still not sufficient to ensure returns to cropping exceed the wage bill, and no change in proportion cropped occurs.

While this is intuitive, the direction of this second derivative is theoretically ambiguous, without adding further structure on the distribution of productivities across space and crops. For example, if all high-productivity locations are already fully planted to crop k, an increase in domestic price for crop k will only influence cropping patterns and water use in locations of relatively lower k-specific productivity. Therefore, the second derivative is ambiguous:

$$\frac{\partial^2 z^f}{\partial \nu^k \partial A_x^{fk}} = \frac{\partial^2 \pi^{fk}}{\partial \nu^k \partial A_x^{fk}} v^k + \sum_{m \neq k} \frac{\partial^2 \pi^{fm}}{\partial \nu^k \partial A_x^{fm}} v^m \lesseqgtr 0$$

In my empirical setting, I can directly test for this interaction effect using proxies for location-specific productivity A_x^{fk}. Moreover, this interaction between location-specific productivity and a policy shock applies to the irrigation technology shock as well, as shown in Appendix D. However, in this case, a negative irrigation cost shock unambiguously leads to larger water losses in agriculturally productive regions, as shown in Figure D.1. In contrast, for both the sector-wide policy shock and the crop-specific policy shock, the effect on total water storage is theoretically ambiguous without imposition of further structure. In my empirical estimation, I treat $1^\circ \times 1^\circ$ grid cells as heterogeneous fields f, and directly measure the effect of changes in τ (policy shocks to aggregate agricultural output), τ_k (crop-specific policy shocks), and δ (costs of extracting water) on water use in a given location. I allow these treatment effects to vary based on heterogeneous proxies of grid cell productivity, A_x^{fk}, to identify the direction and magnitude of this possibly important interaction effect.

5.3 Data

Basic data tracking groundwater withdrawals, surface water usage, aquifer recharge rates and other fundamental indicators of both human and natural influences on water availability have been largely unavailable, at least for significant regions of the world, prior to the GRACE satellite mission. This paper represents, to my knowledge, the first effort to combine this novel source of comprehensive water storage measurements with additional social and biophysical datasets to uncover human-induced changes in water availability. In this section I describe the GRACE mission in more detail, the additional sources of data I use, as well as the processing required to combine these datasets.

GRACE

The Gravity Recovery and Climate Experiment (GRACE) satellite mission, launched in 2002, is a joint partnership between the National Aeronautics and Space Administration

(NASA) in the United States and Deutsche Forschungsanstalt für Luft und Raumfahrt (DLR) in Germany. GRACE is composed of two identical spacecraft, flying 220 km apart in the same orbital plane about 500 km above the Earth. As the pair orbits the Earth, regions with greater mass concentration will exhibit stronger gravitational pull, affecting the lead satellite first and pulling it away from the trailing satellite. As the two continue circling, the trailing satellite is pulled back toward the lead as it passes over the same gravitational anomaly. Two instruments on board are used to generate extremely precise measurements of this changing distance between the satellites while in orbit: an accelerometer accounts for all non-gravitational accelerations, like those due to atmospheric drag, while a K-band ranging instrument measures distance due to remaining gravitational pull. In this way, only changes in acceleration due to gravity are captured and recorded. These distances are accurate up to one micrometer (μm) per second.

GRACE was designed to improve accuracy of models of the geoid – the hypothetical shape of the Earth under the influence of Earth's gravity and rotation alone, without wind and tidal effects – as well as to provide the first ever high-frequency (monthly) measures of variations in Earth's gravity field. Figure 5.4 depicts the two satellites in orbit and an example of gravitational anomalies from GRACE, where anomalies are shown relative to a mathematically perfect sphere.

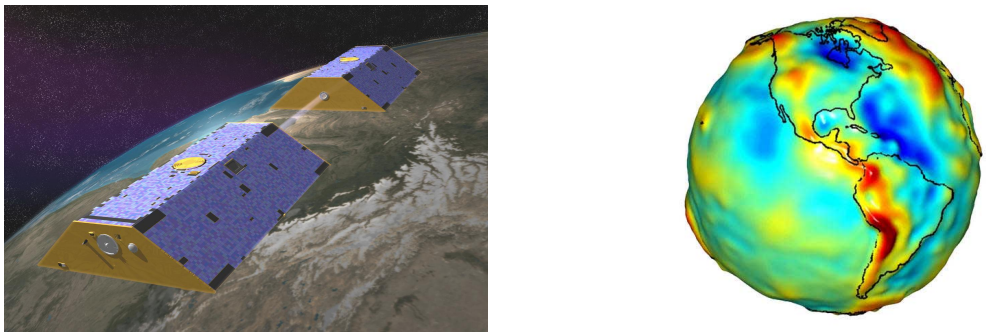

A **B**

Figure 5.4: Gravity Recovery and Climate Experiment

Panel **A** is an artistic representation of the pair of satellites that compose GRACE. Panel **B** is an example gravity anomaly field for the globe from GRACE. Both images were sourced from the Center for Space Research, University of Texas at Austin.

Because water moves in sufficiently large quantities through the hydrologic cycle at a rate far faster than other processes that move mass across the Earth's surface, the month-to-month mass variations uncovered by GRACE are mostly due to changes in water content as it cycles between ocean, atmosphere, continents, glaciers, and polar ice caps (Tapley et al., 2004). This monthly output has been used to study ocean currents (Wahr, Jayne, and Bryan, 2002), measure ground water storage on land (Rodell, Velicogna, and Famiglietti, 2009), and document exchanges between ice sheets

or glaciers and the oceans (Jacob et al., 2012), among many other applications. GRACE "solutions" of these monthly data – a solution converts distances between satellites into estimates of mass – are available at a spatial scale of 220 km × 220km across the entire globe. There are a range of different products available; I use the Goddard Space Flight Center (GSFC) mass concentration solution, which converts time-variable gravity into centimeters of equivalent water height for 41,168 1×1 arc-degree equal-area blocks, called mascons (Luthcke et al., 2013).

I follow the previous literature (e.g. Rodell, Velicogna, and Famiglietti (2009); Richey et al. (2015); Long et al. (2013)) in assuming that month to month changes in mass can be treated as changes in total water storage (TWS), and that the measure of aggregate change in TWS provided by GRACE can be theoretically decomposed into the following elements:

$$\begin{aligned}\Delta TWS = \Delta groundwater + \Delta surface\_water + \Delta soil\_moisture \\ + \Delta snow\_water\_equivalent\end{aligned} \tag{5.11}$$

Nearly all previous literature seeking to tie variations in ΔTWS from GRACE to human activity have focused on isolating the $\Delta groundwater$ component, using administrative data on reservoir levels to partial out changes in surface water, and relying on land use models such as the Global Land Data Assimilation System (GLDAS) or the North American Land Data Assimilation System (NLDAS) to account for soil moisture, snow, and ice (e.g. Famiglietti et al. (2011); Richey et al. (2015)). However, the uncertainty across land surface models is large, making isolation of ground water from aggregate ΔTWS measures error-prone (Long et al., 2013). Moreover, in many parts of the globe, such as the central valley of California, surface water withdrawals are a major source of water use. Therefore, I abstract from decomposition of ΔTWS and directly use this aggregate measure of water storage throughout my empirical estimation.

Gridded ΔTWS data from the GSFC mascon solutions provide me with monthly observations of total water storage, measured as anomalies from the mean gravity field over the period 2004-2014. Covering the years 2003 to 2014, there are 127 observations per grid cell in the GSFC data. Some missing data in 2003 are due to outages while the mission became fully operational, and later holes in the data were caused by solar panels failing to get enough sunlight – these holes occur more often in later years as the batteries aged. Each grid cell in these data is assigned to a hydrologically-determined water basin, of which there are 187 in my estimation sample. One cross-section of these gridded data can be seen below in Figure 5.4. Note that, importantly, I never observe the actual *level* of total water mass at any given point in time and space – because of the nature of the GRACE mission, the data are represented as *anomalies* in the average gravitational field, implying that in all my empirical estimations, I am focusing on changes in the size of anomalies, rather than changes in levels.

Climate, crops, lights

I merge 1° × 1° degree gridded monthly ΔTWS data with three sets of covariates: climate data to capture natural variation in water availability, cropped area to proxy

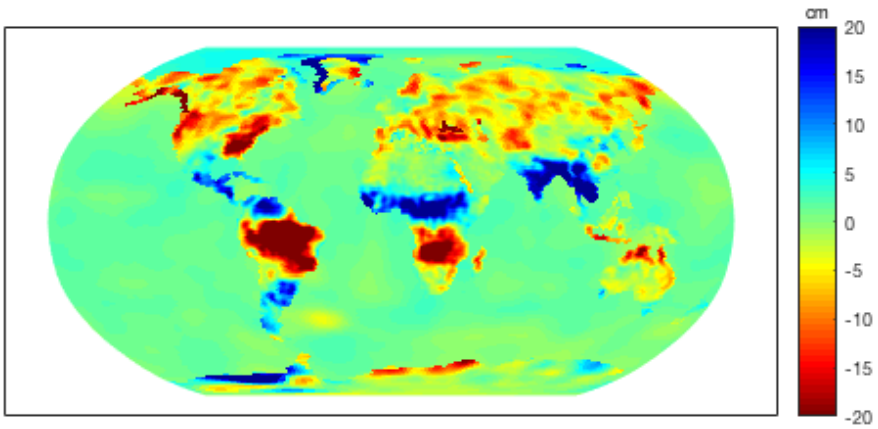

Figure 5.5: Total water storage anomalies, May 2012

Total water storage anomalies are measured in centimeters of equivalent water height. The anomalies shown rely on data from the GSFC mascon solutions of GRACE gravity field anomalies (Luthcke et al., 2013).

for agricultural suitability, and night lights to proxy for the availability of electricity to power irrigation technologies. For climate variables, I use observations of monthly average temperature and cumulative precipitation from Willmott and Matsuura (2014). These temperature and rainfall observations are interpolated from 25,000 weather stations measuring temperature and 27,000 stations measuring rainfall to form a global $0.5^\circ \times 0.5^\circ$ grid. I aggregate these higher-resolution pixels to the lower-resolution ΔTWS grid by simple averaging, where I make any nonlinear transformations used in my empirical specifications (e.g. a quadratic monthly average temperature) at the grid cell level before averaging across space. These data are available from 1900 to 2014.

To measure cropped area globally, I use the global land cover data set created by Monfreda, Ramankutty, and Foley (2008). This high-resolution (approximately 10km $\times$ 10km) product provides cropped area fraction for every grid cell for 175 crops by combining national, state, and country census data with remotely sensed maps of land cover. The most recent year for which it is available is 2000, and I use only 5 crops: rice, wheat, maize, soybean, and cotton. As for the climate data, I take simple averages of crop area fraction to aggregate these pixels up to $1^\circ \times 1^\circ$. Note that these data provide me with proxies for A_x^{fk}, the field and crop specific productivity parameter described in Section 5.2. I have replicated the main results of the paper with an alternative measure of agricultural suitability, the U.N. Food and Agriculture Organization's (FAO) Global Agro-Ecological Zones (GAEZ) Crop Suitability Index, and my findings are very similar.

Finally, I use the National Oceanic and Atmospheric Administration's (NOAA) Global DMSP-OLS Nighttime Lights Time Series (Version 4). These data are derived from satellites that have been circling the Earth multiple times daily since 1970, recording the intensity of lights as seen at night using sensors designed to detect low-light imaging data. NOAA processes raw light intensity data to remove intense sources of

natural light, such that the remaining signal is largely human-caused. These data are available annually at 30 arc-second pixels, excluding very high latitude zones. The value recorded in the data is an integer between 0 (no light) and 63 (full light). This scale implies that a very small fraction of pixels (around 0.1%) are top-coded at 63 – these tend to be wealthy and densely populated areas. These data have been increasingly used to measure economic activity and access to electricity, providing a metric of human prosperity that can be particularly valuable in locations where administrative records are poor (Henderson, Storeygard, and Weil, 2012). I use variation in night lights to capture changes over time in δ^f, the cost of accessing water resources, as described in Section 5.2. The limitations of this proxy should be appropriately noted – while access to electricity enables the adoption of irrigation, it also very likely changes the productivity of the outside option (A_y), as well as other socioeconomic conditions. These results, therefore, should be interpreted with caution and are not designed to be interpreted causally.

Figure 5.6 displays all my covariates, after processing to match all raw data to an identical $1° \times 1°$ grid.

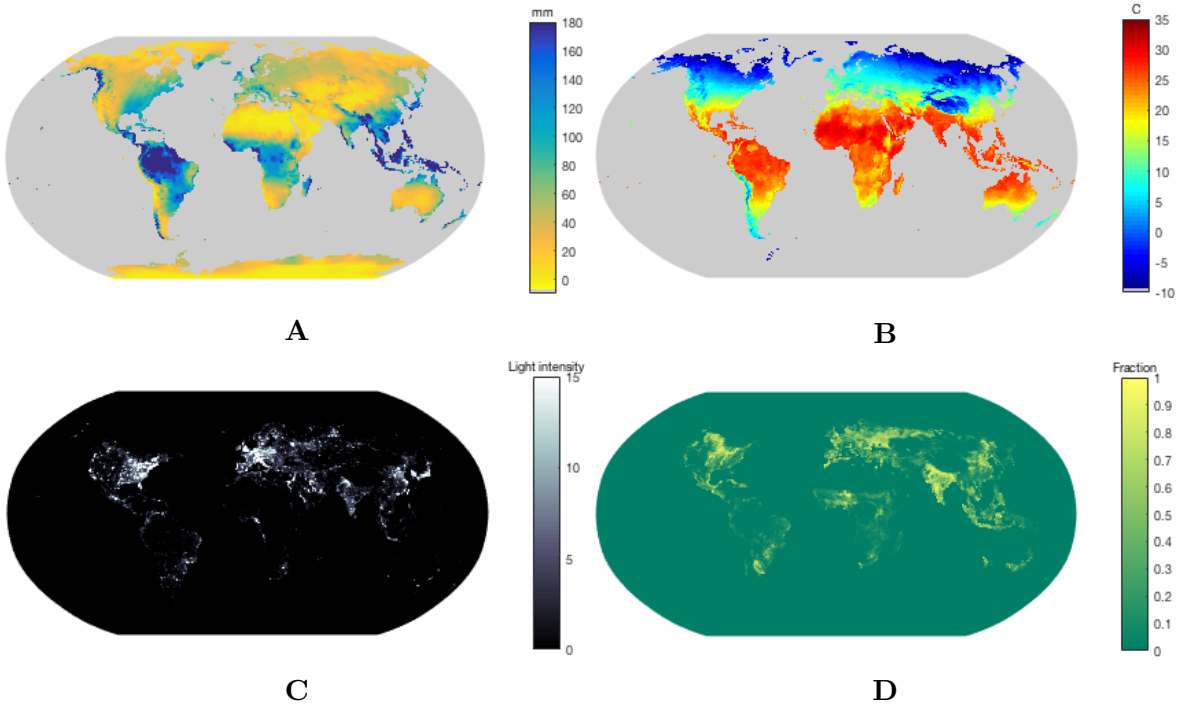

Figure 5.6: Maps of covariates used in regression analysis

Panel **A** shows average monthly precipitation over the sample period, 2003 to 2014, measured in units of millimeters. Panel **B** shows average monthly temperature over the sample period, measured in units of °C. Panel **C** shows average annual nighttime light intensity over the sample period, measured in a unitless indicator of luminosity which ranges from 0 to 63. Panel **D** shows the fraction of each $1° \times 1°$ pixel that is estimated to be planted to any crop in the year 2000.

Distortions to agricultural incentives

To measure political changes to agricultural markets, I use a newly-compiled crop-specific database of distortions in agriculture and food markets across the globe. These internationally comparable data are available from the World Bank's "Distortions to Agricultural Incentives" (DAI) project, which reports the Nominal Rate of Assistance (NRA), a measure of the wedge between current domestic prices and the prices that would prevail without any governmental intervention in the agricultural sector. The DAI database combines detailed price and output data by agricultural product to derive estimates of this wedge under a "small country" assumption. In essence, the NRA is the level of subsidy (if NRA>0) or tax (if NRA<0) imposed on producers of a given product. The NRA is intended to capture output-based domestic taxes and subsidies to farmers, input subsidies or taxes, import tariffs, export subsidies, and government intervention in the foreign exchange market. The NRA is the percentage by which the domestic producer price exceeds the international price of a similar product. NRAs are available for 80 farm products for 82 countries which together account for more than 90% of world's population, farmers, extreme poor, and agricultural GDP. I use these crop-specific values of NRA by country and year to capture variation in τ_k, the crop-specific price distortion described in Equation 5.10. The DAI project also constructs an aggregate NRA across all agricultural products, calculated as the output-weighted average of all covered crop-specific NRA values, plus estimates for the rate of protection on uncovered farm products and other non-crop specific assistance (Anderson, Rausser, and Swinnen, 2013). I use these aggregate NRA values measured by country and year to capture variation in τ, the aggregate agricultural price distortion shown in Figure 5.3. These data are country-by-year-by-crop observations from 1955 to 2011. However, not all countries have data in all years or for all crops – on average a country is observed over 45 years for 12 different crops. The DAI data have recently been used in both studies of political economy of agricultural policy (Anderson, Rausser, and Swinnen, 2013) and of factor misallocation within agriculture (Adamopoulos and Restuccia, 2014).

To merge these administrative records with the rest of my gridded data, I first match the $1^\circ \times 1^\circ$ grid cells from GRACE to countries. Grid cells are assigned to countries based on the location of their centroids. This approximation leads to some lost data, as there are grid cells with some area over land, but with centroids over water bodies; I drop these pixels. While I maintain the monthly resolution of the ΔTWS data, all months within the same year for any given grid cell contain the same NRA value. As described in Sections 5.2 and 5.4, the country-level policy shocks captured by variation in NRA are likely to have heterogeneous effects on water usage based on agricultural productivity of a given grid cell. Therefore, I generate grid cell-specific NRA "treatments" from these country-level observations using cropped area fraction as a means of downscaling to the grid level. Figure 5.7 shows time series of the aggregate NRA metric for 6 example countries, demonstrating the significant spatial and temporal variation in distortionary policy that exists within my sample. Similar time series for wheat and rice, two water-intensive crops, are shown in Figure D.2.

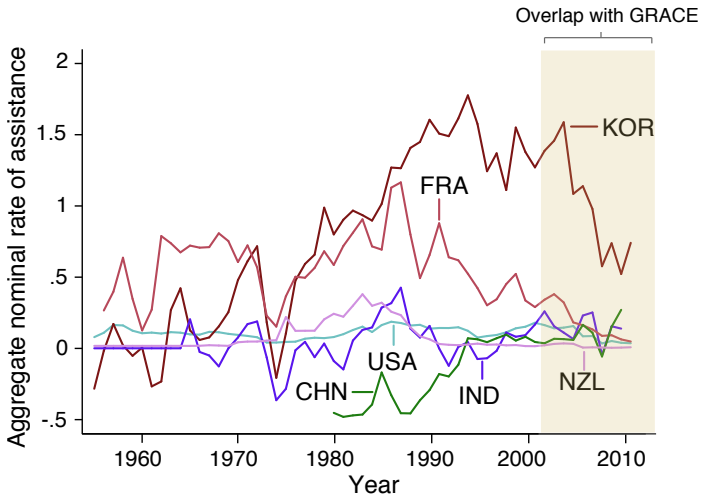

Figure 5.7: Nominal rates of assistance (NRAs) aggregated over all agricultural products

Notes: Time series are shown for 6 countries selected from the full dataset of 82 countries from Anderson (2009). An NRA of 1.5 indicates that after accounting for all distortions imposed by governments, domestic producer gross returns are 150% above counterfactual undistorted gross returns.

5.4 Methods

To my knowledge, this is the first study to link socioeconomic and climatic data to anomalies in the Earth's gravity field as uncovered by GRACE. While many studies have compared time series of GRACE data to other measures of of water availability, such as well gauges (e.g. Long et al. (2013)), I don't believe there has been any previous attempt to use GRACE in more than a descriptive exercise; there have been no studies using these data in a social science context. Before doing so, simple verification that these data – which purportedly capture variations in water storage but also contain any other high-frequency changes in Earth's gravitational field – can be used to statistically uncover meaningful relationships between water availability and climatic and human forcing is required. In this section, I first describe a set of empirical tests designed simply to verify the validity of using GRACE data in linear regression models in combination with other spatial datasets. Second, I outline the empirical approach I use to test the three hypotheses described in Section 5.2, which estimating the effect of changes in agricultural technology and policy on water use.

Using GRACE in linear regression models

Trends in the cross-section

Most existing studies that use GRACE to measure changes in surface and ground water estimate trends in a given location over time, usually relying on qualitative evidence of changes in agricultural production patterns to discuss the extent to which human forces are responsible for either positive or negative observed trends (Richey et al., 2015; Long et al., 2013; Rodell et al., 2006; Famiglietti et al., 2011). To measure these trends globally and relate them quantitatively to other social and physical data, I begin by calculating trends in the monthly total water storage data at every $1^\circ \times 1^\circ$ pixel across the globe with a simple time series regression:

$$\Delta TWS_{it} = \rho_i^W t + \varepsilon_{it} \tag{5.12}$$

Where i indicates a pixel and t is a monthly time variable. ρ_i^W measures the slope of the trend line, in units of cm of equivalent water height per month. To provide initial verification of the usefulness of GRACE data within a linear regression model, I explore the how climatic variables within and outside agricultural areas influence these estimated trends. I perform this test because of clear testable hypotheses. For example, the relationship between ΔTWS, temperature, and precipitation is partly biophysical – hotter temperatures increase evapotranspiration while heavier precipitation can fill surface reservoirs and increase recharge rates in underground aquifers. However, humans also mediate the interaction between temperature, precipitation and ΔTWS in predictable ways. For example, farmers likely respond to dry weather and hot days by increasing irrigation intensities. Therefore, the effects of climate may significantly differ across agricultural and non-agricultural locations. I test for evidence of this likely source of heterogeneity to confirm that I am able to statistically uncover intuitive features that likely drive variation in GRACE. In general, we should see that positive temperature trends lower total water storage, particularly in regions where crops are grown. Likewise, positive trends in precipitation should increase total water storage, with smaller positive effects in regions where crops use and transpire available water resources.

To execute this test, I estimate identical regressions to Equation 5.12, for both temperature and rainfall as dependent variables. Estimated trends are measured in units of degrees Celsius per month, or millimeters of rainfall per month. As a control, I include trends in night lights luminosity, but these results are discussed in more detail below. Call these estimated trends $\hat{\rho}_i^T$, $\hat{\rho}_i^P$, and $\hat{\rho}_i^{NL}$. I conduct a test of the influence of climate variables on ΔTWS by regressing the trends in ΔTWS on the trends in these other covariates. In particular, I interact all covariates with cropped area fraction (CAF), both for an aggregate CAF that includes any major crop, and for a crop-specific CAF for five major crops. The model takes the form:

$$\hat{\rho}_i^W = \beta_1\hat{\rho}_i^T + \beta_2\hat{\rho}_i^P + \beta_3\hat{\rho}^{NL} + \sum_k \alpha_0^k CAF_i^k + \sum_k \alpha_1^k\hat{\rho}_i^T \times CAF_i^k \tag{5.13}$$
$$+ \sum_k \alpha_2^k\hat{\rho}_i^P \times CAF_i^k + \sum_k \alpha_3^k\hat{\rho}_i^{NL} \times CAF_i^k + \gamma_{lat} + \upsilon_i$$

Where k indicates crop (I use five staple crops throughout this project – rice, maize, wheat, soybean and cotton), and CAF^k indicates the fraction of grid cell i that was planted to crop k in 2000. γ_{lat} are latitudinal fixed effects. Standard errors are clustered at the water basin level (there are 187 basins in my data and estimation sample), and the regression is weighted with inverse variance weights, using the variance of $\hat{\rho}_i^W$. In this and in all of my regressions, I limit the sample to grid cells that are between latitudes of -50°N to 60°N, to eliminate confounding effects from ice and glacial melt. While these regression coefficients provide cross-sectional correlations only, I use this specification to simply verify that general trends in ΔTWS respond as expected to climate, crops and night lights. Results are shown in Section 5.5.

Panel model

GRACE data are available at monthly resolution, making possible the estimation of panel data models widely used in applied econometrics. However, most previous research relying on information from GRACE has been focused on lower frequency variation. Therefore, to investigate whether reasonable relationships between climate, irrigation, and cropped area can be uncovered in a model exploiting higher frequency variation, I estimate a standard panel fixed effects model at the grid cell level. In this estimation strategy, I exploit month-to-month variation in temperature and rainfall, as well as annual variation in night lights, to explain changes in monthly ΔTWS. As above, I interact all covariates with cropped area fraction to capture the differential effects that may occur in regions where agriculture dominates the landscape. Because total water storage is a stock variable with nonstationary time series properties,[5] I run all my panel regressions with the first difference of ΔTWS as the dependent variable. Indicating this first difference by $\tilde{\Delta}TWS$, I estimate versions of the following model:

$$\tilde{\Delta}TWS_{imt} = \sum_p \beta_1^p T_{imt}^p + \sum_p \beta_2^p \overline{P}_{imt}^p + \sum_p \beta_3 lights_{imt}^p + \sum_p \alpha_1 T_{imt}^p \times CAF_i \tag{5.14}$$
$$+ \sum_p \alpha_2 \overline{P}_{imt}^p \times CAF_i + \sum_p \alpha_3 lights_{imt}^p \times CAF_i + \psi_i + \gamma_{bm} + \lambda_b^1 t + \lambda_b^2 t^2 + \varepsilon_{icmt}$$

Where p indicates a range of polynomial powers (in my most flexible model, I estimate a 5th-order polynomial), m subscripts indicate month-of-year while t subscripts indicate years, ψ_i are grid cell fixed effects, γ_{bm} are basin-by-month fixed effects to capture

[5] Across a range of panel unit root tests, the GRACE total water storage anomalies fail to reject a unit root.

seasonality, and λ_b^1 and λ_b^2 estimate a basin-specific quadratic time trend. The variable $\overline{P}$ refers to cumulative precipitation over the previous 12 months; I estimate rainfall in this way because there are likely strong lagged effects of rainfall on total water storage. However, changing this specification does not substantially coefficients on any other variables. These fixed effects imply that I identify effects of climate and night lights off of month-to-month changes within a single grid cell, after controlling for basin-level trends and seasonal effects. I show standard errors clustered within basins or within latitudinal bands, and I show results from various sets of temporal fixed effects. Results are shown in Section 5.5: the coefficients of interest are the βs (the impact of each covariate in non-agricultural locations) and the αs (the differential impact of each covariate in fully cropped areas versus uncropped areas).

Estimating the effects of agricultural market distortions on changes in total water storage

After initial verification that relationships between basic climate and human forcing variables and changes in total water storage as measured by GRACE can be feasibly uncovered, I turn to the central focus of this paper: estimating the impact of agricultural policy on water availability. While early studies of trade policy and environmental outcomes conducted cross-country comparisons of trade openness and environmental indicators (e.g. Werner, Brian, and Scott (2001)), cross-sectional regressions are fraught with identification challenges, as trade policy is endogenous to many socioeconomic variables that could also affect the environment. Instead, I follow recent papers in the trade and environment literature in estimating a panel model that facilitates a differences-in-differences comparison with continuous measures of treatment (Martin, 2012; Cherniwchan, 2017). I describe two sets of difference-in-differences regressions, each exploiting distinct temporal variation, but both designed to answer the hypotheses posed in Section 5.2.

Panel model

Recall from Section 5.2 that there are three treatment variables of interest: the cost of accessing irrigation water (δ), the wedge between domestic and global aggregate agricultural output prices (τ), and a crop-specific domestic price wedge (τ_k). Changes in these last two are affected by political forces governing incentives to domestically produce, import, or export agricultural production. To measure the policy treatments, I use the crop-specific and all-crop aggregate Nominal Rate of Assistance (NRA) measures described in Section 5.3, which are defined at the country-by-year level. As described, I assign each grid cell in my data to a country, thereby associating variations in policy distortions over time with all variables measured at the grid cell level – ΔTWS, temperature, rainfall, night lights, and cropping intensity. However, both the aggregate NRA values – which are weighted averages of all agricultural incentives across all food products – as well as the crop-specific NRA values, will impose heterogeneous treatment effects on grid cells within a country, as there is significant variation in within-country

agricultural productivity. As described in Section 5.2, the interaction between these policy levers and location-specific productivity is likely important, but the direction of the effect is theoretically ambiguous. To estimate heterogeneous effects of country-level policy treatments across grid cells, I interact values of τ_k, as measured by NRA, with a corresponding indicator of suitability for agriculture, or suitability for a specific crop in the case of crop-specific NRA treatments. In the results shown here, I use cropped area fraction in 2000 (prior to my sample by two years) as a proxy for suitability, which has the benefit of being strongly correlated with existing cropped area, but not endogenous to year-to-year variation exploited in my estimation.

My main estimating equation is:

$$\begin{aligned}\tilde{\Delta}TWS_{icmt} = &\sum_k \beta_1^k NRA_{ct-1}^k + \sum_k \beta_2^k NRA_{ct-1}^k \times CAF_i^k + \phi_1 lights_{ict} \\ &+ \sum_k \phi_2^k lights_{ict} \times CAF_i^k + \sum_p \alpha_p T_{icmt}^p + \sum_p \gamma_p \overline{P}_{icmt}^p \\ &+ \psi_i + \rho_{bm} + \xi_b^1 t + \xi_b^2 t^2 + \varepsilon_{icmt}\end{aligned} \tag{5.15}$$

Where m indicates month-of-year, c country, and t year of sample. Therefore, ΔTWS_{icmt} indicates total water storage in grid cell i located in country c during month m of year t and $\tilde{\Delta}TWS_{icmt} = \Delta TWS_{icmt} - \Delta TWS_{icm-1,t}$ is the first difference in the total water storage anomaly. When abstracting away from crop-specific policies, I estimate Equation 5.15 with a single policy shifter NRA_{ct-1} in place of the crop-specific shifters NRA^k. I cluster standard errors either at the water basin level or across latitudinal bands, and I show results using various sets of temporal fixed effects. The variable $\overline{P}$ is again cumulative rainfall over the past 12 months; I include this to account for storage of rainfall in soil, surface water and groundwater. Note that I used lagged values of the policy variables NRA^k. This is for two reasons. First, it seems reasonable to assume that production would not respond immediately in the year of a policy change to a new incentive structure. Second, as discussed below, using lagged values makes the requirements for causal identification less stringent.

The parameters of interest are the β coefficients, which, if interpreted causally, measure the impact of agricultural price distortions on local water storage. The identifying assumption to causally estimate the effect of crop-specific NRAs is that there are no unobserved shocks that are correlated both with crop-specific policies and with water loss *in the locations where the affected crop is prominent*. This assumption may be violated – policy changes are by no means exogenously determined. Identification would be threatened if, for example, protection for certain crops was achieved through lobbying from growers producing that crop and simultaneously facing water shortages. However, using lagged policy changes makes such endogeneity less likely; month-to-month changes in water today, conditional on basin-specific trends, cannot drive policies that were determined in previous years. Therefore, a violation would require that there be some underlying political force that, *conditional on basin-specific trends and grid cell time invariant characteristics*, is both correlated with water loss and with political shifts specific to the crops grown where water is being lost.

While imperfect, this is a similar strategy to Martin (2012), who compares changes in emissions in industries affected by a (lagged) change in trade policy to changes in emissions in industries left unaffected. My estimation approach also closely follows Cherniwchan (2017), who compares plant-level emissions before and after NAFTA in industries affected and unaffected by the liberalization, in states with higher and lower trade costs (a triple difference).

Recall the first hypothesis generated in Section 5.2: that a decrease in the cost of accessing irrigation water – i.e. a decrease in δ – will increase total water use through expanding agricultural production and through increasing water used per unit of output. Just like the policy shocks, Figure D.1 shows how the effect of irrigation costs is likely to differ across agriculturally productive and unproductive regions. While I cannot directly test this hypothesis empirically in a causal manner, as I do not have data on exogenous changes in irrigation technology, Equation 5.15 includes luminosity as measured in the night lights dataset as a means of capturing availability of electricity to power irrigation. Because access to electricity dramatically lowers the cost of accessing groundwater, and can aid the diversion of surface water, it serves as a powerful proxy for costs of obtaining and using irrigation water. Interpretation of the ϕ coefficients can provide suggestive evidence either for or against this basic prediction of the effects of irrigation costs on total water use.

Long differences

While year-to-year variation in agricultural policy may observably influence short-run changes in water storage, longer run trends in water availability are the main concern in resource management. Short-run losses or gains in water storage may affect immediate social outcomes, but the magnitudes of long-run trends are what lie behind the commonly discussed concerns of desertification, aquifer drawdown, subsidence, and lake disappearance, among others. Therefore, I estimate a second "long differences" empirical model in which differences are taken over the entire sample period, following the strategy outlined in Burke and Emerick (2016). That is, I estimate:

$$\begin{aligned}\Delta TWS_i = &\sum_k \beta_1^k \Delta NRA_i^k + \sum_k \beta_2^k \Delta NRA_i^k \times CAF_i^k + \beta_3 \Delta lights_i \\ &+ \sum_k \beta_4^k \Delta lights \times CAF_i^k + \alpha \Delta T_i + \gamma \Delta P_i + \nu_{lat} + \varepsilon_i\end{aligned} \tag{5.16}$$

Where differences Δ are taken as follows:

$$\Delta X = \frac{1}{3}\sum_{t=2009}^{2011} X_t - \frac{1}{3}\sum_{t=2003}^{2005} X_t$$

I again cluster standard errors at the water basin level, and I show estimation with latitudinal fixed effects, continent fixed effects, and basin fixed effects. This specification better addresses the policy-relevant dynamics of water loss, but suffers much

more challenging identification problems because trends are compared across grid cells. Causal identification of β coefficients requires that there are no unobserved trends in variables that correlate both with water storage and with crop-specific policy. Because this assumption is harder to support, I emphasize the high frequency differences results in Sections 5.6 and 5.7.

5.5 Detection of a human footprint in GRACE

Time series trends

The first step in my analysis is to estimate grid cell-specific trends in total water storage, and in my main covariates (temperature, rainfall and night lights). These trend coefficients are plotted in Figure 5.8. Panel (a) shows a band of decreasing rainfall in the tropics near the equator, with a band of increasing rainfall in the latitudes just above the equator. Panel (b) shows strong positive temperature trends at northern latitudes, with some cooling at lower latitudes. In Panel (c), we see that night lights are generally trending upward in places where average intensity is already high, such as the US and EU. Finally, panel (d) shows the highly heterogeneous trends in total water storage. Some features are plausibly linked to known human stresses on water availability, such as the large negative trends over the Ogallala aquifer, the Central Valley of California, northern India and Northeastern China. Other features are likely linked to climate change, such as ice melting in Greenland and over the West Antarctic ice sheet. However, there is significant heterogeneity across space in the sign and magnitude of these trends, much of which may be entirely unrelated to human activity. A fundamental goal of this paper is to uncover the extent to which climatic and agricultural drivers are responsible for the variation observed in Panel (d).

Cross-sectional trend regression

As described above, my first set of results serve to verify the that climatic footprints can be uncovered in GRACE measures of total water storage via linear regression. Table 5.1 shows the results of estimating Equation 5.13. Recall that these results are cross-sectional comparisons of the estimated trends described above; I show results both with and without latitudinal fixed effects, and I omit grid cells at very high and low latitudes.

Table 5.1 confirms basic intuition regarding the influence of climatic variables on total water storage. In columns (1) and (4), we see that warming temperature trends are correlated with declines in ΔTWS. The coefficient in Column (4) implies that the average temperature trend of 0.002°C per month corresponded with a decline in ΔTWS of 3mm over the entire 12 year period. We can also see in columns (1) and (4) that rainfall has the predicted effect of increasing ΔTWS: positive rainfall trends are correlated with positive ΔTWS trends. The average rainfall trend of 0.017mm per month corresponds with an increase in ΔTWS of 13.7mm over the 12 year period. Moreover, the climate interactions with aggregate cropped area fraction (CAF) and

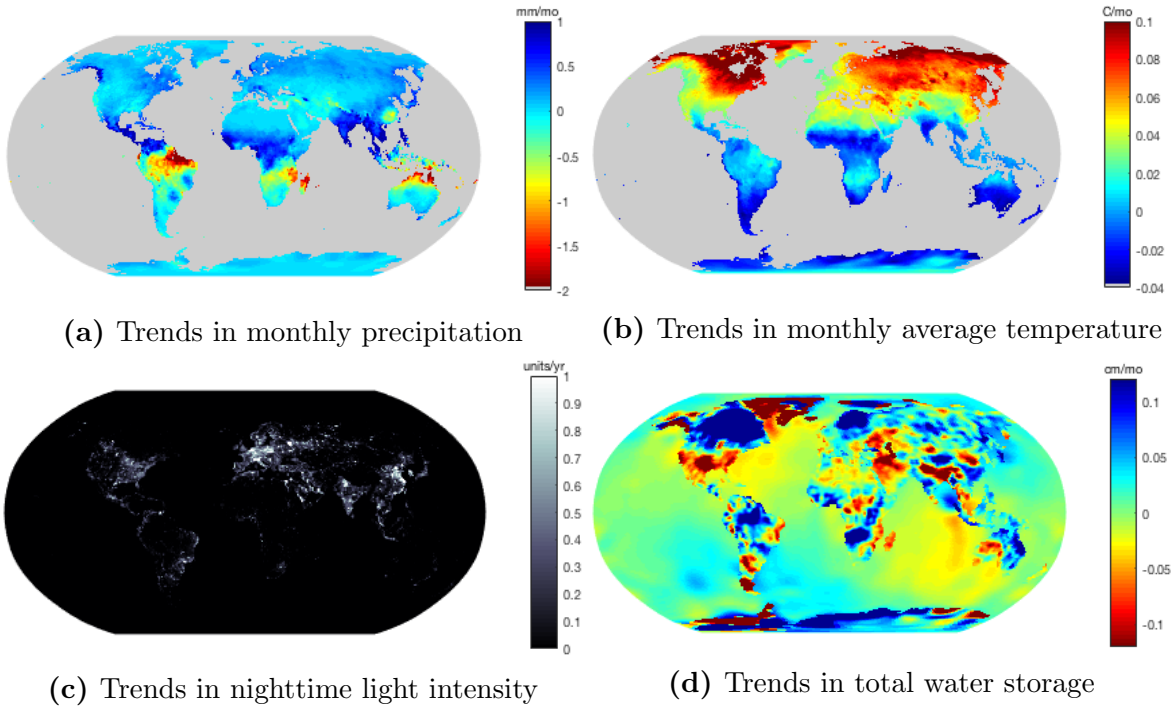

(a) Trends in monthly precipitation

(b) Trends in monthly average temperature

(c) Trends in nighttime light intensity

(d) Trends in total water storage

Figure 5.8: Maps of estimated trends in covariates

Colors indicate the estimated temporal trend in each variable, and are estimated at the grid-cell level using monthly observations for climate and total water storage and annual observations for night lights. All regressions include the years 2003 to 2014.

with crop-specific CAF generate results of predictable signs. For example, in columns (2) and (5) we see that while temperature may have a minimal effect on ΔTWS in a grid cell with no crops, the effect of warming is incredibly damaging when a pixel is fully cropped (i.e. CAF=1). The magnitude of this interaction term on temperature implies that for a fully cropped pixel, the average warming trend over the period caused a loss in ΔTWS of 2.5cm. For a grid cell at the equator, this amounts to 31Gm^3 of lost water, or 0.4% of the estimated annual use of water in crop production worldwide (Mekonnen and Hoekstra, 2011). Likewise, the benefits of increased rainfall are mitigated in fully cropped regions, as presumably some of the gains in water availability are consumed by crops, either naturally or through irrigation. Finally, the crop-specific interactions are also encouraging. As discussed in detail below, rice and wheat are the most intensively irrigated crops globally. In general, Table 5.1 shows that locations with high cropping intensities of rice and wheat respond more negatively to warming trends and to increases in rainfall. All other crop interactions are estimated jointly in Columns (3) and (6), with full results shown in Table D.1; other crop interactions with temperature and rainfall are also negative or insignificant, aside from soybean.

	(1) OLS	(2) OLS	(3) OLS	(4) Lat. FE	(5) Lat. FE	(6) Lat. FE
Trend in temp. (°C/mo)	-0.345**	-0.129	-0.218**	-0.107	0.080*	0.003
	(0.163)	(0.085)	(0.102)	(0.087)	(0.044)	(0.052)
Trend in precip. (mm/mo)	0.070***	0.070***	0.063***	0.056***	0.072***	0.060***
	(0.017)	(0.023)	(0.018)	(0.008)	(0.008)	(0.007)
Crop area fraction (CAF)	-0.026	0.010		-0.001	0.029***	
	(0.016)	(0.023)		(0.010)	(0.011)	
Temp. trend × CAF		-10.866***			-8.931***	
		(2.335)			(0.873)	
Precip. trend × CAF		-0.080			-0.153***	
		(0.091)			(0.054)	
CAF rice			0.457***			0.435***
			(0.064)			(0.032)
CAF wheat			-0.093			-0.040*
			(0.059)			(0.023)
Temp. trend × CAF rice			-7.436			-6.091
			(11.911)			(5.684)
Temp. trend × CAF wheat			-35.833***			-30.900***
			(7.240)			(3.747)
Precip. trend × CAF rice			-0.492**			-0.587***
			(0.215)			(0.158)
Precip. trend × CAF wheat			-0.602			-0.213
			(0.435)			(0.248)
Observations	12,050	12,050	12,057	12,050	12,050	12,057
R-squared	0.027	0.058	0.098	0.441	0.461	0.490
Latitude FE	NO	NO	NO	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table 5.1: Trends in the climate and in electrification explain cross-sectional variation in total water storage depletion

Coefficients shown are from a cross-sectional trend-on-trend regression for 12,050 grid cells over the period 2003 to 2014. The outcome variable is the grid cell-specific trend coefficient on change in total water storage anomalies (ΔTWS), and the explanatory variables are estimated trends in temperature, rainfall, and night lights luminosity. "CAF" indicates the cropped area fraction of each grid cell, and serves as a proxy for crop-specific agricultural suitability. Grid cells above 60° and below -50°C latitude are omitted. All regressions are precision weighted, using the variance of the trend in total water storage. Columns (3) and (6) have controls for night lights, maize, soybean and cotton CAF and CAF interactions, but these coefficients are omitted for display clarity – see Table D.1 for the full set of results. Standard errors in parentheses are clustered at the water basin level.

Panel model

I also verify the usefulness of the GRACE data in a standard panel fixed effects framework where climate variables are modeled nonlinearly, as described in Equation 5.14. The effects of climate on ΔTWS are shown in Figure 5.9. Panel (a) shows that temperature generally decreases ΔTWS, with stronger effects in locations that are intensively cropped, particularly for extremely hot months. The negative effect of warming at high temperatures is corroborated by robustness tests shown in Appendix D, and while the coefficients are not shown in Table D.2, the temperature response function looks

similar in this regression, where agricultural policy variables are the focus and climate variables are used purely as controls. The turning point of the nonlinear temperature function is 30°C, the value at which cereal yields fall precipitously with rising temperature (Schlenker and Roberts, 2009). This may be due to farmers using irrigation as an adaptation strategy in response to hot days, but further tests would be required to pin down this mechanism. Even without further investigation into the temperature response, the negative relationship between temperature and ΔTWS, and the fact that crop intensity exacerbates this response, corroborates the cross-sectional evidence above.

The relationship between cumulative rainfall and ΔTWS is shown in Panel (b) of Figure 5.9. There is a positive correlation between rainfall and ΔTWS for levels of cropped area fraction above average, and it is interestingly stronger for locations that are heavily cropped. This may be due to the fact that crops are generally planted in locations where soil retention of rainwater is higher, where groundwater recharge rates are higher, and even where surface water storage may be more developed. It is beyond the scope of this paper to uncover the mechanism behind this effect. Understanding the determinants of spatially heterogeneous climate response functions in GRACE data, and in particular how these response functions depend on agricultural activity, is a fruitful area for future research.

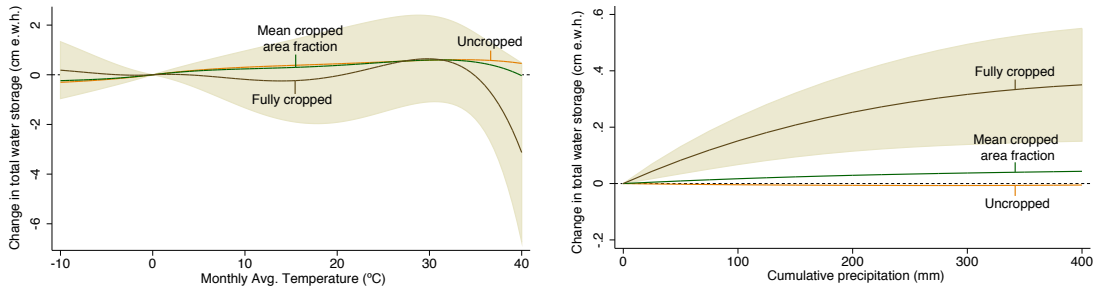

(a) Monthly average Temperature **(b)** Cumulative precipitation

Figure 5.9: Nonlinear impact of temperature and rainfall on total water storage

Results shown are from a panel model with grid cell fixed effects, basin-by-month fixed effects, and basin-specific quadratic trends. "Uncropped" indicates a cropped area fraction (CAF) of zero in 2000, "mean cropped area fraction" indicates the average cropping intensity across all grid cells in the sample, "fully cropped" indicates a cropped area fraction of 1, and "cm e.w.h" indicates units of centimeters of equivalent water height. Standard errors are clustered at the water basin level.

5.6 Electrification and total water storage in agricultural areas

The first hypotheses posed in Section 5.2 concerns the influence of irrigation costs on total water storage. In the cross-sectional trend-on-trend regression, I uncover estimates consistent with the conceptual framework: decreasing costs of irrigation, as proxied by increasing luminosity in nighttime lights, lowers total water storage anomalies. This finding is particularly strong in locations of high agricultural productivity (Table 5.1).

In Figure 5.10, these results are shown to be robust in the main specification of interest, the panel fixed effects model described in Equation 5.14. The general downward sloping relationship is consistent with the first hypothesis: I find that night lights are inversely correlated with ΔTWS, with particularly strong effects in agriculturally productive areas. Increasing from a night lights intensity of zero to an intensity of 20 (for example, think of an un-electrified agricultural region in Africa and a grid-connected agricultural region in India) is associated with a loss of approximately 2cm of ΔTWS per month in a location that is fully cropped. This is a substantial effect: across an entire year, this amounts to a loss of approximately 4% of total global agricultural water use annually (Mekonnen and Hoekstra, 2011).

Results in Table 5.1 show that cross-sectional trend regressions are consistent with the panel fixed effects model. Positive trends in night lights appear to negatively correlate with ΔTWS, at least when latitude fixed effects are included (columns (4)–(6)). The coefficient in column (4) suggests that the average trend in night lights of 0.005 per month is associated with a loss in ΔTWS of 9.8mm over the entire sample period. While columns (2) and (5) show that night lights are less damaging to ΔTWS in cropped areas, this coefficient flips signs to become negative and highly statistically significant in regions growing the world's most water-intensive crops, rice and wheat, as shown in column (6). While providing evidence consistent with my conceptual framework, this result is also shows the potential for using night lights luminosity in tandem with cropped area data as a powerful proxy for the presence of irrigation technologies.

5.7 The impact of agricultural policies on total water storage

In this section, I provide evidence supporting my second and third testable hypotheses discussed above in Section 5.2. First, I demonstrate that policy distortions which increase the domestic producer price of agricultural output – that is, increases in τ – decrease anomalies in total water storage. This is consistent with the hypothesis that these distortions shift the composition of the economy toward agriculture, causing an increase in the use of water, the factor which agriculture uses relatively intensively. Columns (1)–(4) in Table 5.2 show this result: these are estimates from a simplified version of Equation 5.15 in which the independent variable of interest is the aggregate nominal rate of assistance (NRA), a measure constructed by the DAI as a production

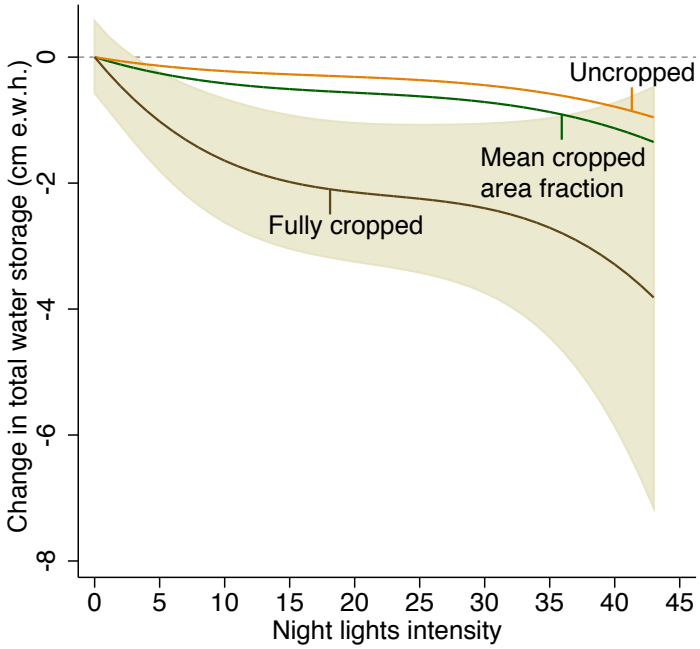

Figure 5.10: Nonlinear impact of night lights luminosity on total water storage

Results shown are from a panel model with grid cell fixed effects, basin-by-month fixed effects, and basin-specific quadratic trends. "Uncropped" indicates a cropped area fraction (CAF) of zero in 2000, "mean cropped area fraction" indicates the average cropping intensity across all grid cells in the sample, "fully cropped" indicates a CAF of 1, and "cm e.w.h" indicates units of centimeters of equivalent water height. Standard errors are clustered at the water basin level.

weighted average of NRAs across all agricultural goods, and where the NRA treatment is held constant across all grid cells within a given country. The negative coefficient on aggregate NRA implies that an increase in distortions favoring agriculture is associated with lower total water storage; these coefficients are negative in all specifications and statistically significant when the most restrictive set of fixed effects are included.

My preferred specification is column (3), which includes grid cell fixed effects, basin-by-month fixed effects, and basin-specific quadratic trends. While the basin-year fixed effects added in column (4) may help account for basin-specific annual shocks not captured by the quadratic trends, these fixed effects may remove a significant portion of the relevant variation in policy. In column (3), the coefficient of -0.5 implies that an annual increase of 100 percentage points in the nominal rate of assistance across all agricultural products is associated with a monthly loss in total water storage of 0.5 cm. To compare magnitudes, this is roughly the same size effect as changing monthly average temperature from 30°C to 40°C for a grid cell with average cropped area fraction. For a grid cell at the equator, this loss of water is approximately 6.2 Gm^3 – over a full year this amounts to just over 1% of agriculture's total global water footprint annually (Mekonnen and Hoekstra, 2011).

	(1) Basin FE	(2) Basin FE	(3) Grid FE	(4) Grid FE	(5) Basin FE	(6) Basin FE	(7) Grid FE	(8) Grid FE
Aggregate NRA	-0.252	-0.332	-0.485*	-1.088*	-0.114	-0.172	-0.283	-0.724
	(0.158)	(0.236)	(0.271)	(0.629)	(0.155)	(0.230)	(0.269)	(0.593)
Aggregate NRA × CAF					-1.057**	-1.310**	-1.493*	-2.988**
					(0.482)	(0.554)	(0.803)	(1.350)
Night lights intensity	0.002	0.001	0.037*	0.003	-0.000	-0.002	0.040*	0.008
	(0.005)	(0.005)	(0.022)	(0.019)	(0.005)	(0.005)	(0.022)	(0.019)
Night lights × CAF	-0.010	-0.007	-0.146**	-0.085**	0.001	0.007	-0.155**	-0.103**
	(0.010)	(0.011)	(0.062)	(0.042)	(0.011)	(0.012)	(0.062)	(0.045)
R-squared	0.403	0.404	0.405	0.407	0.403	0.404	0.405	0.407
Basin trends	YES		YES		YES		YES	
Basin Year FE		YES		YES		YES		YES

Standard errors clustered at the water basin level. Obs. in all regressions = 677,027. *** p<0.01, ** p<0.05, * p<0.1

Table 5.2: Impact of aggregate agricultural subsidies on total water storage

Observations are 1° × 1° grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. All regressions include fifth-order polynomials in temperature and precipitation, and 12 months of lagged precipitation. Standard errors in parentheses are clustered at the water basin level.

As discussed in Section 5.2 and Figure D.1, these effects of policy on total water storage are likely to differ across locations with higher and lower crop suitability. While the direction of this effect is theoretically ambiguous, I empirically test for such heterogeneity by estimate a version of Equation 5.15 that interacts the aggregate NRA variable with a grid-cell specific cropped area fraction (CAF) value, from the year 2000. Figure 5.11 shows the result of this regression for the specification with basin-by-month fixed effects and basin-specific quadratic trends (also shown in columns (5)–(8) of Table 5.2), including the estimated correlation with electrification for comparison. The results display striking heterogeneity across regions by agricultural suitability. The coefficient on aggregate NRA is slightly negative and insignificant, meaning that for locations with low agricultural suitability, the NRA treatment has little to no effect on water storage. However, the strongly negative and statistically significant coefficient on the interaction term with cropped area fraction shows that for locations with high agricultural productivity, increases in distortions that favor agriculture substantially lower total water storage. The magnitude of the result in Figure 5.11 implies that a policy-induced increase of 100 percentage points in nominal rates of assistance lowers monthly ΔTWS by 1.8cm for locations fully planted to crops. This is over 3 times larger than the effect of a uniform industry-wide treatment τ, and represents just under 3% of agriculture's global water footprint annually. This result suggests that distorting general agricultural incentives in locations with high crop suitability leads to substantial loss of water resources.

While Figure 5.11 captures some subnational heterogeneity in policy distortions,

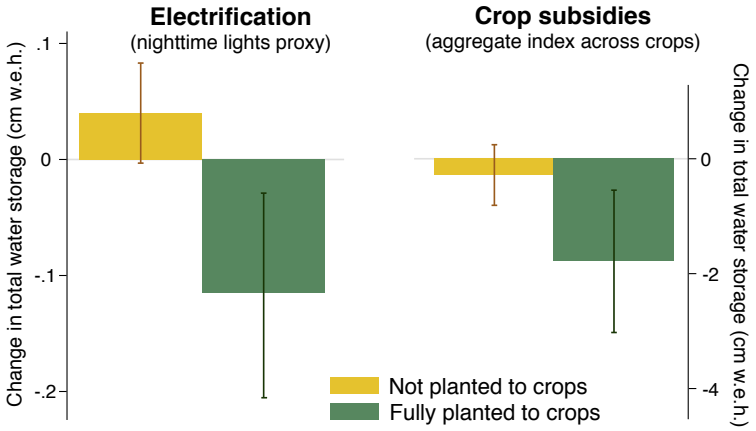

Figure 5.11: Impact of aggregate agricultural subsidies on total water storage: Heterogeneity by agricultural suitability

Results shown are from a panel model with grid cell fixed effects, basin-by-month fixed effects, and basin-specific quadratic trends. Magnitudes in yellow show coefficients corresponding to locations where cropped area fraction was zero in 2000; magnitudes in dark green show coefficients corresponding to locations where the cropped area fraction was 1 in 2000. Standard errors are clustered at the basin level.

using an aggregate NRA for all agricultural products masks any heterogeneity in treatment and response that may exist across different crops. Addressing the fact that policies and water intensity of production both vary substantially across crops enables me to provide evidence to support the third hypothesis from Section 5.2: all else equal, policy-induced increases in the domestic producer price for crops that are relatively water-intensive should cause increases in water use. This effect, as above for the aggregate-level policy shocks, should differ across locations based on agricultural suitability.

Figure 5.12 shows this result, while the full table of estimates is in Appendix D. Policy distortions favoring production of wheat and rice have a strong, negative effect on total water storage in locations with high suitability for these two crops. This result is robust to specification choice, and is statistically significant with standard errors clustered at the basin level. In contrast, distortions to maize, soybean and cotton have no identifiable effect on total water storage in locations where those crops are grown. The magnitudes of these effects imply that a 100 percentage point increase in producer prices for rice is associated with a 5cm loss in ΔTWS for a location planted to rice. The analogous value for wheat is 5.4cm; these effects are approximately 4 times larger than the aggregate NRA effect. Maize has no clear effect, while cotton follows the predicted pattern of larger ΔTWS losses for locations fully planted to cotton, but this result is estimated with large uncertainty. Increased NRAs for soybean have a strong negative effect on ΔTWS, but the heterogeneity based on agricultural suitability is inverted: water losses are large and insignificant in regions *not* suitable for soybean. While this is not inconsistent with the theoretical framework in Section 5.2, it is unclear why this

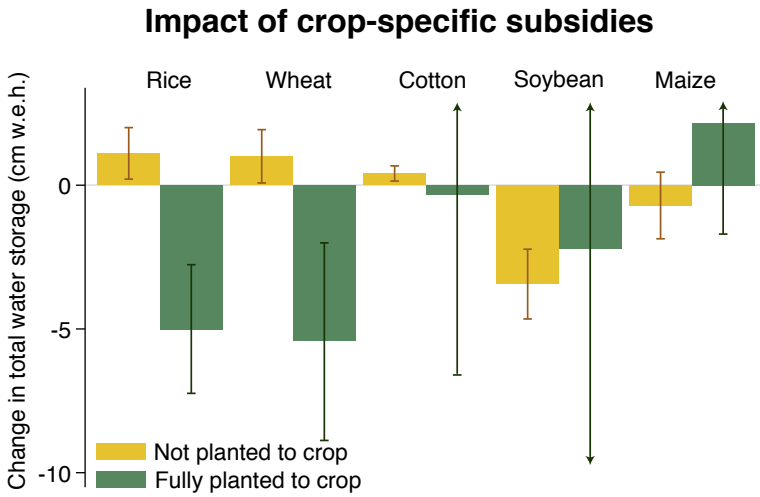

Figure 5.12: Impact of crop-specific agricultural subsidies on total water storage: Heterogeneity by agricultural suitability

Results shown are from a panel model with grid cell fixed effects, basin-by-month fixed effects, and basin-specific quadratic trends. Magnitudes in yellow show coefficients corresponding to locations where cropped area fraction was zero in 2000; magnitudes in dark green show coefficients corresponding to locations where the cropped area fraction was 1 in 2000. Standard errors are clustered at the basin level.

behavior deviates so substantially from that for other crops. As shown in Appendix D, this result holds in most robustness checks and in most subsamples of the data; I leave investigation of this finding to future work.

Recall that the theoretical framework in Section 5.2 suggests that increases in output prices for specific crops should increase water use only when the affected crop is relatively water-intensive. However, water intensity of production by crop is difficult to measure and likely varies substantially across space, depending on local soil characteristics, climate, terrain, and agricultural practices. Available datasets generate estimates of spatially heterogeneous water intensities using agronomic and hydrologic models of varying degrees of complexity. To my knowledge, the highest spatial resolution and most recent product is described in Mekonnen and Hoekstra (2011). This extensive database provides gridded water footprints by crop, separately identifying green water use (use of rainwater stored in the soil as soil moisture), blue water use (use of surface water or ground water), and gray water (the amount of freshwater required to assimilate pollutants used in the production of the crop, such as applied nitrogen fertilizers). The authors rely on a water balance model that accounts for local climate and soil conditions as well as fertilizer application rates to create estimates of what they call a crop-specific water "footprint" or "virtual water content" metric. This metric is expressed in water volume per unit of product that is consumed (i.e. not returned to the local land or water bodies) through the entire process of producing the crop.

These water footprint estimates provide insight into which crops we should expect to

most influence ΔTWS. Because my regressions do not distinguish between components of ΔTWS, as listed in Equation 5.11, agricultural production may be affecting my measure of total water storage both through use of green water – increased production consumes more of available soil moisture – and blue water – increased production in irrigated production systems draws on surface and groundwater resources. However, green water use is inherently constrained by the climate, while blue water use can respond more elastically, as the stock of available surface or groundwater can be drawn upon when production incentives change. Therefore, the crops with the highest blue water footprint are likely those that should exhibit the strongest response with respect to changes in NRA. Mekonnen and Hoekstra (2011) estimate that wheat and rice have the largest blue water footprints of all crops.[6] Together, these two crops account for 45% of global consumptive use of blue water. The total blue water footprint of the five crops I study, according to Mekonnen and Hoekstra (2011), are: 204 Gm^3/yr (wheat), 202 Gm^3/yr (rice), 75 Gm^3/yr (cotton), 51 Gm^3/yr (maize), and 12 Gm^3/yr (soybean).

These blue water footprint estimates indicate that the results in Figure 5.12 and in Table D.2 conform closely with the hypothesis that favorable incentives for water intensive crops lead to water losses. Producer incentives for wheat and rice affect ΔTWS very similarly – NRAs for these crops have strong, statistically significant negative effects on water storage that are robust to various fixed effects specification choices. Their magnitudes are similar, particularly in my preferred specification in column (3) of Table D.2. These effects are large when viewed in the context of global *blue* water footprints alone: the effect of a 100 percentage point increase in the NRA for rice over a pixel fully cropped to rice is a loss of water storage that amounts to 82% of global agricultural blue water use annually.

Maize and soybean, both of which have low blue water footprints and are minimally irrigated throughout the world, show no robust or statistically significant effect on ΔTWS. Cotton is heavily irrigated in many places, and has been associated anecdotally with massive freshwater withdrawals, like that of the Aral Sea, despite having a moderate global blue water footprint. While the interaction term with cropped area fraction for cotton is negative, it is statistically insignificant – this may be due to substantial heterogeneity in production practices across the world.[7] However, it could also be that cotton is not grown in large enough contiguous areas to be reflected in my relatively coarse $1^\circ \times 1^\circ$ grid cells, and further investigation of this result is required.

Unfortunately, because few countries have data for all five of these staple crops, the regressions shown in Table D.2 are limited to grid cells in Australia, Brazil, China, Columbia, India, the U.S. and Zambia. Moreover, because the subsidy data end in 2011, my sample is restricted to 2003-2011 for all regressions including NRAs. I show results in Appendix Table D.4 in which I estimate each crop independently, so as to include far more countries in each regression. These results are qualitatively similar for

[6] Wheat and rice also have the largest total water footprints, estimated by Mekonnen and Hoekstra (2011) at 1087 Gm^3/yr and 992 Gm^3/yr, respectively.

[7] Note that cotton's interaction term is also insignificant in a regression including only cotton NRAs, which makes the sample size much larger and covers many more countries. See Table D.4.

the major water consuming crops wheat and rice, and exhibit negative but statistically insignificant interaction terms for soybean and cotton.

Long differences model

To address the concerns of long-run loss of freshwater resources, I estimate the long differences model in Equation 5.16. A benefit of this model is that while year-to-year variation in NRA is somewhat limited within my sample, there may be more substantial variation to exploit across the full time period through estimating a long difference. However, this is a relatively underpowered regression, as a single difference is taken for each grid cell, and a cross-sectional regression is run across grid cells. Due to insufficient variation in soybean and cotton long differences, I only estimate this model for rice, maize and wheat. Results are shown in Table D.6. NRAs for all three crops demonstrate the predicted behavior, as does the estimated effect for aggregate NRA: the effect of an upward trend in producer prices for a given crop is associated with a gradual loss in total water storage over grid cells that are heavily cropped with that particular crop.

The magnitudes of these long differences estimates imply that for rice, an increase in the NRA of 1 (100 percentage points) over the full sample period is associated with a loss in ΔTWS of 259cm, or about 3.1cm/month, in locations that are fully cropped with rice. To compare, recall that the effect in the panel model of an increase in the NRA of 100 percentage points was approximately 5cm/month. Thus, these effects suggest that in the short run the response is significantly larger than it is in the long run. Likewise, for wheat the long difference estimate is essentially zero, as opposed to a 5.4cm/month loss per 100 percentage point increase in NRA in the panel model. For maize, the sign flips and we see a significant negative effect in the long run. Given the identification challenges of this model described in Section 5.4, I am hesitant to place much weight on the long differences estimates. Moreover, results across different specifications shown in Table D.6 demonstrate the sensitivity of these magnitudes to different sets of fixed effects.

Finally, note that night lights, as in the panel model, are inversely related to ΔTWS in locations that are heavily cropped. Increasing night lights by 1 unit over the entire sample period is associated with a decrease in ΔTWS in cropped areas of 1.2cm, or $<$0.01cm/month. The analogous panel model estimate was an effect of 0.35cm/month, again suggesting smaller effects in the long run than in the short run.

A focus on Asia

Political distortions are unlikely to have the same effect across space, given heterogeneity in production systems, natural resource endowments, climate, and political and institutional environments. To explore some of this heterogeneity, I focus on Asia, where agricultural production is rapidly rising, massive water infrastructure projects have recently been completed or are underway (e.g. the Three Gorges Dam, the world's largest power station), and where total irrigation water use is massive (Mekonnen and

	(1) OLS	(2) Lat. FE	(3) Continent FE	(4) Basin FE
NRA rice × CAF rice	-510.837***	-276.070***	-438.763***	-232.244***
	(102.503)	(48.104)	(110.404)	(86.744)
NRA maize × CAF maize	-89.918	-74.022***	-124.748**	-96.046**
	(60.234)	(22.499)	(54.619)	(45.797)
NRA wheat × CAF wheat	-22.812	-6.457	-19.102	-25.054
	(20.353)	(4.606)	(18.573)	(19.448)
NRA rice	3.262	17.517***	8.325	-7.981
	(4.632)	(3.564)	(6.241)	(6.451)
NRA maize	29.343***	20.998***	13.574	19.370**
	(6.324)	(2.416)	(11.013)	(8.325)
NRA wheat	11.140**	10.499***	29.707**	14.822*
	(4.658)	(1.413)	(12.637)	(8.006)
Night lights intensity	0.159	-0.222	0.283	0.933
	(0.619)	(0.222)	(0.597)	(0.733)
Night lights × CAF	0.317	-0.959***	-0.031	-2.671**
	(0.957)	(0.352)	(1.152)	(1.333)
Observations	3,882	3,882	3,882	3,882
R-squared	0.252	0.710	0.282	0.457
Lat FE	NO	YES	NO	NO
Continent FE	NO	NO	YES	NO
Basin FE	NO	NO	NO	YES

*** p<0.01, ** p<0.05, * p<0.1

Table 5.3: Impact of agricultural subsidies on total water storage: Long differences estimation

Results are from a long differences specification over the period 2003 to 2011. Observations are $1° \times 1°$ grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. Standard errors are clustered at the basin level.

Hoekstra (2011) estimate that the Indus and Ganges river basins in China, India and Pakistan together account for 25% of global surface and ground water use in agriculture). Note that because there are insufficient numbers of basins within Asia alone, I cluster my standard errors at the grid level; this should be improved in future work by estimating Conley standard errors. There results should therefore be interpreted cautiously.

Results in Table 5.4 confirm that Asian countries respond similarly to the full global sample, but with more consistency across crops and specifications. Columns (1) and (2) suggest that aggregate subsidies have a larger effect in these countries than the global average effect; in particular, the effect of aggregate NRA on ΔTWS in non-agricultural areas is quite large. This is also true for rice and wheat distortions, where the coefficient on uninteracted NRA is negative for both crops and where interaction terms suggest that distortions have negative impacts on water storage in locations that are heavily cropped. Using my preferred specification in column (3), the total effect for locations

fully planted to wheat is worse in the Asian sample than in the aggregate sample (loss of 6.1cm for a 100 percentage point increase in NRA vs. a 5.4cm loss), while the total effect for locations fully planted to rice is slightly better in Asia than in aggregate (loss of 3cm versus 5cm). Effects of cotton distortions on ΔTWS are much stronger in Asia than in the global sample, which is sensible given that the crop is heavily irrigated in India and China and accounts for a significant share of agricultural output for some subregions (Zaveri et al., 2016).

In addition, in Asian countries night lights correspond strongly with ΔTWS; Table 5.4 shows a robust negative relationship between ΔTWS and night lights intensity in heavily cropped locations that is slightly larger than the effect in the full sample. This is reasonable, as access to electricity is often cited as a constraint to irrigation throughout Asia, and electric pumps have often been blamed for significant loss of groundwater (Rodell, Velicogna, and Famiglietti, 2009).

5.8 Discussion

Global scale dynamics increasingly govern local water management challenges and solutions. However, for decades, data limitations have limited researchers' ability to empirically ground hypotheses linking large-scale processes, such as climate change and international trade, to depletion of freshwater resources. Both the unfolding of climate change and the structure of international trade are themselves shaped by policies undertaken heterogeneously across the globe. While not explicitly designed to influence water resources, these policies have potentially substantial but currently unmeasured impacts on rapidly diminishing freshwater resources.

In this paper, I use a novel satellite dataset to provide what are, to my knowledge, the first plausibly causal empirical estimates of the relationships between the policies which shape international trade in water intensive goods and rates of domestic water depletion. In particular, I develop a conceptual framework that ties international trade in water-intensive agricultural goods to an implicit global market for water, in which domestic water use is intricately linked to policies which govern the production patterns in explicit markets for agricultural output. My theoretical framework adapts a classic trade and environment model to the context of water use under distortionary agricultural policy, and demonstrates how policy-induced changes to domestic agricultural producer prices influence depletion of domestic water resources. I take this model to grid-level data on changes in total water storage made available by the GRACE satellites, a unique mission that uses measures of changes in Earth's gravitational field to derive estimates of water storage across the globe. My empirical analysis provides the first evidence that changes in the explicit market for agricultural goods have substantial implications for the implicit market for water resources.

In a panel fixed effects framework in which crop and location-specific policy shocks are assumed to be exogenously determined, I find, consistent with my conceptual framework, that policies which increase the domestic price of agricultural output lead to substantial depletion of domestic water resources stored in aquifers, soil, surface water,

	(1) Aggregate	(2) Aggregate	(3) All Crops	(4) All Crops
Aggregate NRA	-2.225***	-12.732***		
	(0.105)	(0.631)		
Aggregate NRA × CAF	-0.164	-1.469*		
	(0.356)	(0.804)		
NRA rice × CAF rice			-1.456***	-3.745***
			(0.540)	(0.675)
NRA maize × CAF maize			0.935	3.879**
			(1.213)	(1.655)
NRA wheat × CAF wheat			-6.032***	-2.578**
			(0.879)	(1.138)
NRA soybean × CAF soybean			2.444	-0.493
			(2.161)	(4.520)
NRA cotton × CAF cotton			-5.282***	-6.350***
			(1.951)	(2.060)
NRA rice			-1.554***	17.787***
			(0.093)	(2.183)
NRA maize			3.367***	-6.051***
			(0.175)	(1.281)
NRA wheat			-0.124	-5.719***
			(0.082)	(0.402)
NRA soybean			-6.320***	-24.328***
			(0.219)	(1.274)
NRA cotton			0.688***	7.227***
			(0.046)	(0.777)
Night lights intensity	0.112***	0.065***	0.054**	0.059***
	(0.022)	(0.018)	(0.022)	(0.021)
Night lights × CAF	-0.514***	-0.288***	-0.486***	-0.260***
	(0.065)	(0.046)	(0.059)	(0.045)
Observations	205,671	205,671	83,300	83,300
R-squared	0.322	0.327	0.306	0.320
Basin trends	YES	NO	YES	NO
Basin Year FE	NO	YES	NO	YES
Basin Month FE	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table 5.4: Impact of crop-specific agricultural subsidies on total water storage in Asia

Observations are $1° \times 1°$ grid cells. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. All regressions include fifth-order polynomials in temperature and precipitation, and 12 months of lagged precipitation. Aggregate NRA effects were estimated independently from the crop-specific effects. Columns (3) and (4) include only India and China due to data limitations. Standard errors in parentheses are clustered at the grid level.

snow and ice. This effect is particularly severe for water-intensive crops, such as rice and wheat, as well as for the locations within country borders that are most suitable

for agricultural production. Consistent with prior work estimating crop-specific water demand, I find no effect of policies which increase incentives to produce corn and soybean on total water storage. While these estimates are the first globally-representative empirical measurements of the effect of agricultural subsidies on domestic water use, my combination of the GRACE dataset with socioeconomic and climatic information also represents the first statistical attribution of GRACE-derived water losses to human activity. I show that commonly discussed drivers of water depletion, such as climate changes, agricultural activity, and proxies for access to irrigation are statistically recoverable in GRACE. I hope that these findings open the door to future work mobilizing the GRACE dataset to answer questions on the economics of global scale water management.

Appendix A

Appendix: Social and economic impacts of climate

Computations for climate impact attribution (Table 2.1)

Most entries in Table 2.1 were taken directly from the text of individual studies (see footnote in table). Here we describe novel calculations for values in cases where attribution was not done by the authors, but either original data were available or summary statistics and estimation results reported in the original study were sufficient to calculate attribution values. In all cases, even when analyses examine the impacts of multiple different climate variables, our calculations focus on temperature. We take this approach because the magnitude of the effect of temperature dominates in most studies included in the table, and because historical and future trends in temperature are much more certain than are trends in other climate variables.

A.1 Effects of the current climate distribution

Our general approach to compute the total impact of the current climate distribution (Table 2.1, Column 5) is the same for all studies. We take the following steps:

1. **Calculate $\hat{\mathbf{y}}^{\mathbf{actual}}$:** Apply the estimated empirical relationship in each study the actual climate observed in the authors' data, predicting outcomes under actual climate.

2. **Calculate $\hat{\mathbf{y}}^{\mathbf{optimal}}$:** Apply the estimated empirical relationship in each study to some "optimal" counterfactual climate, predicting outcomes under a counterfactual climate.

3. **Compare:** Calculate $\hat{\mathbf{y}}^{\mathbf{actual}} - \hat{\mathbf{y}}^{\mathbf{optimal}}$, normalizing the difference by predicted levels of outcomes under the "optimal" counterfactual climate.

For example, with linear models we set optimal temperature to be the minimum temperature experienced for a given panel unit (Figure A.1 A), and for quadratic response functions we set optimal temperature to be the temperature for a given panel unit that maximizes the predicted level of the social outcome (Figure A.1 B). For binned models we set the optimum temperature to be the omitted bin value.

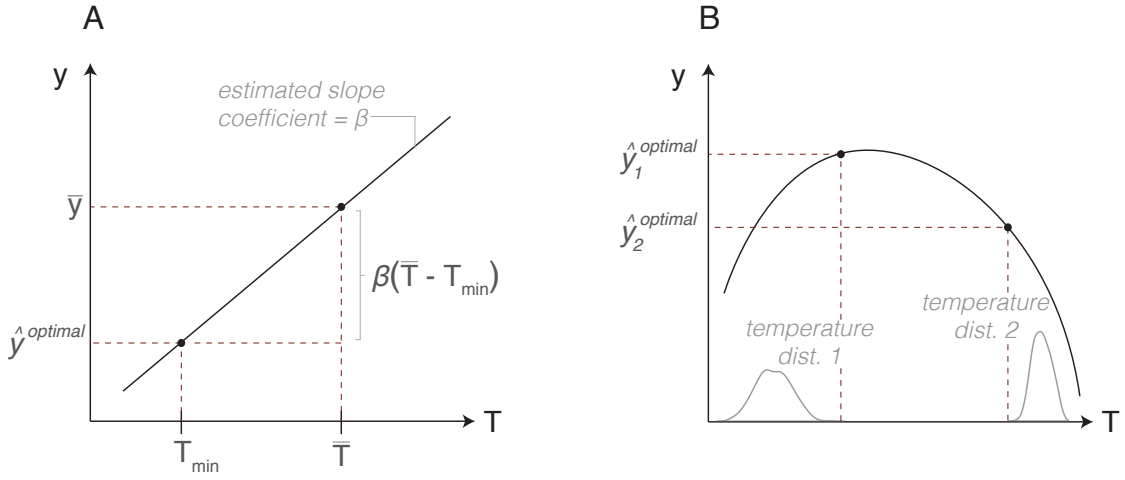

Figure A.1: Identifying optimal counterfactual temperatures

This figure illustrates how optimal counterfactual temperatures are identified with (A) a linear dose-response function and (B) a quadratic dose-response function.

It is important to note that while we describe these computations from the point of view of an "optimal" temperature versus the actual realized temperature, it is equivalent to think of this exercise as the value of fully adapting to temperatures within a population's historical experience. That is, predicted outcomes for a population continuously experiencing an optimal temperature are equivalent to those realized under a flat dose-response function, where changes in temperature have no impact on social outcomes.

Calculation details for each study marked with $(*)$ in Column 5 of Table 2.1 are analyzed are below.

1. **Schlenker & Roberts (2009) - Maize yields in the U.S.** Schlenker and Roberts (2009). We reanalyze the authors' original data using their main degree-days specification, for maize yields only, where i indicates county, s indicates state, and t indicates year:

$$log(yield_{it}) = \beta_1 DD_{0-29} + \beta_2 DD_{>29} + \delta_1 precip_{it} + \delta_2 precip_{it}^2 + c_i + d_s \times t + d_s \times t^2 + \varepsilon_{it}$$

Where DD_{0-29} and $DD_{>29}$ are growing season degree days as described in Schlenker and Roberts (2009) with a single threshold value of 29°C, $precip_{it}$ is total precipitation, c_i are county fixed effects, and $d_s \times t$ and $d_s \times t^2$ are state-specific time trends.

We estimate $\hat{\beta}_1 = 0.00019$ and $\hat{\beta}_2 = -0.006$, consistent with the authors' original result. We use these coefficients to recover predicted values for every county-year observation using observed values of DD_{0-29} and $DD_{>29}$, and exponentiate these values to obtain predicted levels of yields in units of bushels/acre; this is $\hat{y}_{it}^{actual}$. We then compute the counterfactual damage value $\hat{y}^{optimal}$ by setting $DD_{0-29} = 184 \times 29$ and $DD_{>29} = 0$, since in Schlenker and Roberts (2009) there are 184 days in the growing season for maize and because 29°C is the optimal growing season temperature based on the estimated response function (Schlenker and Roberts, 2009). Conceptually, this simulation mimics an experimental setting in which maize is grown in a greenhouse where temperature is set at 29°C continuously throughout the growing season.

We calculate the difference in damages between counterfactual and actual climate distributions as $\hat{\Delta} y_{it} = \hat{y}_{it}^{optimal} - \hat{y}_{it}^{actual} = \hat{y}^{optimal} - \hat{y}_{it}^{actual}$, where the last equality comes from the fact that optimal climate is 29°C for all counties in all years so that predicted outcomes under optimal climate do not change over space or time. In each year, we normalize this damage by the level of predicted yields in the optimal climate, to get the ratio $\frac{\hat{\Delta} y_{it}}{\hat{y}^{optimal}}$, which is the fraction of potential optimal yields that are lost annually due to realized climate. Our reported value is the weighted average of this ratio, where weights in each year across counties are the county's area planted to corn:

$$fractional\_yield\_losses = \frac{1}{\tau} \sum_t \left[\frac{\sum_i \frac{\hat{\Delta} y_{it}}{\hat{y}^{optimal}} \times A_{it}}{\sum_i A_{it}} \right]$$

Where A_{it} is the area planted to maize in each county-year, and τ is the total number of years in the panel (1950-2005). This value is the average (across all counties and all years in the sample) share of optimal bushels/acre that are lost each year due to realized temperatures.

2. **Burke et al. 2009 - Civil conflict in Sub-Saharan Africa** Burke et al. (2009). We reanalyze the authors' original data using their Model 3, for war incidence only:

$$war\_incidence_{it} = \beta_1 T_{it} + \beta_2 T_{it-1} + \delta_1 precip_{it} + \delta_2 precip_{it-1} + \gamma X_{it} + c_i + d_i \times t + \varepsilon_{it}$$

Where i indicates country, t indicates year, T is annual average temperature, $precip$ is total precipitation, X includes all controls in Model 3 in Burke et al. (2009), and $d_i \times t$ are country-specific time trends.

From this estimated relationship (we estimate $\hat{\beta}_1 = 0.0489$ and $\hat{\beta}_2 = 0.0206$, consistent with the authors' original result), we follow a similar process as for Schlenker and Roberts (2009), but set the optimal temperature in the counterfactual (called $T_i^{optimal}$) as the minimum temperature each country is observed to experience in the sample, since the estimated relationship in Burke et al. is linear (see Figure A.1 A). We calculate $\hat{y}_{it}^{actual} = \hat{\beta}_1 T_{it}^{actual}$ and $\hat{y}_{it}^{optimal} = \hat{\beta}_1 T_i^{optimal}$. The difference in war incidence between actual and counterfactual climate distributions is $\hat{\Delta} y_{it} = \hat{y}_{it}^{actual} - \hat{y}_{it}^{optimal} = \hat{y}_{it}^{actual} - \hat{y}_i^{optimal}$. While the model includes lags, because the lag coefficient is often large but never statistically significant, and does not change the magnitude of the contemporaneous coefficient (either in our reanalysis or in Burke et al.'s reported results), we report estimates using the contemporaneous effect only.

To normalize these damages, we need to calculate the baseline risk at the optimal (within-country minimum) temperature. This is not straightforward, because the fixed effects estimation identifies only the slope, and not the level, of the linear response function. To calculate baseline risk, we exploit the fact that the OLS hyperplane passes through the sample mean. Because this is a linear model, we can back out the level of y predicted at $T_i^{optimal}$ as $\bar{y}_i - \hat{\beta}_1 \times (\bar{T}_i - T_i^{optimal})$, as shown in Figure A.1 A. Using this normalization, relative damages in each country-year are $\frac{\hat{\Delta} y_{it}}{\bar{y}_i - \hat{\beta}_1 \times (\bar{T}_i - T_i^{optimal})}$, where $\bar{y}_i$ and $\bar{T}_i$ are mean war incidence and mean annual temperature over the sample for country i. Because this is a linear probability model, when this denominator value is predicted as negative, we set it to zero. Our reported value is the average increase in the annual incidence of war under realized temperatures, relative to baseline risk under optimal temperatures:[1]

$$fractional\_excess\_war\_risk = \frac{\sum_i \sum_t \hat{\Delta} y_{it}}{T \times \sum_i \left[\bar{y}_i - \hat{\beta}_1 \times (\bar{T}_i - T_i^{optimal}) \right]}$$

3. **Ranson 2014 - Crime in the U.S.** Ranson (2014). We reanalyze the author's original data for rates of rape, murder, and aggravated assault using a linear approximation of the author's nonparametric regression on monthly average maximum temperature, as the estimated nonparametric responses in the original paper are very close to linear over most of the support (e.g. see Figure 3 panel n in the main text). We follow a nearly identical procedure to that discussed above for Burke et al. 2009. We include lags, as Ranson does, county-by-month-of-year (c_{im}) and county-by-year (d_{it}) fixed effects. Following Ranson, we weight

[1] Note that we do not take the ratio before averaging because the denominator is zero for some countries. Thus, the interpretation of our final number is the cumulative elevated risk due to temperature, as a percent of the cumulative risk at the optimal temperature, accumulated across all countries in the sample. This is analogous to the "Annual Conflict Risk" (ACR) measure used in Hsiang, Meng, and Cane (2011), i.e. average probability of war in a randomly selected country.

the regression by county population.

$$crime\_rate_{imt} = \beta_1 T_{i,m,t} + \beta_2 T_{i,m-1,t} + \delta_1 precip_{i,m,t} + \delta_2 precip_{i,m-1,t} + c_{im} + d_{it} + \varepsilon_{it}$$

Where i indicates country, m indicates month-of-year, t indicates year, T is monthly maximum temperature, and *precip* is total monthly precipitation. We estimate that $\hat{\beta}_{1,rape} = 0.028, \hat{\beta}_{1,murder} = 0.0028, \hat{\beta}_{1,assault} = 0.351$. In this calculation, we define optimal temperature as each county's minimum value of observed maximum temperature for each month over all years of the sample, because the observations are at the county-month level (e.g. county i will have a January optimal that is distinct from its July optimal). We also account for cumulative effects of the contemporaneous and one-month-lagged temperature effect, as these lagged variables are often statistically significant in both our reanalysis and in Ranson's original study. Thus, each month's optimal temperature exposure is the current month's optimal temperature in addition to the previous month's optimal temperature. We use the same normalization discussed above to get the average risk of crime, relative to the level of crime predicted at each county's minimum temperature. We population weight this average to generate an estimate of the excess crime risk for an average American:

$$fractional\_excess\_crime\_risk = \frac{\sum_t \sum_i \hat{\Delta} y_{i,m,t} \times w_{i,m,t}}{\tau \times \sum_i \left((\bar{y}_{i,m,t} - [\hat{\beta}_1 (\bar{T}_{i,m,t} - T^{optimal}_{i,m,t}) + \hat{\beta}_2 (\bar{T}_{i,m-1,t} - T^{optimal}_{i,m-1,t})]) \times w_{i,m,t} \right)} \tag{A.1}$$

Where τ is the total number of monthly observations in the panel and $w_{i,m,t}$ is the population weight.[2]

4. **Burke et al. 2015 - Global GDP growth** Burke, Hsiang, and Miguel (2015b). We reanalyze the authors' original data using their preferred specification:

$$g_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \delta_1 precip_{it} + \delta_2 precip_{it}^2 + c_i + d_i \times t + d_i \times t^2 + \varepsilon_{it}$$

Where g_{it} is growth in GDP per capita in country i and year t, T_{it} is country average annual temperature, $precip_{it}$ is total country annual precipitation, c_i are country fixed effects, and $d_i \times t$ and $d_i \times t^2$ are country-specific quadratic time trends. We estimate $\hat{\beta}_1 = 0.127$ and $\hat{\beta}_2 = -0.00049$, consistent with the authors' original result.

To get an optimal temperature for each country, we use $T_i^{optimal} = \underset{T_{it}: t \in S}{argmax} \left[\hat{\beta}_1 T_{it} + \hat{\beta}_2 T_{it}^2 \right]$, where S is the set of years included in the original data. This value $T_i^{optimal}$ is the observed temperature for country i that minimizes growth rate damages, given the estimated relationship (see Figure A.1 B). We compute $\hat{g}_{it}^{actual} = \hat{\beta}_1 T_{it}^{actual} + \hat{\beta}_2 (T_{it}^{actual})^2$ and $\hat{g}_{it}^{optimal} = \hat{\beta}_1 T_{it}^{optimal} + \hat{\beta}_2 (T_{it}^{optimal})^2$. The difference in growth rates between counterfactual and actual climates is then $\hat{\Delta} g_{it} = \hat{g}_{it}^{optimal} - \hat{g}_{it}^{actual}$.

[2] As with our calculation for Burke et al. (2009), note that here we do not take the ratio before averaging because the denominator is zero for some counties.

This $\hat{\Delta} g_{it}$ is the difference growth rates across the two climate scenarios, measured in GDP per capita for each country-year. We transform this value into total dollars by multiplying by the level of GDP per capita (Y_{it}) and the population (pop_{it}) in country i and year t; we then sum to get global losses in dollars in year t:

$$unearned\_dollars\_of\_global\_GDP_t = \sum_i \hat{\Delta} g_{it} \times Y_{it} \times pop_{it}$$

We average across all years and divide by global GDP in each year to get the average change in the global growth rate:

$$\begin{aligned} Avg\_growth\_rate\_adjustment &= \frac{1}{\tau} \sum_t \frac{unearned\_dollars\_of\_global\_GDP_t}{global\_GDP_t} \\ &= \frac{1}{\tau} \sum_t \left[\frac{\sum_i \hat{\Delta} g_{it} \times Y_{it} \times pop_{it}}{\sum_i Y_{it} \times pop_{it}} \right] \end{aligned}$$

5. **Deschenes & Greenstone 2011 - Mortality and energy consumption in the U.S.** Deschênes and Greenstone (2011). We do not have access to the original data, so we compute the effects of climate using the summary statistics and coefficient estimates reported in the paper. The authors use the binned temperature specifications shown below (where superscripts m and e indicate the mortality and energy regression coefficients, respectively). Let $\tilde{T}_j = \mathbb{1}(T \in \Omega^j)$, where $\Omega^j = [\underline{T}^j, \bar{T}^j)$ – that is, $\tilde{T}_j$ is an indicator function equal to 1 if temperature is in the set $[\underline{T}^j, \bar{T}^j)$, and $\tilde{P}_k$ is defined analogously for precipitation:

$$mortality\_rate_{it} = \sum_j \beta_j^m \tilde{T}_{itj} + \sum_k \delta_k^m \tilde{P}_{itk} + c_i^m + d_t^m + \gamma_{st}^m + \varepsilon_{it}^m$$

$$\log(energy\_demand_{st}) = \sum_j \beta_j^e \tilde{T}_{stj} + \sum_k \delta_k^e \tilde{P}_{stk} + \mathbf{X}_{st}\beta^e + c_s^e + \gamma_{dt}^e + \varepsilon_{it}^e$$

Where c_i are county fixed effects, γ_{st} are state-by-year fixed effects, γ_{dt} are census division-by-year fixed effects, and $\mathbf{X}_{st}$ is a vector of state-level covariates, including population and GDP (see Deschênes and Greenstone (2011) for details).

We use summary statistics on the average annual number of days that temperature falls into each bin across all panel units, as well as the estimated coefficients for each bin, to compute the total effect of climate. The number of days in each bin are taken from Figure 1 in Deschênes and Greenstone (2011). Note that because these average values are population-weighted across all the county-year observations in the sample, our estimates of the average excess mortality risk and excess energy demand will both be population-weighted averages. The mortality

coefficients $\hat{\beta}_j^m$ are taken from Figure 2 and the residential energy use coefficients $\hat{\beta}_j^e$ are from Figure 3. Because the energy demand specification is log-linear, we report the log-transformation of the cumulative damages from days in every bin to get the average percent change in energy consumption due to temperature:

$$fractional\_change\_in\_energy\_demand = \exp\left(\sum_j \hat{\beta}_j^e \times \overline{\tilde{T}_j}\right) - 1$$

Where $\overline{\tilde{T}_j}$ is the mean value of $\tilde{T}_j$ across the sample. Mortality rates are not logged in the authors' specification. Thus, we compute the cumulative impacts across all bins, relative to the the predicted level at the optimum. As above, we estimate the predicted level of mortality at the optimum using the fact that the OLS hyperplane passes through the sample means.

$$fractional\_excess\_mortality\_risk = \frac{\sum_j \hat{\beta}_j^m \times \overline{\tilde{T}_j}}{\overline{y} - \sum_j \hat{\beta}_j^m \times \overline{\tilde{T}_j}}$$

Where $\overline{y}$ is the average mortality rate in the sample. As the original article did not provide the mean all-age mortality rate, we download the publicly available version of the outcome data and multiply each age-specific mortality rate in these data by the age-group population weights used in the original article to get a mean mortality rate of 859 per 100,000.

6. **Zhang et al. 2016 - TFP in China** Zhang et al. (2016). We follow an identical approach as above for energy consumption in Deschênes and Greenstone (2011), as we do not have access to the original data and Zhang et al. use a binned temperature response function with the log of TFP on the left hand side. We use summary statistics for the average number of days in each bin shown in Figure 2, and TFP coefficients from Figure 4 of the original article. We calculate total losses of TFP as a percent of optimal:

$$fractional\_reduction\_of\_TFP = \exp\left(\sum_j \hat{\beta}_j \times \overline{\tilde{T}_j}\right) - 1$$

A.2 Effects of climate change to date

Very few papers in the literature conduct warming-to-date attribution exercises (calculations in Table 2.1, Column 6). Nonetheless, for two papers where we have sufficient data, we compute estimates for the effect of warming temperature trends since 1980, following the approach outlined in Lobell, Schlenker, and Costa-Roberts (2011). We do

not consider trends in any other variables, although some other papers do report impacts of trends in variables such as precipitation (Lobell, Schlenker, and Costa-Roberts, 2011) or pollution (Auffhammer, Ramanathan, and Vincent, 2006).

The approach outlined below to compute the total impact of recent warming trends is analogous to that employed above to measure the impact of the current climate distribution. Now, however, our two quantities are $\hat{\mathbf{y}}^{\mathbf{actual}}$ – the predicted social outcome under actual climate (including any warming that has occurred) – and $\hat{\mathbf{y}}^{\mathbf{detrended}}$ – the predicted social outcome under a counterfactual de-trended climate, where warming trends are removed. We then compare the difference between these quantities and normalize in an identical manner to the calculations in Section 1 of this supplement.

1. **Burke et al. (2009) - Civil conflict in Sub-Saharan Africa** Burke et al. (2009). These data are country-year observations covering the years 1980-2006. For each country in the data, we estimate a linear trend in average annual temperature and predict temperatures in each year, calling these predicted values T^*_{it}. E.g., denote the predicted country-level temperature in country i during 1981 using this linear fit as $T^*_{i,1981}$.

   We then create a de-trended temperature residual for every country-year as follows, which is normalized to temperature in 1981 (see Figure B.3):

$$T\_detrended_{it} = T_{it} - T^*_{it} + T^*_{i,1981}$$

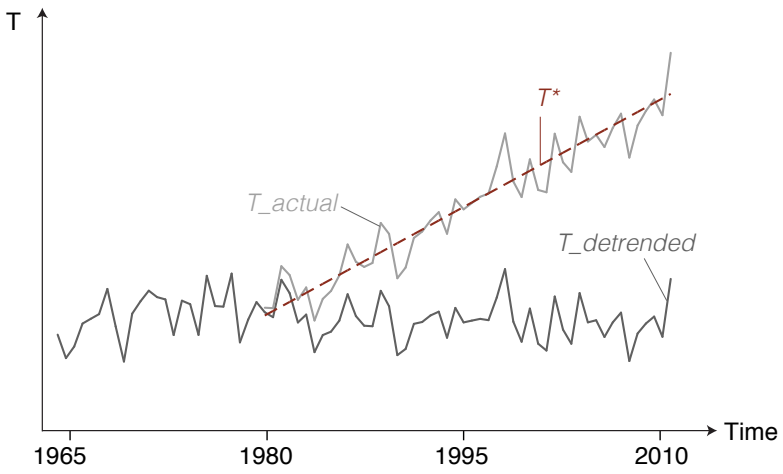

Figure A.2: Identifying counterfactual de-trended temperatures

We predict conflict levels using actual and de-trended temperature, using the coefficient estimates from our re-analysis of the originally reported empirical model. Excess conflict risk due to the trend, relative to the de-trended counterfactual, is the difference between these two predictions, which simplifies to $\hat{\Delta} y_{it} = \hat{\beta} \times (T_{it} - T\_detrended_{it})$. Integrating over countries and years provides total additional risk born due to recent warming. We report this number as a

percent change relative to the total risk in the de-trended climate (recall Figure A.1 A) and report the following total damages due to recent warming, relative to a counterfactual de-trended climate, as:

$$excess\_conflict\_risk\_from\_warming = \frac{\sum_i \sum_t \hat{\Delta} y_{it}}{\sum_i \sum_t \left[\bar{y}_i - \hat{\beta}_1 \times (\bar{T}_i - T\_detrended_{it}) \right]}$$

Note that the values we report are averages over the 20+ years of warming in the sample obtained from Burke et al. (2009). Comparing outcomes just for years at the end of the sample leads to higher estimated impacts of warming, as the trend has been generally linear since 1980.

2. **Burke et al. (2015) - Global GDP growth** Burke, Hsiang, and Miguel (2015b). These data are country-year observations covering the years 1960-2010. We estimate country-specific linear trends in temperature only using data after 1980 (inclusive) to generate T^*_{it} and $T\_detrended_{it}$, as above. We then calculate $\hat{\Delta} g_{it} = \left[\hat{\beta}_1 T^{actual}_{it} + \hat{\beta}_2 (T^{actual}_{it})^2 \right] - \left[\hat{\beta}_1 T\_detrended_{it} + \hat{\beta}_2 (T\_detrended_{it})^2 \right]$.

   As described in Section 1.4 of this supplement, we transform this value into total dollars of GDP, averaged across all years, and divide by global GDP in each year to get the average damages to the global average income growth rate:

$$\begin{aligned} Avg\_growth\_rate\_adjustment\_from\_warming &= \frac{1}{\tau} \sum_t \frac{unearned\_dollars\_of\_global\_GDP_t}{global\_GDP_t} \\ &= \frac{1}{\tau} \sum_t \left[\frac{\sum_i \hat{\Delta} g_{it} \times Y_{it} \times pop_{it}}{\sum_i Y_{it} \times pop_{it}} \right] \end{aligned}$$

Again, note that the values we report are averages over the damages each year from 1980 to 2010.

Appendix B

Appendix: Climate and suicide in India

B.1 Materials and methods

Data appendix

I compiled suicide and climate data at the state level for the years 1967-2013, and agricultural yield and climate data at the district level for the period 1956-2000. Summary statistics for key variables of interest are provided in Table B.1.

Variable	Mean	(Std. Dev.)	Min.	Max.	N
Suicide data: 1967 - 2013 (32 states)					
Suicide rate (deaths per 100,000)	11.4	(11.9)	0	73.23	1472
Growing season daily degree days > 20°C	5.32	(2.93)	0	12.23	1645
Nongrowing season daily degree days > 20°C	3.85	(2.83)	0	9.56	1645
Growing season precip. (mm)	1186.18	(626.85)	111.16	4461.3	1598
Nongrowing season precip. (mm)	439.12	(361.48)	5.76	2148.4	1598
Agricultural data: 1956 - 2000 (271 districts)					
Log annual yield (Rupees per ha)	3.92	(0.72)	-1.87	6.45	11289
Growing season daily degree days > 20°C	6.7	(2.6)	0	14.99	11780
Nongrowing season daily degree days > 20°C	4.56	(1.79)	0	9.51	11780
Growing season precip. (mm)	870.55	(467.27)	10.74	4663.99	11780
Nongrowing season precip. (mm)	205.18	(185.12)	0.83	1577.09	11780

Table B.1: Summary statistics

Suicide data are from India's National Crime Records Bureau and are reported annually at the state level. Yield data are from Duflo and Pande (2007) and are reported annually at the district level, valued in constant rupees. Growing season is June-September, nongrowing season contains all other months. Precipitation is measured cumulatively. See below for details on the degree days variables.

Note that for estimation throughout the article, I use cumulative degree days, as described in Section 3.6, which sums the daily degree day values across an entire season. When reporting standardized effects in the main text, I use the within-state standard deviations in cumulative degree days. The growing season within-state standard deviation of cumulative degree days is 51 in my suicide sample and 44 in my yield sample.

Suicide data

I use annual suicide data as reported by the Indian National Crime Records Bureau (NCRB) at the state or union territory (UT) level from 1967 to 2013. States and UTs included in the data: Adaman & Nicobar Islands, Andhra Pradesh, Arunachal Pradesh, Assam, Bihar, Chandigarh, Chhattisgarh, Dadra & Nagar Haveli, Daman & Diu, Delhi, Goa, Gujarat, Haryana, Himachal Pradesh, Jammu & Kashmir, Jharkhand, Karnataka, Kerala, Lakshadweep, Madhya Pradesh, Maharashtra, Manipur, Meghalaya, Mizoram, Nagaland, Orissa, Puducherry, Punjab, Rajasthan, Sikkim, Tamil Nadu, Tripura, Uttar Pradesh, Uttaranchal, and West Bengal. I calculate suicide rates as the number of total suicides per 100,000 people, with population data linearly interpolated between Indian censuses.

Deaths in general are under-reported in India (Burgess et al., 2014), and the suicide data provided by the NCRB are particularly problematic in this regard. The data are aggregated from district police reports; because attempted suicide was a criminal offense punishable under the Indian Penal Code until 2014, there is likely to be significant under-reporting of suicide as a cause of death. As evidence of this, the NCRB reports 135,000 suicides in India in 2010, while data from a nationally-representative cause of death survey calculates the value at 187,000 (Patel et al., 2012). This under-reporting is likely uncorrelated with temperature and precipitation, implying my estimates of the response of suicide to climate provide lower bounds on the true marginal effect.

The evolution over time and space of state level suicide rates in India during my sample period is shown in Figure B.1; darker shades indicate higher suicide rates. As a point of reference, suicide rates in the United States are currently approximately 12.5 per 100,000. There is clear spatial heterogeneity, with southern India experiencing the highest suicide rates and largest increases over time. These geographic differences can be seen in more detail for a subset of states in Figure B.2. My empirical strategy accounts for this geographic heterogeneity by relying on within-state variation in order to avoid conflation of climate impacts with unobservables, such as cultural norms, political structures, and religious influences. Moreover, I account for spatially varying time trends, due to clearly distinct patterns over time across India (Figure B.2).

Agriculture data

I use agricultural data from Duflo and Pande (2007). These are district-level annual yield records for major crops (rice, wheat, sugar, sorghum, millet and maize) between 1956 and 2000, compiled from Indian Ministry of Agriculture reports and other official sources. These data cover 271 districts in 13 major agricultural states: Andhra Pradesh,

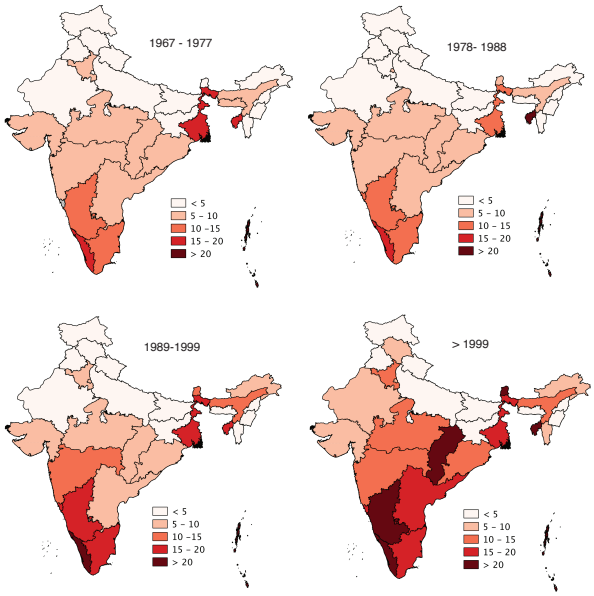

Figure B.1: Evolution of suicide rates across space and time

This figure shows states colored by the average annual suicides per 100,000 people.

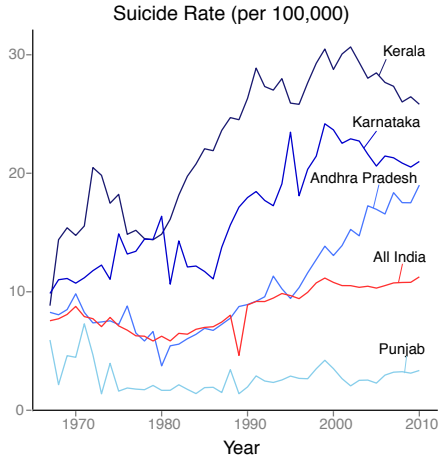

Figure B.2: Time series of suicide rates in four selected states

Suicide rates are annual suicide deaths per 100,000 people. The red line shows the average suicide rate across all of India.

Bihar, Gujarat, Haryana, Karnataka, Madhya Pradesh, Maharashtra, Orissa, Punjab, Rajasthan, Tamil Nadu, Uttar Pradesh, and West Bengal. Districts are defined by the 1961 political boundaries. As noted in Burgess et al. (2014) and Duflo and Pande (2007), these data importantly omit Kerala and Assam, two large agricultural producers that I find also have high rates of suicide. Rather than reporting quantities of each crop,

these data provide log annual yield values of a production-weighted index across all crops measured in constant Indian rupees, where prices are fixed at their 1960-1965 averages (see Duflo and Pande (2007) for details).

Climate data

Climate data are generally available at higher spatial and temporal resolution than social outcome data. Although suicides and yields are only measured annually, if the relationship between these outcomes and temperature is nonlinear, daily climate data are required, as annual average temperatures obscure such nonlinearities (Hsiang, 2016). While existing studies on temperature and suicide in the epidemiology, sociology, or meteorology literatures do not explore nonlinearities, there are two reasons why they are likely to occur. First, the growing literature on climate and interpersonal conflict reviewed by Burke, Hsiang, and Miguel (2015a) often identifies nonlinearities in the effect of temperature on violent crime. If we view suicide as a type of violence against oneself, it is possible that a similar relationship exists in this context. Second, Schlenker and Roberts (2009), among others, have identified a strongly nonlinear response of staple crop yields to temperature. If suicide in India is indeed related to agricultural productivity, then capturing this nonlinearity is important.

For daily temperature data, I use the National Center for Environmental Prediction and National Center for Atmospheric Research (NCEP/NCAR) gridded daily reanalysis product, which provides observations in an irregular grid that is approximately $1^\circ \times 1^\circ$ (Kalnay et al., 1996). These data include daily mean temperature for each grid over my entire sample period. To convert daily temperature into annual observations without losing intra-annual variability in daily weather, I use the agronomic concept of degree days. Degree days are calculated as follows, where T^* is a selected cutoff temperature value and T is a realized daily temperature value:

$$D^{T^*}(T) = \begin{cases} 0 & if \ T \leq T^* \\ T - T^* & if \ T > T^* \end{cases} \tag{B.1}$$

Because there are multiple grid cells per state, I aggregate grid-level degree day values $D^{T^*}(T)$ to state-level observations using an area-weighted average (see Table B.5 for robustness checks using weights based on population and area planted with crops). When these state-level degree day values are summed over days within a year, e.g. from day t to τ, regressing an annual outcome on cumulative degree days $\sum_{t=1}^{\tau} D^{T^*}(T_t)$ imposes a piecewise linear relationship in daily temperature, in which the outcome response has zero slope for all temperatures less than T^*. While a body of literature identifies biologically-determined cutoffs T^* for yields of variety of major crops, there is no empirical support to draw on in selecting T^* for suicides. Thus, while I use $T^* = 20^\circ$C throughout the study, I also show robustness for a range of plausible cutoffs based on the distribution of my temperature data (see Tables B.6 and B.7), and in Figure 3 of the main text I estimate a flexible piecewise linear function using four different degree day cutoffs simultaneously to impose minimal structure on the response function.

Due to the fact that reanalysis models are less reliable for precipitation data (Auffhammer et al., 2013), and because nonlinearities in precipitation that cannot be captured with a polynomial appear to be less consistently important both in the violent crime literature (Burke, Hsiang, and Miguel, 2015a) and in the agriculture literature (Schlenker and Roberts, 2009), I use the University of Delaware monthly cumulative precipitation data to complement daily temperature observations (Willmott and Matsuura, 2014). These data are gridded at a $0.5°\times0.5°$ resolution, with observations of total monthly rainfall spatially interpolated between weather stations. Again, I aggregate grids up to states and districts using area-based weights, after calculating polynomial values at the grid-level.

Regression methods

To identify the impact of temperature and precipitation on annual suicide rates, I estimate a multivariate panel regression using ordinary least squares, in which the identifying assumption is the exogeneity of within-state annual variation in degree days and cumulative precipitation. Heterogeneity in suicide rates and in temporal trajectories across states, due to an interplay between unobservable cultural, political and economic factors, implies that cross-sectional variation in climate is endogenous. Thus, I use state and year fixed effects with state-specific time trends to control for time-invariant state-level unobservables, national-level temporal shocks and regional trends.

Without precedent for the functional form of suicide's relationship to climate, my primary estimation approach employs a flexible piecewise linear specification with respect to temperature and a cubic polynomial function of cumulative precipitation. To capture the distinct impact of economically meaningful climate variation, I separately identify the temperature and precipitation response functions by agricultural seasons. My empirical model takes the general form:

$$suicide\_rate_{it} = \sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}g_s\sum_{m\in s}P_{imt} + \delta_i + \eta_t + \tau_i t + \varepsilon_{it} \qquad \text{(B.2)}$$

Where $suicide\_rate_{it}$ is the number of suicides per 100,000 people in state i in year t, $s \in \{1, 2\}$ indicates the season (growing and nongrowing), and $k = 1, ..., \kappa$ indicates a set of degree day cutoffs that constrain the piecewise linear response. In my most flexible model I let $\kappa = 7$ with degree day intervals of 5°C, and in my simplest model I let $\kappa = 2$ and estimate a standard degree day model with just one kink point and two piecewise linear segments. DD_{idt}^{k} is the degree days in bin k (e.g. degree days between 10°C and 20°C) on day d in year t in state i, and P_{imt} is cumulative precipitation during month m in year t in state i. I estimate $g(\cdot)$ as a cubic polynomial. State fixed effects δ_i account for time-invariant unobservables at the state level, while year fixed effects η_t account for India-wide time-varying unobservables. In most specifications, I include state-specific time trends $\tau_i t$ to control for differential trends in suicide driven by time-varying unobservables. My identifying assumption is that, conditional on these fixed effects and trends, variations in daily temperature and monthly rainfall are as good as

randomly assigned. Robustness to different fixed effects specifications is shown in the supplementary tables below.

Separately for each season, Equation B.2 allows me to identify $\hat{\beta}_{ks}$, the estimated change in the annual suicide rate induced by one day in bin k becoming 1°C warmer. This annual response to a daily forcing variable is similar to that estimated and described in Deryugina and Hsiang (2017). The polynomial response function for precipitation generates marginal effects of one additional millimeter of rainfall, again estimated seasonally. Due to likely correlation between errors within states, I cluster standard errors at the state level. This strategy assumes spatial correlation across states in any time period is zero, but flexibly accounts for within-state, across-time correlation.

Mechanism tests

With ideal data, I would estimate separate response functions for farmers and non-farmers to isolate the importance of an agricultural channel. Because my data do not provide the occupation of suicide victims prior to 2001 (and because using only post-2001 data at the state level leaves me statistically under-powered), I utilize a variety of other methods to investigate the validity of the oft-cited agricultural mechanism. The primary approach I take is to compare the significance and magnitude of each coefficient β_{ks} in Equation B.2 across seasons. Temperatures and rainfall in June through September have been shown to be most critical for agricultural productivity (Burgess et al., 2014), and thus should dominate the climate-suicide relationship if the agricultural channel is important. In a similar exercise, Fetzer (2014) and Blakeslee and Fishman (2017) demonstrate that monsoon-season precipitation impacts civil conflict and interpersonal crime in India, respectively, more than precipitation outside the growing season. Just as they use these findings as evidence of an agricultural channel through which climate affects crime and conflict, I use my results to identify the presence of an agricultural channel for suicide.

An additional method for examining mechanisms is to "pattern match" response functions (Burke, Hsiang, and Miguel, 2015a). For example, Hidalgo et al. (2010) show that the nonlinear relationship between agricultural income and rainfall in Brazil is nearly a perfect inverse of the relationship between land-invasion risk and rainfall. Similarly, Hsiang and Meng (2014) match the responses of conflict and income to the timing of the El Niño Southern Oscillation (ENSO), arguing that the results provide support for an income channel. I follow this approach by estimating Equation B.2 using the log value of yield as the dependent variable in place of suicide rates. My estimating equation for the yield regression is:

$$\begin{aligned} log\_yield_{ct} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\sum_{d\in s}DD_{cdt}^{k} + \sum_{s=1}^{2}g_s\sum_{m\in s}P_{cmt} \\ &+ \delta_c + \eta_t + \tau_i t + \varepsilon_{ct} \end{aligned} \tag{B.3}$$

Where the subscript c now indicates district, as my agriculture data are at the district-by-year level. δ_c are district fixed effects, η_t are year fixed effects, and $\tau_i t$ are

state-specific linear trends. Standard errors are clustered at the district level. I use the response functions uncovered in Equations B.2 and B.3 to identify matching patterns between suicide and yield.

Finally, I look for further support of economic motives by exploring spatial heterogeneity of impacts. For temperature shocks, I estimate a model that allows each of India's 32 states and union territories to have a distinct suicide rate response function:

$$\begin{aligned} suicide\_rate_{it} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa} \beta_{ks}^{i} \times \mathbb{1}[state = i] \sum_{d \in s} DD_{idt}^{k} + \sum_{s=1}^{2} g_s \sum_{m \in s} P_{imt} \\ &+ \delta_i + \eta_t + \tau_i t + \varepsilon_{it} \end{aligned} \tag{B.4}$$

I then look at correlations between these state-level temperature responses β_{ks}^{i} and analogous state-level temperature responses for log yield.

Adaptation

Figure 3 in the main text shows results from four sets of tests for adaptation. The exact specification for each regression model is shown below; all models use $\kappa = 2$ with a degree day cutoff of 20°C and include state-specific linear trends.

- **Fig. 3 A: Heterogeneity by long-run average climate**
  I calculate the average degree days over the entire period for each state in the sample, and assign each state to a tercile of high, middle or low average degree days based on the national distribution. Let $avg\_degday\_tercile_i$ indicate the tercile of state i. I estimate:

$$\begin{aligned} suicide\_rate_{it} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa} \beta_{ks} \times \sum_{d \in s} DD_{idt}^{k} + \sum_{s=1}^{2}\sum_{k=1}^{\kappa} \beta_{ks}^{2} \times \mathbb{1}[avg\_degday\_tercile_i = 2] \sum_{d \in s} DD_{idt}^{k} \\ &+ \sum_{s=1}^{2}\sum_{k=1}^{\kappa} \beta_{ks}^{3} \times \mathbb{1}[avg\_degday\_tercile_i = 3] \sum_{d \in s} DD_{idt}^{k} + \sum_{s=1}^{2} g_s \sum_{m \in s} P_{imt} \\ &+ \delta_i + \tau_i t + \varepsilon_{it} \end{aligned}$$

  Note that in this regression, the first tercile is omitted, such that coefficients β_{ks}^{2} and β_{ks}^{3} are effects for the 2^{nd} and 3^{rd} terciles, relative to the 1^{st} tercile.

- **Fig. 3 B: Heterogeneity by average income**
  I use cross-sectional gross domestic product (GDP) per capita data for each state for the year 2010 from Gennaioli et al. (2014) to assign states to terciles of the

national income distribution. I estimate:

$$\begin{aligned} suicide\_rate_{it} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\times\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}^{2}\times\mathbb{1}[avg\_income\_tercile_i=2]\sum_{d\in s}DD_{idt}^{k} \\ &+\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}^{3}\times\mathbb{1}[avg\_income\_tercile_i=3]\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}g_s\sum_{m\in s}P_{imt} \\ &+\delta_i+\tau_i t+\varepsilon_{it} \end{aligned}$$

- **Fig. 3 C: Heterogeneity by temporal subsamples**
  I estimate:

$$\begin{aligned} suicide\_rate_{it} = &\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\times\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}^{2}\times\mathbb{1}[period=1983-1997]\sum_{d\in s}DD_{idt}^{k} \\ &+\sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}^{3}\times\mathbb{1}[period=1997-2013]\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}g_s\sum_{m\in s}P_{imt} \\ &+\delta_i+\tau_s t+\varepsilon_{it} \end{aligned}$$

  Note that in this regression, the period 1967-1982 is omitted, such that coefficients β_{ks}^{2} and β_{ks}^{3} are effects relative to this earlier time period.

- **Fig. 3 D: Heterogeneity by frequency of climate variation**
  The "panel" response is estimated as follows, with $\kappa = 2$ and a degree day cutoff of 20°C:

$$suicide\_rate_{it} = \sum_{s=1}^{2}\sum_{k=1}^{\kappa}\beta_{ks}\sum_{d\in s}DD_{idt}^{k} + \sum_{s=1}^{2}g_s\sum_{m\in s}P_{imt} + \delta_i + \tau_i t + \varepsilon_{it}$$

  The "long difference" estimate is discussed below.

Long differences estimation

My main estimation strategy exploits year-to-year variation in temperature and precipitation. To test whether there are adaptive behaviors that are infeasible in response to such short-run climate shocks, but become feasible at longer time scales, I estimate a "panel of long differences" empirical model in addition to the standard panel regression, the results of which are shown in Figure 3, panel D of the main text. This strategy follows closely the approach outlined in Burke and Emerick (2016).

I first construct a moving average of the suicide rate and climate variables with a window of 5 years, over the entire sample. I then calculate the 10-year change in this average at four points in my sample: 1970, 1980, 1990 and 2000. That is, I collapse my data to 4 observations for each state in my data, where each observation measures the 10-year change in suicide rates and climate variables for each decade, and where these

changes are "smoothed" by taking 5-year averages at the end points. I then estimate the effect of changes in average degree days and precipitation on changes in average suicide rates. This model takes the following form:

$$\Delta suicide\_rate_{i\tau} = \sum_{s=1}^{2} \beta_s \Delta DD^k_{is\tau} + \sum_{s=1}^{2} \gamma_s \Delta P_{is\tau} + \delta_i + \nu_\tau + \varepsilon_{i\tau} \tag{B.5}$$

Where δ_i are state fixed effects, ν_τ are fixed effect for each of the four decadal starting points in my sample, s indicates the growing and nongrowing seasons, k indicates the degree day cutoff, and Δ indicates the 10-year change in each variable. I report results both including and excluding the state fixed effect δ_i and the decadal starting point fixed effect ν_τ. Results are shown in Table B.9.

Impacts of recent climate trends

To compute estimates of the effect of warming temperature trends since 1980, I follow the approach outlined in Carleton and Hsiang (2016) and Lobell, Schlenker, and Costa-Roberts (2011). I do not consider trends in precipitation, as my estimates for suicide impacts of precipitation were highly uncertain. Nor do I consider the impacts of warming outside the growing season.

For each state in my data, I estimate a linear trend in growing season degree days above 20°C for the years 1980-2013. Let the predicted value of degree days in state i in year t, as estimated by the trend, be indicated by DD^*_{it}. I then create a de-trended degree days residual that is normalized to temperature in 1980, for every state-year (see Figure B.3):

$$DD\_detrended_{it} = DD_{it} - DD^*_{it} + DD^*_{i,1980}$$

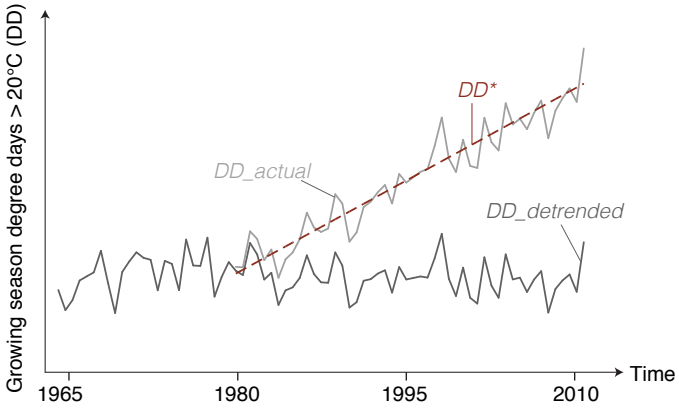

Figure B.3: Identifying counterfactual de-trended temperatures

This figure is adapted from Carleton and Hsiang (2016).

I predict suicide rates using actual and de-trended growing season degree days, using coefficient estimates from the model in Table 1 of the main text which includes both state trends and year fixed effects (column 3). The elevated risk of suicide attributable to the trend, relative to the de-trended counterfactual, is the difference between these two predictions, which simplifies to $\Delta s_{it} = \hat{\beta} \times (DD_{it} - DD\_detrended_{it})$, where $\hat{\beta} = 0.008$ as estimated in my preferred empirical model. I multiply Δs_{it} — the increase in the suicide rate attributable to warming — by the population in each state and each year to recover the additional number of suicide deaths. Figure 4 *B* in the main text displays these additional deaths in each year; integrating over states and years gives the cumulative effect of temperature trends for all of India over the entire period since 1980.

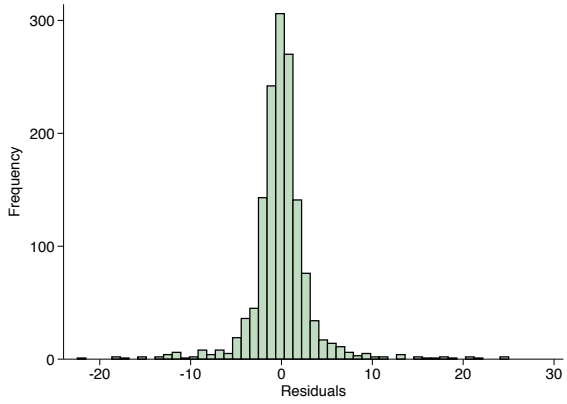

Figure B.4: Distribution of residuals

Histogram shows the distribution of residuals from regression model in Equation B.2.

Assumptions behind ordinary least squares

Throughout this article, I estimate the effect of climate on suicide using ordinary least squares (OLS). It is also possible to model suicide events using nonlinear count models, such as Poisson regression or negative binomial regression, and these approaches may be preferable to OLS when the conditional distribution of the dependent variable is poorly approximated by a normal distribution. While other analyses have modeled causes of suicide using count models (e.g. Haukka et al. (2008)), I use OLS for two reasons: the relative weakness of assumptions required for consistent estimation of causal effects, and its ease of interpretation.rupees

Distributional assumptions on either the disturbances or the outcome variable, such as normality, are not required in order for OLS regression coefficients to consistently estimate a true population parameter (Wooldridge, 2010). However, normality of the disturbances is an assumption used to estimate critical values for inference in finite

samples. As count data are aggregated to coarser levels of spatial and temporal aggregation, it becomes more likely that the conditional distribution of the outcome variable approximates a normal distribution. In my case, my state-by-year observations are relatively coarse measures. Reassuringly, the residuals from my main regression model very closely approximate a normal distribution, as shown in Figure B.4.

In contrast, the assumptions imposed by count models can be much more restrictive. For example, modeling the data generating process as Poisson imposes the restriction that the mean and variance of the outcome variable are identical (as shown in Table B.1, this is not the case in my data). Moreover, the coefficients derived from count models are much more difficult to interpret than those derived from OLS. I therefore follow the literature on climate and mortality (e.g. Deschênes and Greenstone (2011)), as well as the literature studying the socioeconomic drivers of suicide in aggregate panel data (e.g. Andres (2005)), and use OLS with the state-by-year suicide rate as an outcome variable.

B.2 Supplementary tables

	Suicides per 100,000			*100×Log yield (rupees/ha)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Growing seas. degree days below threshold (°C)	0.003***	0.000	0.004***	0.013	-0.019	-0.003
	(0.001)	(0.001)	(0.001)	(0.009)	(0.018)	(0.013)
Growing seas. degree days (°C)	0.007***	0.009**	0.008**	-0.017***	-0.020*	-0.019*
	(0.002)	(0.004)	(0.003)	(0.006)	(0.010)	(0.010)
Nongrowing seas. degree days below threshold (°C)	-0.001	-0.009*	-0.003*	0.002	0.007	0.001
	(0.001)	(0.004)	(0.002)	(0.003)	(0.005)	(0.004)
Nongrowing seas. degree days (°C)	-0.002*	0.002	0.001	0.010***	0.018***	0.010*
	(0.001)	(0.003)	(0.003)	(0.004)	(0.006)	(0.006)
Growing seas. precip. (cm)	0.115	0.251	0.183	6.048***	4.255***	4.422***
	(0.147)	(0.176)	(0.152)	(0.712)	(0.638)	(0.653)
Growing seas. precip.^{2} (cm^2)	-0.026	-0.027	-0.022	-0.870***	-0.653***	-0.681***
	(0.020)	(0.019)	(0.017)	(0.133)	(0.117)	(0.122)
Growing seas. precip.^{3} (cm^3)	0.001*	0.000*	0.000	0.030***	0.023***	0.024***
	(0.000)	(0.000)	(0.000)	(0.005)	(0.004)	(0.004)
Nongrowing seas. precip. (cm)	-0.031	0.097	0.033	1.496**	2.974***	3.103***
	(0.206)	(0.221)	(0.237)	(0.596)	(0.775)	(0.746)
Nongrowing seas. precip.^{2} (cm^2)	0.013	0.003	0.014	0.282	-0.300	-0.365
	(0.030)	(0.035)	(0.033)	(0.323)	(0.401)	(0.377)
Nongrowing seas. precip.^{3} (cm^3)	-0.001	-0.001	-0.001	-0.071*	-0.018	-0.011
	(0.001)	(0.001)	(0.001)	(0.037)	(0.050)	(0.045)
Observations	1,434	1,434	1,434	11,289	11,289	11,289
R-squared	0.908	0.893	0.916	0.840	0.842	0.849
State Trends	YES		YES	YES		YES
Year FE		YES	YES		YES	YES

Table B.2: Effect of heat exposure and precipitation on suicide rates and yield values, by agricultural season

Temperature coefficients represent the effect of one day becoming 1°C warmer on the annual suicide rate (suicide deaths per 100,000 people) or annual yield (log rupees/ha), where temperature effects are differentially estimated for days below 20°C and above 20°C and for the growing and nongrowing seasons in India. Precipitation coefficients represent the effect of seasonal cumulative rainfall increasing by 1cm on the annual suicide rate and annual yield. Columns (1)–(3) include state fixed effects and report standard errors clustered at the state level. Columns (4)–(6) include district fixed effects and report standard errors clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000			
	(1) OLS	(2) State FE	(3) State FE + State Trends	(4) State & Yr FE + State Trends
Growing seas. degree days below threshold (°C)	0.0046	-0.0040***	0.0026***	0.0037***
	(0.0053)	(0.0010)	(0.0008)	(0.0008)
Growing seas. degree days (°C)	-0.0020	0.0175***	0.0066***	0.0079**
	(0.0040)	(0.0035)	(0.0023)	(0.0031)
Nongrowing seas. degree days below threshold (°C)	-0.0038	-0.0036	-0.0009	-0.0027*
	(0.0025)	(0.0040)	(0.0010)	(0.0016)
Nongrowing seas. degree days (°C)	0.0159***	0.0029	-0.0020*	0.0014
	(0.0040)	(0.0026)	(0.0011)	(0.0026)
Growing seas. precip. (cm)	0.1083	0.3407*	0.1150	0.1826
	(0.8147)	(0.1703)	(0.1465)	(0.1522)
Growing seas. precip.^{2} (cm^2)	-0.0071	-0.0510**	-0.0264	-0.0218
	(0.0789)	(0.0216)	(0.0196)	(0.0171)
Growing seas. precip.^{3} (cm^3)	-0.0001	0.0009***	0.0006*	0.0004
	(0.0013)	(0.0003)	(0.0003)	(0.0002)
Nongrowing seas. precip. (cm)	1.1920*	0.0673	-0.0312	0.0327
	(0.6845)	(0.2126)	(0.2060)	(0.2367)
Nongrowing seas. precip.^{2} (cm^2)	0.0406	-0.0000	0.0133	0.0142
	(0.1395)	(0.0326)	(0.0300)	(0.0330)
Nongrowing seas. precip.^{3} (cm^3)	-0.0065	-0.0005	-0.0009	-0.0009
	(0.0066)	(0.0013)	(0.0011)	(0.0011)
Observations	1,434	1,434	1,434	1,434
R-squared	0.478	0.871	0.908	0.916
State FE		YES	YES	YES
State Trends			YES	YES
Year FE				YES

Table B.3: Robustness of the suicide degree day model to various fixed effects specifications

Regression includes annual data for 32 Indian states between 1967 and 2013. Growing season is June-September, nongrowing season contains all other months. Degree day cutoff is 20°C. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	Log yield (rupees per ha)			
	(1) OLS	(2) District FE	(3) District FE + State Trends	(4) District & Yr FE + State Trends
Growing seas. degree days below threshold (°C)	0.2557***	-0.0116	0.0133	-0.0027
	(0.0377)	(0.0348)	(0.0090)	(0.0127)
Growing seas. degree days (°C)	-0.0376**	0.1453***	-0.0173***	-0.0191*
	(0.0149)	(0.0114)	(0.0060)	(0.0097)
Nongrowing seas. degree days below threshold (°C)	-0.0425***	-0.0784***	0.0019	0.0012
	(0.0107)	(0.0056)	(0.0034)	(0.0044)
Nongrowing seas. degree days (°C)	0.0007	0.0020	0.0101***	0.0100*
	(0.0124)	(0.0075)	(0.0038)	(0.0056)
Growing seas. precip. (cm)	0.6421	8.8834***	6.0476***	4.4222***
	(2.6620)	(0.8961)	(0.7121)	(0.6529)
Growing seas. precip.^{2} (cm^2)	-0.1911	-1.4395***	-0.8697***	-0.6812***
	(0.4707)	(0.1730)	(0.1332)	(0.1222)
Growing season precip.^{3} (cm^3)	0.0155	0.0483***	0.0296***	0.0238***
	(0.0181)	(0.0065)	(0.0048)	(0.0044)
Nongrowing seas. precip. (cm)	26.9092***	14.8560***	1.4962**	3.1032***
	(2.3122)	(1.3260)	(0.5958)	(0.7457)
Nongrowing seas. precip.^{2} (cm^2)	-7.6716***	-6.4903***	0.2816	-0.3646
	(1.3907)	(1.0047)	(0.3230)	(0.3765)
Nongrowing season precip.^{3} (cm^3)	0.6476***	0.5919***	-0.0710*	-0.0114
	(0.1599)	(0.1419)	(0.0372)	(0.0450)
Observations	11,289	11,289	11,289	11,289
R-squared	0.225	0.666	0.839	0.848
District FE		YES	YES	YES
Year FE				YES
State Trends			YES	YES

Table B.4: Robustness of the yield degree day model to various fixed effects specifications

Regression includes annual district-level data for 13 Indian states between 1956 and 2000. Growing season is June-September, Nongrowing season contains all other months. Degree day cutoff is 20°C. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000			*100×Log yield (rupees/ha)*		
	(1) area weighted	(2) crop weighted	(3) pop weighted	(4) area weighted	(5) crop weighted	(6) pop weighted
Growing seas. degree days below threshold (°C)	0.0026***	0.0032***	0.0032***	0.0133	0.0180*	0.0177*
	(0.0008)	(0.0004)	(0.0005)	(0.0090)	(0.0097)	(0.0099)
Growing seas. degree days (°C)	0.0066***	0.0050***	0.0058***	-0.0173***	-0.0171***	-0.0170***
	(0.0023)	(0.0016)	(0.0019)	(0.0060)	(0.0060)	(0.0060)
Nongrowing seas. degree days below threshold (°C)	-0.0009	-0.0010	-0.0009	0.0019	0.0014	0.0011
	(0.0010)	(0.0009)	(0.0009)	(0.0034)	(0.0040)	(0.0040)
Nongrowing seas. degree days (°C)	-0.0020*	-0.0021**	-0.0021**	0.0101***	0.0103***	0.0104***
	(0.0011)	(0.0001)	(0.0001)	(0.0038)	(0.0038)	(0.0038)
Growing Seas. precip (cm)	0.1150	0.0159	0.1689	6.0476***	6.0430***	6.0390***
	(0.1465)	(0.0625)	(0.1722)	(0.7121)	(0.7154)	(0.7153)
Growing seas. precip.^{2} (cm^2)	-0.0264	-0.0127	-0.0340	-0.8697***	-0.8687***	-0.8680***
	(0.0196)	(0.0076)	(0.0240)	(0.1332)	(0.1336)	(0.1336)
Growing seas. precip.^{3} (cm^3)	0.0006*	0.0005**	0.0007*	0.0296***	0.0296***	0.0295***
	(0.0003)	(0.0002)	(0.0004)	(0.0048)	(0.0049)	(0.0049)
Nongrowing seas. precip (cm)	-0.0312	-0.2630**	-0.0460	1.4962**	1.4964**	1.4982**
	(0.2060)	(0.1222)	(0.2160)	(0.5958)	(0.5982)	(0.5980)
Nongrowing seas. precip.^{2} (cm^2)	0.0133	0.0498	0.0195	0.2816	0.2829	0.2818
	(0.0300)	(0.0300)	(0.0319)	(0.3230)	(0.3238)	(0.3238)
Nongrowing seas. precip.^{3} (cm^3)	-0.0009	-0.0025*	-0.0013	-0.0710*	-0.0712*	-0.0710*
	(0.0011)	(0.0014)	(0.0014)	(0.0372)	(0.0372)	(0.0372)
Observations	1,434	1,387	1,434	11,289	11,289	11,289
R-squared	0.9083	0.9141	0.9086	0.8401	0.8401	0.8401

Table B.5: Robustness of the degree day model to different weighting schemes for aggregation of climate data

Regressions in columns (1)–(3) include annual data for 32 Indian states between 1967 and 2013, and in columns (4)–(6) include annual data for all districts in 13 Indian states between 1956 and 2000. Growing season is June-September, nongrowing season contains all other months. Degree day cutoff is 20°C. Columns (1)–(3) include state fixed effects and report standard errors clustered at the state level. Columns (4)–(6) include district fixed effects and report standard errors clustered at the district level. All regressions include state-specific linear trends. Crop weights for each grid cell are cropped area fraction; population weights for each grid cell are total population in 2010. *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000		
	(1)	(2)	(3)
	15°C	20°C	25°C
Growing seas. degree days below threshold (°C)	0.0026***	0.0026***	0.0024**
	(0.0009)	(0.0008)	(0.0009)
Growing seas. degree days (°C)	0.0062***	0.0066***	0.0087**
	(0.0021)	(0.0023)	(0.0033)
Nongrowing seas. degree days below threshold (°C)	-0.0009	-0.0009	-0.0009
	(0.0013)	(0.0010)	(0.0010)
Nongrowing seas. degree days (°C)	-0.0017	-0.0020*	-0.0027
	(0.0010)	(0.0011)	(0.0022)
Growing seas. precip (cm)	0.1079	0.1150	0.1376
	(0.1450)	(0.1465)	(0.1511)
Growing seas. precip.^{2} (cm^2)	-0.0258	-0.0264	-0.0283
	(0.0195)	(0.0196)	(0.0200)
Growing seas. precip.^{3} (cm^3)	0.0006*	0.0006*	0.0006*
	(0.0003)	(0.0003)	(0.0003)
Nongrowing seas. precip (cm)	-0.0293	-0.0312	-0.0303
	(0.2059)	(0.2060)	(0.2085)
Nongrowing seas. precip.^{2} (cm^2)	0.0132	0.0133	0.0135
	(0.0301)	(0.0300)	(0.0305)
Nongrowing seas. precip.^{3} (cm^2)	-0.0009	-0.0009	-0.0009
	(0.0011)	(0.0011)	(0.0011)
Observations	1,434	1,434	1,434
R-squared	0.9083	0.9083	0.9084

Table B.6: Robustness of the suicide degree day model to different degree day cutoffs

Regressions include annual data for 32 Indian states between 1967 and 2013. Growing season is June-September, nongrowing season contains all other months. All regressions include state-specific linear time trends. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	100 × Log yield (rupees per ha)		
	(1) 15°C	(2) 20°C	(3) 25°C
Growing seas. degree days below threshold (°C)	0.0113	0.0133	0.0272***
	(0.0093)	(0.0090)	(0.0082)
Growing seas. degree days (°C)	-0.0164***	-0.0173***	-0.0270***
	(0.0058)	(0.0060)	(0.0078)
Nongrowing seas. degree days below threshold (°C)	0.0090*	0.0019	-0.0029
	(0.0054)	(0.0034)	(0.0027)
Nongrowing seas. degree days (°C)	0.0052*	0.0101***	0.0251***
	(0.0029)	(0.0038)	(0.0058)
Growing seas. precip (cm)	6.1603***	6.0476***	5.6173***
	(0.7099)	(0.7121)	(0.6877)
Growing seas. precip.^{2} (cm^2)	-0.8844***	-0.8697***	-0.8136***
	(0.1327)	(0.1332)	(0.1297)
Growing Season precip.^{3} (cm^3)	0.0300***	0.0296***	0.0279***
	(0.0048)	(0.0048)	(0.0047)
Nongrowing seas. precip (cm)	1.4283**	1.4962**	1.6163***
	(0.5975)	(0.5958)	(0.5969)
Nongrowing seas. precip.^{2} (cm^2)	0.3065	0.2816	0.1888
	(0.3244)	(0.3230)	(0.3204)
Nongrowing seas. precip ^{3} (cm^3)	-0.0736**	-0.0710*	-0.0598
	(0.0371)	(0.0372)	(0.0371)
Observations	11,289	11,289	11,289
R-squared	0.8387	0.8388	0.8395

Table B.7: Robustness of the yield degree day model to different degree day cutoffs

Regressions include annual data for all districts in 13 Indian states between 1956 and 2000. Growing season is June-September, nongrowing season contains all other months. All regressions include state-specific linear time trends. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Linear trends	(2) Quad. trends	(3) Linear trends + Year FE	(4) Quad. trends + Year FE
Growing seas. degree days below threshold (°C)	0.0026***	0.0035***	0.0037***	0.0046***
	(0.0008)	(0.0011)	(0.0008)	(0.0008)
Growing seas. degree days (°C)	0.0066***	0.0064**	0.0079**	0.0082**
	(0.0023)	(0.0024)	(0.0031)	(0.0031)
Nongrowing seas. degree days below threshold (°C)	-0.0009	-0.0009	-0.0027*	-0.0018
	(0.0010)	(0.0010)	(0.0016)	(0.0013)
Nongrowing seas. degree days (°C)	-0.0020*	-0.0015	0.0014	0.0024
	(0.0011)	(0.0011)	(0.0026)	(0.0025)
Growing seas. precip. (cm)	0.1150	0.1030	0.1826	0.2166
	(0.1465)	(0.1447)	(0.1522)	(0.1502)
Growing seas. precip.^{2} (cm^2)	-0.0264	-0.0268	-0.0218	-0.0253
	(0.0196)	(0.0192)	(0.0171)	(0.0172)
Growing seas. precip.^{3} (cm^3)	0.0006*	0.0006*	0.0004	0.0005*
	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Nongrowing seas. precip. (cm)	-0.0312	0.0125	0.0327	0.0421
	(0.2060)	(0.1994)	(0.2367)	(0.2426)
Nongrowing seas. precip.^{2} (cm^2)	0.0133	0.0049	0.0142	0.0152
	(0.0300)	(0.0299)	(0.0330)	(0.0343)
Nongrowing seas. precip.^{3} (cm^2)	-0.0009	-0.0006	-0.0009	-0.0009
	(0.0011)	(0.0011)	(0.0011)	(0.0012)
Observations	1,434	1,434	1,434	1,434
R-squared	0.9083	0.9091	0.9163	0.9173
Linear State Trends	YES		YES	
Quad. State Trends		YES		YES
Year FE			YES	YES

Table B.8: Robustness of the suicide degree day model to different time-varying controls

Regressions include annual data for 32 Indian states between 1967 and 2013. Growing season is June-September, nongrowing season contains all other months. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Deg. days 20°C	(2) Deg. days 20°C	(3) Deg. days 25°C	(4) Deg. days 25°C
Growing seas. degree days (°C)	0.023	0.020	0.037	0.023
	(0.021)	(0.020)	(0.031)	(0.026)
Nongrowing Seas. degree days (°C)	-0.012	-0.004	-0.013	0.008
	(0.012)	(0.011)	(0.016)	(0.017)
Growing seas. precip (cm)	-0.844**	-0.731	-0.829**	-0.723*
	(0.340)	(0.438)	(0.313)	(0.418)
Nongrowing seas. precip (cm)	-0.378	-0.020	-0.360	-0.027
	(0.452)	(0.472)	(0.417)	(0.456)
Observations	116	116	116	116
R-squared	0.408	0.479	0.408	0.478
State FE	YES	YES	YES	YES
Time Period FE		YES		YES

Table B.9: Panel of long differences

Dependent variable in all regressions is the decadal difference in the smoothed suicide rate, where the data are organized as a 4-period panel of 10-year differences. Periods are 1970-1980, 1980-1990, 1990-2000 and 2000-2010. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Baseline Model	(2) Irrigation & Temp	(3) Irrigation & Precip
Growing seas. degree days (°C)	0.0056***	0.0056**	0.0049**
	(0.0020)	(0.0028)	(0.0021)
Nongrowing seas. degree days (°C)	-0.0016	-0.0021*	-0.0021*
	(0.0011)	(0.0012)	(0.0012)
Growing seas. precip (cm)	-0.0789	-0.0754	-0.0741
	(0.0475)	(0.0508)	(0.0555)
Nongrowing seas. precip (cm)	-0.0082	-0.0341	-0.0343
	(0.0819)	(0.0998)	(0.1003)
Irrigated		-21.8881***	-23.5120***
		(3.6741)	(2.3590)
Irrigated × growing season degree days (°C)		-0.0026	
		(0.0034)	
Irrigated × growing season precip. (cm)			-0.0105
			(0.0729)
Observations	1,434	1,332	1,332
R-squared	0.908	0.907	0.907

Table B.10: Heterogeneity in the degree days model by irrigation prevalence

Regressions include annual data for 32 Indian states between 1967 and 2013. Growing season is June-September, nongrowing season contains all other months. Degree day cutoff is 20°C. The variable "irrigated" is a time-invariant dummy classifying each state as irrigated or not. For this classification, I use Indian Ministry of Agriculture data to classify states as irrigated if their share of cropped area that is under irrigation ever exceeds 50% during my sample period. All regressions include state-specific linear time trends. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000		*100×Log yield (rupees/ha)*	
	(1)	(2)	(3)	(4)
	June-Sep. season	State-specific season	June-Sep. season	State-specific season
Growing seas. degree days (°C)	0.0066***	0.0072**	-0.0173***	-0.0177***
	(0.0023)	(0.0032)	(0.0060)	(0.0060)
Nongrowing seas. degree days (°C)	-0.0020*	-0.0023*	0.0101***	0.0100***
	(0.0011)	(0.0012)	(0.0038)	(0.0037)
Growing seas. precip. (cm)	0.1150	0.2327	6.0476***	6.1183***
	(0.1465)	(0.2391)	(0.7121)	(0.6711)
Growing seas. precip.^{2} (cm^2)	-0.0264	-0.0355	-0.8697***	-0.8810***
	(0.0196)	(0.0261)	(0.1332)	(0.1265)
Growing seas. precip.^{3} (cm^3)	0.0006*	0.0007	0.0296***	0.0300***
	(0.0003)	(0.0004)	(0.0048)	(0.0046)
Nongrowing seas. precip. (cm)	-0.0312	-0.1927	1.4962**	1.6372**
	(0.2060)	(0.1697)	(0.5958)	(0.7688)
Nongrowing seas. precip.^{2} (cm^2)	0.0133	0.0296	0.2816	-0.0189
	(0.0300)	(0.0333)	(0.3230)	(0.5408)
Nongrowing seas. precip.^{3} (cm^3)	-0.0009	-0.0010	-0.0710*	-0.0389
	(0.0011)	(0.0011)	(0.0372)	(0.0821)
Observations	1,434	1,434	11,289	11,289
R-squared	0.908	0.908	0.840	0.840

Table B.11: Robustness of the degree day model to state-specific growing season definitions

In the main text, I define the growing season to be the months of June through September, based on the average arrival and withdrawal dates of the southwest monsoon. However, the monsoon arrives and withdraws differentially throughout India, first arriving in the southeast in late May, and reaching the northwest of the country by the middle of July. Withdrawal occurs in reverse, with the rains first ending in the northeast in early September, but continuing in the southeast until December. Because my approximation of this timing is coarse, in this table I demonstrate robustness of my main results to an alternative definition of the growing season, in which each state is described by a state-specific growing season, the dates of which are obtained from the Indian Meteorological Department. Regressions in columns (1)–(2) include annual data for 32 Indian states between 1967 and 2013, and in columns (3)–(4) include annual data for all districts in 13 Indian states between 1956 and 2000. The growing season is defined as June-September in columns (1) and (3), and is defined individually by state using data from the India Meteorological Department on average monsoon arrival and withdrawal dates in columns (2) and (4). The nongrowing season contains all other months. The degree day cutoff is 20°C, and all regressions include state-specific linear time trends. Standard errors are clustered at the state level in columns (1)–(2) and at the district level in columns (3)–(4). *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000		
	(1)	(2)	(3)
Growing seas. degree days (°C)	0.0162***	0.0059***	0.0073**
	(0.0036)	(0.0021)	(0.0031)
Nongrowing seas. degree days (°C)	0.0033	-0.0014	0.0013
	(0.0025)	(0.0011)	(0.0024)
Drought: 20th percentile	-0.3160	0.0417	-0.0439
	(0.3940)	(0.3430)	(0.4440)
Surplus: 80th percentile	-0.3770	-0.3740	0.2940
	(0.2750)	(0.2330)	(0.3370)
Observations	1,472	1,472	1,472
R-squared	0.869	0.908	0.916
State Trends		YES	YES
Year FE			YES

Table B.12: Effect of drought and excessive rainfall on suicide

Regressions include annual data for 32 Indian states between 1967 and 2013. Temperature coefficients represent the effect of one day becoming 1°C warmer on the annual suicide rate, for days above 20°C in the growing and nongrowing seasons in India. Drought is defined as an indicator equal to one when annual rainfall is in the 20th percentile or below, while surplus is equal to one when annual rainfall is above the 80th percentile, where percentiles are state-specific. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

	Suicides per 100,000	
	(1) Main model	(2) AR model
Lagged suicide rate (suicides/100,000)		0.3195***
		(0.0740)
Growing seas. degree days below threshold (°C)	0.0037***	0.0036***
	(0.0008)	(0.0009)
Growing seas. degree days (°C)	0.0079**	0.0067*
	(0.0031)	(0.0033)
Growing seas. precip. (cm)	0.1826	0.1031
	(0.1522)	(0.1849)
Growing seas. precip.^{2} (cm^2)	-0.0218	-0.0151
	(0.0171)	(0.0215)
Growing seas. precip.^{3} (cm^3)	0.0004	0.0004
	(0.0002)	(0.0003)
Observations	1,434	1,400
R-squared	0.9163	0.9270

Table B.13: Robustness of the suicide degree day model to inclusion of a lagged dependent variable

Regressions include annual data for 32 Indian states between 1967 and 2013. Growing season is June-September, nongrowing season contains all other months. All regressions include linear state-specific trends and year fixed effects. Standard errors are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1.

B.3 Supplementary figures

Degree days across seasons

One concern with using the pattern matching approach I show in Figure 3.1 is that temperature may impact suicide during the growing season months only, but for reasons unrelated to agriculture. In particular, there is some evidence that suicide is directly impacted by heat through a psychological mechanism (Deisenhammer, Kemmler, and Parson, 2003). However, this direct impact is not identifiable in the nongrowing season, despite the presence of many hot days during this period (Figure B.5). Across a variety of robustness checks (Tables B.3-B.13), coefficients on high temperatures in the off-season are consistently close to zero and statistically insignificant, suggesting no strong psychological mechanism is at play.

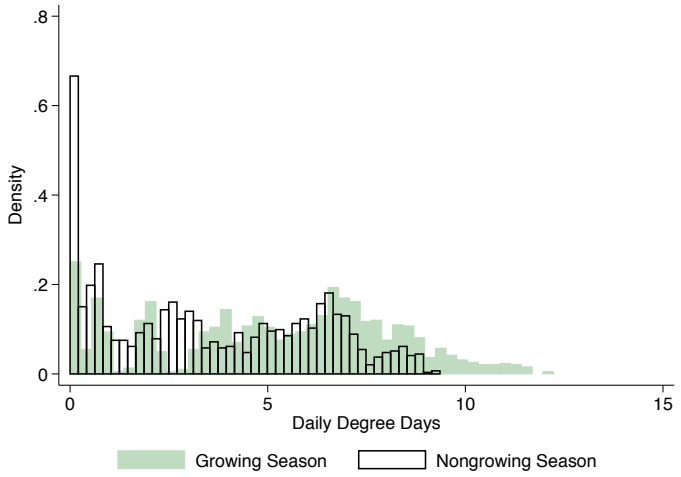

Figure B.5: Distribution of cumulative degree days above 20°C in the growing and nongrowing seasons

This figure shows the distribution of daily degree days above 20°C for the growing and nongrowing seasons, using daily mean temperature for 32 of India's states between 1967-2013. The growing season is June through September, while the nongrowing season is all other months.

Robustness of piecewise linear response

Figure B.6 shows the robustness of my piecewise linear estimation strategy for temperature to a higher order of flexibility. The dotted lines show the response function when estimating Equation 3.2 and setting $\kappa = 7$, while the solid lines, as in the main text, show the response function when setting $\kappa = 4$. Temperatures below 10°C are not shown, although are included in the regression.

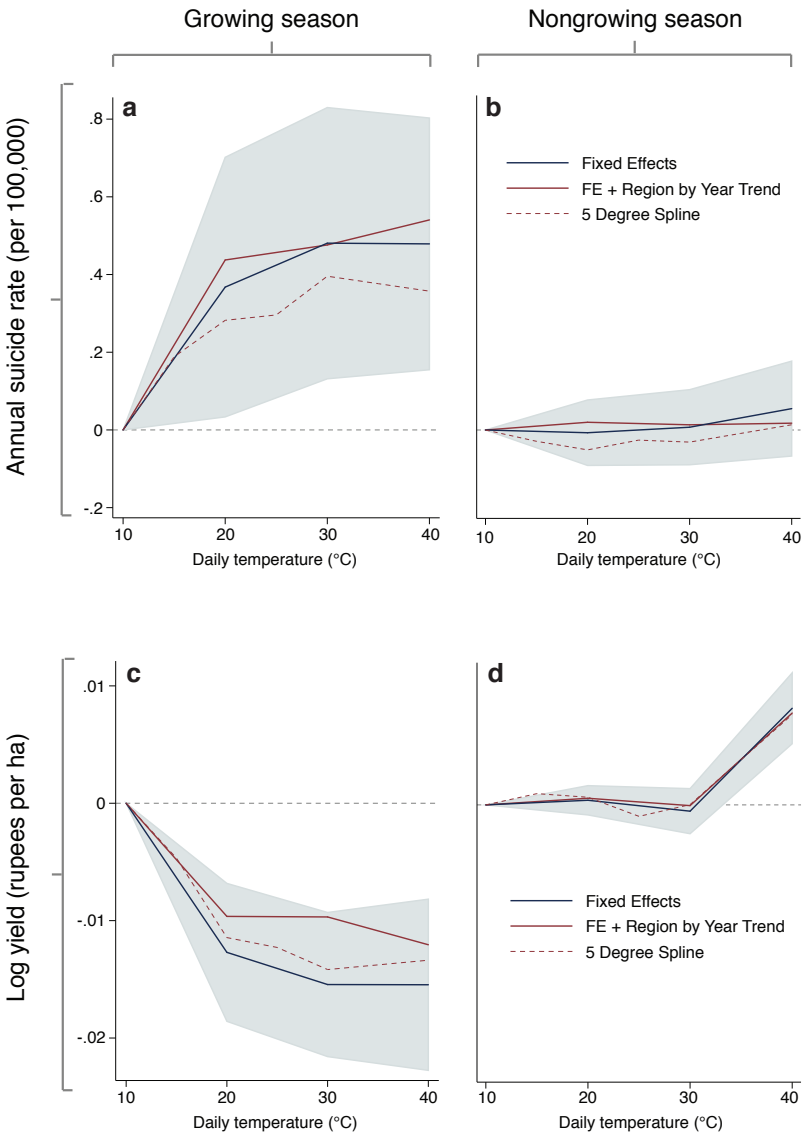

Figure B.6: Nonlinear relationship between temperature and suicide rates, and between temperature and yield

a and **b** show the response of annual suicides per 100,000 people to growing season (June through September) and nongrowing season (all other months) temperatures, respectively. Panels **c** and **d** show the response of annual log yield, valued in rupees per hectare, to growing season (June through September) and nongrowing season (all other months) temperatures, respectively. The fixed effects regression includes year fixed effects, while the FE + Region by Year Trend regression includes year fixed effects and linear regional time trends. The "5 Degree Spline" model estimates a linear spline with knots at every 5°C interval. All graphs are centered at zero.

Monthly estimation of temperature and precipitation effects

With my main specification, I am unable to reject that rainfall has no effect on suicide rates. This result may be due to my need to characterize monsoon rainfall at the

state level, as there can be important within-state differences in monsoon arrival and withdrawal (Burgess et al., 2014). The higher-resolution district-level agricultural data, in contrast, suffer far less from this problem. Figure B.7 suggests measurement error may be at play: this plot of *monthly* rainfall effects illustrates a consistently negative, yet often insignificant, impact of rainfall on suicide rates during the main growing season months.

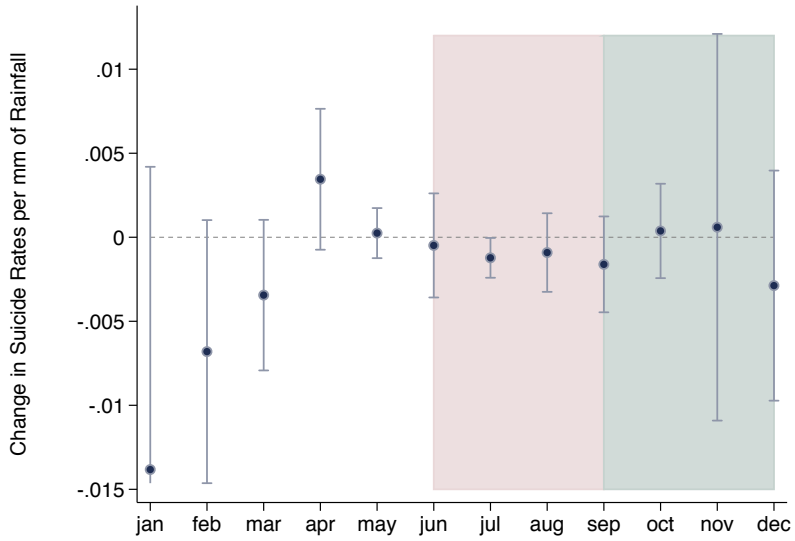

Figure B.7: Within-season effects of precipitation on suicide rates

Points represent the marginal effect of one mm of rainfall in each month on suicide deaths per 100,000 people. Shaded pink areas represent the growing season months and shaded green areas represent the harvesting season months, although some states continue to grow crops through October and November.

Lagged effects

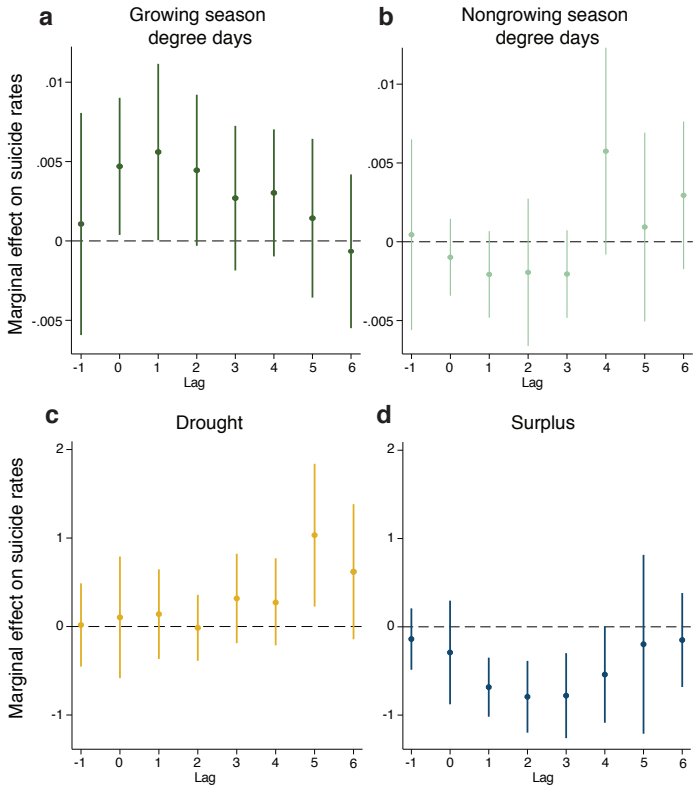

Figure B.8: Lagged effects of temperature and rainfall on suicide rates

Points represent the marginal effect of degree days (panels **a** and **b**), an indicator for drought (panel **c**), or an indicator for surplus rainfall (panel **d**) on the annual number of suicides per 100,000 people. The x-axis corresponds to the number of annual lags. All coefficients shown in panels **a** and **b** were estimated jointly in a degree days model with a cutoff of 20°C and a cubic polynomial in precipitation; all coefficients shown in panels **c** and **d** were estimated jointly in a degree days model with a cutoff of 20°C where indicators for drought (annual rainfall below state-specific 20^{th} percentile) and surplus (annual rainfall above state-specific 80^{th} percentile) were used instead of continuous rainfall. Standard errors are clustered at the state level, and 95% confidence intervals are shown around each coefficient.

Appendix C

Appendix: Global mortality consequences of climate change

C.1 Adaptation costs: Empirical implementation

In Section 4.2, we outline a conceptual framework that derives a cost function for the costs of adaptation under climate change. This cost function, while depending on many unobservable choice variables available to individual agents, is shown to be recoverable using only observable parameters. The key to this derivation is that the response of observable outcomes – such as the mortality rate – to characteristics of the long-run climate – such as long-run average temperature – is a sufficient statistic for the beneficial effects of all compensatory investments. This is shown in the main text in Equation 4.11. The final expression for total change in adaptation costs due to a change in the climate from climate $\boldsymbol{C}_0$ to climate $\boldsymbol{C}_1$ is:

$$\boldsymbol{A}(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) - \boldsymbol{A}(\boldsymbol{b}^*(\boldsymbol{C}_0, Y)) = -\int_{\boldsymbol{C}_0}^{\boldsymbol{C}_1} VSL \frac{\partial \tilde{f}(\boldsymbol{b}^*(\boldsymbol{C}, Y))}{\partial \boldsymbol{C}} d\boldsymbol{C} \tag{C.1}$$

Recall that our empirical model of $f(\cdot)$ is (Equation 4.15 in the main text):

$$\begin{aligned} M_{aict} = & \sum_{k \in K} \gamma_{0,a}^k T_{it}^k \\ & + \sum_{k \in K} \gamma_{1,a}^k T_{it}^k \times TMEAN_s \\ & + \sum_{k \in K} \gamma_{2,a}^k T_{it}^k \times \log(GDPpc)_s \\ & + g^c(P_{it}) + \alpha_{ai} + \delta_{act} + \varepsilon_{ait} \end{aligned} \tag{C.2}$$

with s referring to ADM1-level (e.g., state or province), and $k \in K$ indicates the term in the nonlinear function of temperature. M_{iact} is the mortality rate for age group a in county i in country c in year t, and $TMEAN$ is our approximation of the climate, $\boldsymbol{C}$.

We estimate two objects related to Equation C.1, based on our empirical model in Equation C.2. The first, which we call an "upper bound" on adaptation costs, corresponds directly to Equation C.1. It is an upper bound on the costs incurred by populations to adapt to a climate that warms from an initial climate of $\boldsymbol{C}(t_0)$ and by the end of the 21^{st} century endures climate $\boldsymbol{C}(t_1)$. As discussed in the main text, treating this bound as the total costs assumes all agents in all climates face a common adaptation cost function $A(b)$, such that the optimal $b*$ chosen is identical for two locations facing the same climate. We approximate this integral as:

$$\begin{aligned} A(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) - \boldsymbol{A}(\boldsymbol{b}^*(\boldsymbol{C}_0, Y)) &= -\int_{\boldsymbol{C}(t_0)}^{\boldsymbol{C}(t_1)} \frac{\partial \tilde{f}}{\partial \boldsymbol{C}} d\boldsymbol{C} \\ &\approx -\sum_{\boldsymbol{C}=\boldsymbol{C}(t_0)}^{\boldsymbol{C}(t_1)} \frac{\partial \tilde{f}}{\partial \boldsymbol{C}} (\boldsymbol{C}(t) - \boldsymbol{C}(t-1)) \\ &\approx -\sum_{t=t_0}^{t_1} \sum_{k \in K} \gamma_{1,a}^k \tilde{T}_{it}^k \times (TMEAN_{it} - TMEAN_{it-1}) \end{aligned}$$

Note that to account for the difference between $f(\cdot)$ and $\tilde{f}(\cdot)$, the annual measure of realized temperature T_{it}^k is substituted in the final line for $\tilde{T}_{it}^k$, the expectation of realized temperature (we use the same Bartlett kernel to estimate $\tilde{T}_{it}^k$ as we do for $TMEAN_{it}$ — this kernel is described in Section C.6).

As discussed in Section 4.2, we also estimate a "lower bound" on adaptation costs. This lower bound is the integral from $\boldsymbol{C}(t_0)$ to $\boldsymbol{C}(t_1)$ along the tangency between the indifference curve at the initial climate, and the unobservable cost function. We approximate this integral as follows:

$$\begin{aligned} A^{LB}(\boldsymbol{b}^*(\boldsymbol{C}_1, Y)) - \boldsymbol{A}^{LB}(\boldsymbol{b}^*(\boldsymbol{C}_0, Y)) &= -\int_{\boldsymbol{C}(t_0)}^{\boldsymbol{C}(t_1)} \left.\frac{\partial \tilde{f}}{\partial \boldsymbol{C}}\right|_{t_0} d\boldsymbol{C} \\ &\approx -\sum_{t=t_0}^{t_1} \sum_{k \in K} \gamma_{1,a}^k \tilde{T}_{i0}^k \times (TMEAN_{it} - TMEAN_{it-1}) \end{aligned}$$

We calculate both of these bounds for each impact region, age group, and year. In our damage function results, we show an average value across the lower and upper bounds.

C.2 Data summary: EU

Code	Country	Number of NUTS2 regions	Years
AT	Austria	9	1990-2014 (not 1995)
BE	Belgium	11	1990-2014
BG	Bulgaria	6	1990-2014
CH	Switzerland	7	1991-2014
CY	Cyprus	1	1993-2014 (before 1993 is not by age-group)
CZ	Czech Republic	8	1992-2014
DE	Germany	50	2002-2014 (of these 2 regions from 2011-2014)
DK	Denmark	5	2007-2014
EE	Estonia	1	1990-2014
EL	Greece	4	1990-2014 (13 for 2013)
ES	Spain	19	1990-2014
FI	Finland	5	1990-2014
FR	France	22	1990-2014 (additional 4 regions for 2014)*
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 (of these 2 have age-divided from 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 (under 5 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 (of these 4 from 2000 , 2 from 2002, 5 for just 2014)

** NOTE: climate data is only available till 2012 onwards.

*** NOTE: one region (PT30) has no temperature data.

C.3 Covariate data

Country	ISO code	Mortality sample	Income sample
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 C.1: Years covered in mortality sample and available sub-national income data

C.4 Spatial units for projection

We create a set of boundaries that define the spatial units onto which we extrapolate temperature-mortality sensitivities derived from our estimation, and for which we create location-specific projected damages of climate change. To do so, we utilize politically defined regions, as they form a better scale for analysis than regular grids due to their use in collecting socioeconomic data. Moreover, these regions are generally more 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 the Global Administrative Region dataset (Global Administrative Areas, 2012) to delineate boundaries, but require fewer than the approximately 295,000 spatial units present in that dataset. We thus create a set of 24,378 agglomerated regions that allow for greater

comparability and computational feasibility than unagglomerated regions. We establish a set of criteria to create these regions that makes them approximately comparable with respect to population, and internally consistent with respect mean temperature, diurnal temperature range, and mean precipitation. A map of these regions is shown in Figure C.1.

Figure C.1: Map of the 24,378 impact regions

These impact regions provide the spatial unit onto which location-specific predictions are projected. We use a clustering algorithm to form these regions, such that they are roughly similar in total population, and so that they are internally homogenous with respect to mean temperature, diurnal temperature range, and mean precipitation.

We develop an algorithm which agglomerates administrative units from GADM2 into regions with approximately equal amounts of population and climate variability, and which are spatially compact.We first allot region targets to each country, based on population density and climatic variability. The population weighted target is $20000^{P_i}/\sum_i P_i$, for country populations P_i.

The climate weighted target is $20000^{A_i V_i}/\sum_i A_i V_i$ for areas A_i and $V_i = \frac{Var[T_i]}{E[Var[T]]} + \frac{Var[D_i]}{E[Var[D]]} + \frac{Var[P_i]}{E[Var[P]]} + \frac{Var[Q_i]}{E[Var[Q]]}$ where T_i is mean temperature, $E[T] = 8\,^\circ\mathrm{C}$, D_i is diurnal temperature range, $E[D] = 2.1\,^\circ\mathrm{C}$, P_i is precipitation in the wettest month, $E[P] = 250\,\mathrm{mm}$, Q_i is precipitation in the driest month, and $E[Q] = 26\,\mathrm{mm}$ (Hijmans et al., 2005). The final target region count for each country is the average of the population and climate weighted targets.

The target regions, relative to the available administrative levels, are shown in Figure C.2. For most countries, there is no available administrative division for our preferred resolution, as shown in Figure C.3.

For those countries for which the target number of regions is between the total country region count at any administrative level and half that count, we take the closest administrative level. Otherwise, the agglomeration algorithm is applied.

For the agglomeration algorithm, we calculate a number of attributes at the highest available administrative level within a given country. As the agglomerations are performed, the attributes of the new agglomerated region are generated from its component regions. These attributes are as follows:

- Contained regions ($\# = M$)
- Neighbors ($\# = N$)

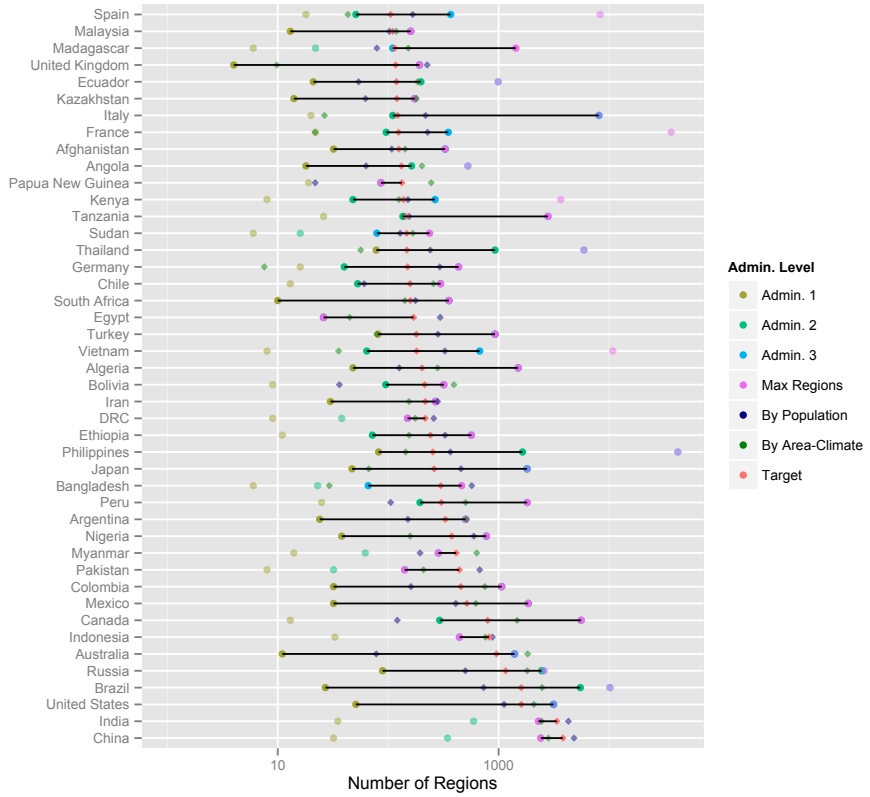

Figure C.2: Number of impact regions by country

Countries with over 1000 target regions ♦, based on their population ♦ and climatic ♦ targets. The range between administrative region counts above and below these targets are shown in black.

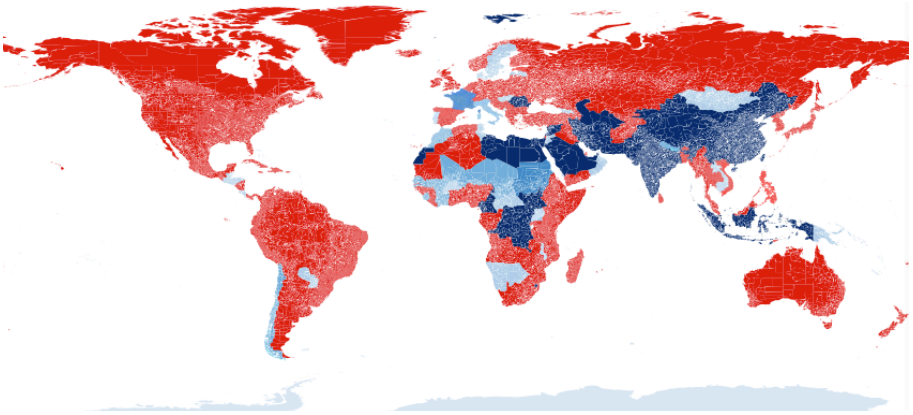

Figure C.3: Spatial scale of impact regions relative to administrative units

Let C be a country's target region count. Countries in dark blue have $C \geq$ Max Regions; lightest blue have $C \leq 1$; other shades of blue have an administrative region with $C/2 \leq N \leq C$. All others (red) need agglomeration.

- Population (P), from Bright et al. (2012) and area (A)
- Socioeconomic and Climatic traits, e.g., income, urban fraction, temperatures, biomes ($\{T\}$)
- Containing region centroids (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, these consist of:
   - The combination of R with each of its neighbors within S.
   - The next higher administrative region, S (e.g., all counties within the same state).
   - If neither of the above is available (e.g., an island state), the combination of R and the closest also at the first administrative level.
2. Each proposed agglomeration is scored, and this is compared to the score for the unagglomerated region. For a region R_i containing subregions indexed by j. The scores consists of a weighted sum of the following:

Attribute	**E**xpression	**W**eight
Area	$-(\sum_j {}^{A_j}/_{A_0})^2$, where A_0 is the average US county area	0.01
Population	$-(\sum_j {}^{P_j}/_{P_0})^2$, where P_0 is the average US county population	1
Dispersion	$-Var[Lat] - Var[Lon \cos E[Lat]]$	10
Other traits	$-\sum_T {}^{Var[T_i]}/_{T_0}$, 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	$-M^N/6\sqrt{M}$	1

3. The best possible agglomeration proposed by any region is identified, as determined by the smallest negative difference.
4. The regions within the agglomeration are merged, and new properties are applied to the new region.
5. This process repeats until the desired number of regions is reached.

C.5 Climate projection data

The NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP) dataset constitutes downscaled climate projections from 21 General Circulation Models (GCMs) participating in the Coupled Model Intercomparison Project Phase 5 (CMIP5) under two greenhouse

gas emission scenarios: Representative Concentration Pathway (RCP) 4.5 and 8.5 (Thrasher et al., 2012). The CMIP5 project (Taylor, Stouffer, and Meehl, 2012) provided the GCM simulations most heavily utilized in the Fifth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC AR5). There are two primary limitations when implementing the original GCM outputs into the current climate impact assessment. First, the relatively coarse resolution (about several to one degrees) of GCMs has limited their ability to capture small-scale climate patterns, which render them unsuitable for local climate impact assessment. Second, the GCM outputs exhibit large local bias compared with observations.

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 in a historical period. NEX-GDDP uses a climate dataset from the Global Meteorological Forcing Dataset (GMFD) for Land Surface Modeling developed by the Terrestrial Hydrology Research Group at Princeton University (Sheffield, Goteti, and Wood, 2006). The daily maximum temperature, daily minimum temperature, and daily precipitation at 0.25×0.25 degree resolution during the period of 1950-2005 are used in the downscaling process. A relationship between daily 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 outputs 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 GMFD data.

In our study, we use daily averaged temperature and daily precipitation in RCP4.5 and RCP8.5 scenarios from the NEX-GDDP dataset, where the daily averaged temperature is approximated as the mean of daily maximum and daily minimum temperature. The list of the 21 GCMs is shown in Table C.2. To assess the climate impact in administrative regions, the daily data at pixel level are spatially aggregated to "impact regions". For impact regions large enough to cover one or more grid cells, the aggregation is done by averaging the area-weighted values on these grids. If the impact region is small and contains no grid cell, the aggregated value of this region equals to the value at the nearest grid cell. We used median values of the 21 climate projections in each scenario and applied them in the mortality analysis.

ACCESS1-0	CSIRO-MK3-6-0	MIROC-ESM
bcc_csm1-1	GFDL-CM3	MIROC-ESM-CHEM
BNU-ESM	GFDL-ESM2G	MIROC5
CanESM2	GFDL-ESM2M	MPI-ESM-LR
CCSM4	inmcm4	MPI-ESM-MR
CESM1-BGC	IPSL-CM5A-LR	MRI-CGCM3
CNRM-CM5	IPSL-CM5A-MR	NorESM1-M

Table C.2: List of CMIP5 models included in NEX-GDDP

SMME and model surrogates

This section details the surrogate model mixed ensemble (SMME), which assigns probability weights to each CMIP5 model, as well as to surrogate models constructed using pattern scaling. the SMME method first divides the unit interval [0,1] into a set of bins. For this analysis, the bins are centered at the 1th, 6th, 11th, 16th, 33th, 50th, 67th, 82th, 89th, 94th, and 99th percentiles. Bins are narrower in the tails to ensure samples are created for portions of the GMST PDF 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 MAGGIC6 (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 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 $20\% \div 4 = 5\%$.

C.6 Determining the temporal dynamics of adaptation

The income covariate mediates the rate of income-based adaptation. If the income covariate is held at historical levels, no income-based adaptation is used. At the other extreme, if the contemporaneous income is applied in each year, then changes in income translate into immediate changes in sensitivity. Some benefits of income are expected to take many years to manifest, as richer governments and citizens invest in adaptive capital and enjoy greater health. We use a weighted average of recent year incomes ($z_{i,t-s}$), according to a Bartlett kernel, to calculate the effective level of income-based adaptation ($\bar{z}_{id}$):

$$\bar{z}_{it} = \frac{\sum_{s=1}^{L}(L-s+1)z_{i,t-s}}{\sum_{s=1}^{L}(L-s+1)}$$

To find a plausible length for the Bartlett kernel, we study changes in the response of mortality for people over 65 to temperature in the United States.

To do so, we estimate the same coefficients at each year. To remove the year fixed effect, we estimate the coefficient for the difference between each pair of years:

$$y_{it} - y_{i,t-1} = \alpha + \sum_k \beta_{kt}(T_{it}^k - T_{i,t-1}^k) + \text{controls} + \epsilon_{it}$$

Where y_{it} is the death rate for region i in period t, and T_{it}^k is the pixel average of the mean temperature raised to the power k for each region i. The controls are precipitation and precipitation squared. This produces a series of coefficients, β_{kt}, and their standard errors, σ_{kt}. The estimated coefficients are shown in Figure C.4.

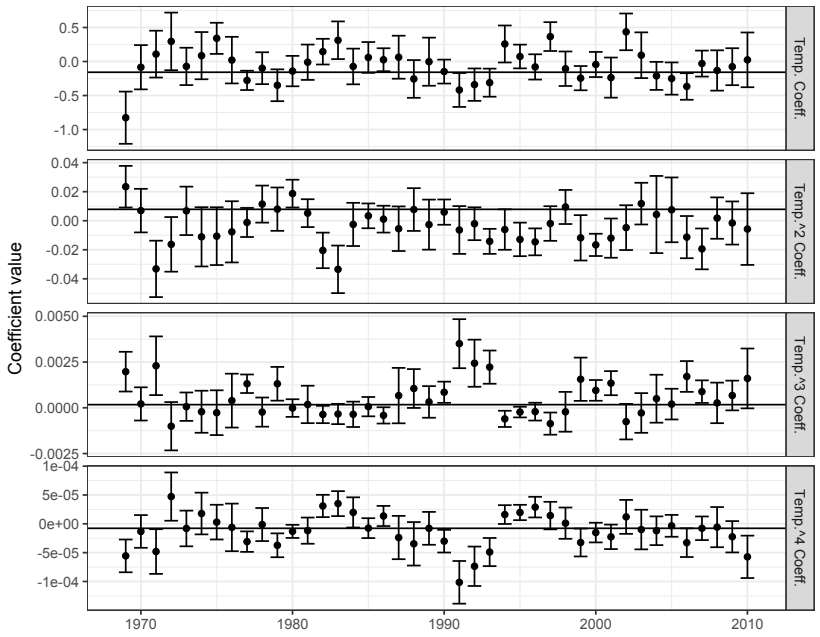

Figure C.4: Adaptation kernel estimates over time

Coefficients and 95% confidence intervals estimated for each year first-difference. The horizontal line is the estimated mean across all years, including a year fixed effect.

We use a Bayesian model to estimate the length of the Bartlett kernel that best explains these coefficients. Under the model, each coefficient is a draw from a Gaussian distribution with a mean that varies with the covariate:

$$\beta_{kt} \sim \mathcal{N}(\theta_k + \phi_k z_t, \tau_k + \sigma_{kt})$$

In this model, θ_k and ϕ_k correspond to the uninteracted and income interacted coefficients from our standard model, respectively. τ_k is a hyper-parameter which controls the rate of pooling of the data, so that if it is zero, inverse-variance weighting is used across individual year estimates.

The covariate z_t is calculated as a Bartlett kernel over up to 25 years of delayed income. National real income data is from the U.S. Bureau of Economic Analysis. The kernel is characterized by an unknown 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, as shown in figure C.5. This corresponds to the maximum likelihood estimated value.

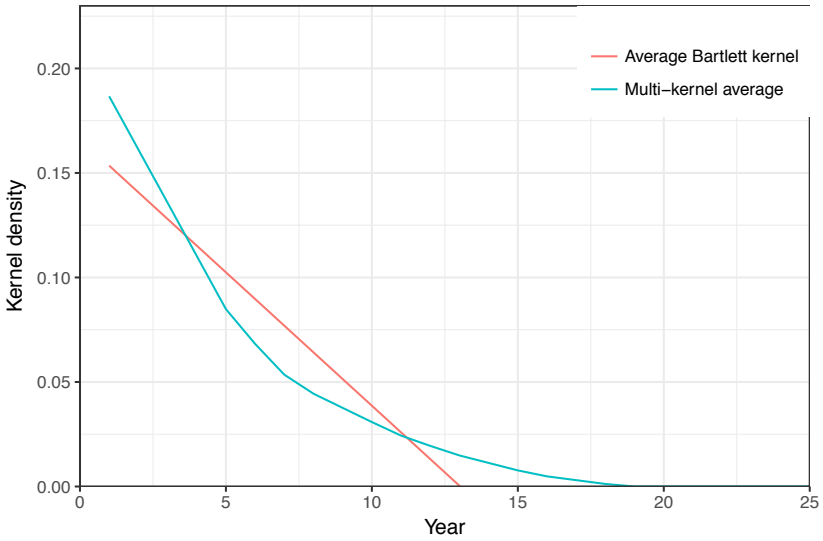

Figure C.5: Kernel weights used for speed of adaptation

Kernel weights shown for the estimated Bartlett kernel, and for the multi-kernel average. The multi-kernel average consists of the weights for all Bartlett kernel lengths (1 to 25) according to the estimated posterior probabilities.

C.7 Heterogeneity in the all-age mortality-temperature response function

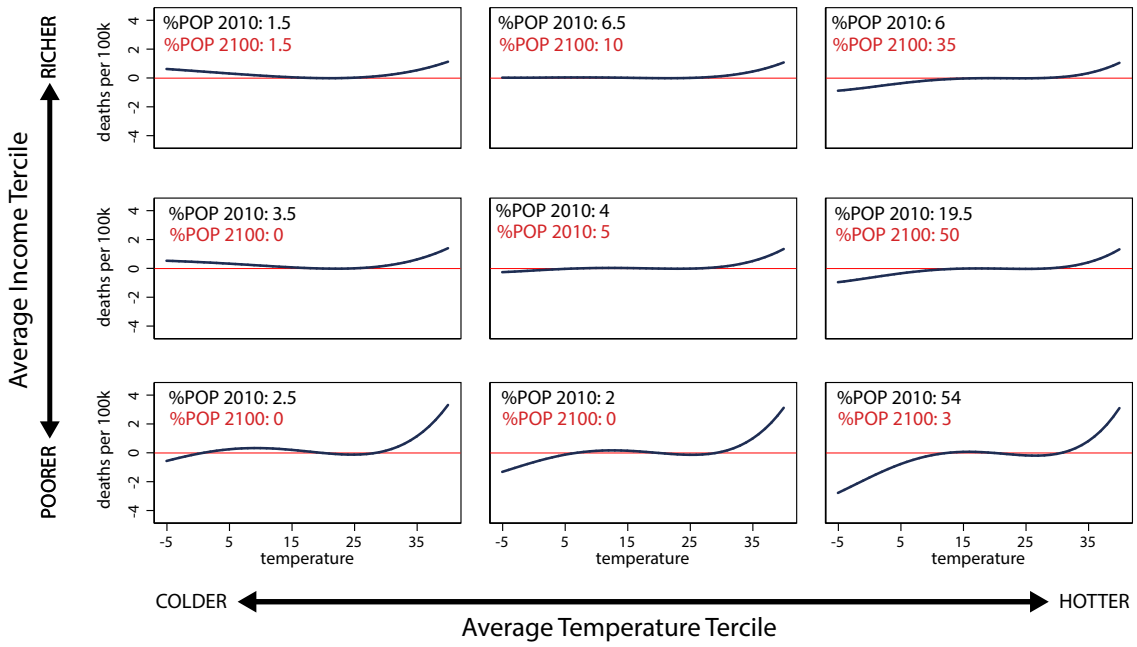

Figure C.6: Heterogeneity in the mortality-temperature relationship

Each panel represents a predicted all-age response function 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 poor, cold regions of our sample, while those in the upper right apply to the wealthy, 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. 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. The all-age response functions shown here are population-weighted averages of estimated coefficients in each of three age categories.

C.8 Leave-one-country-out cross-validation of the adaptation model

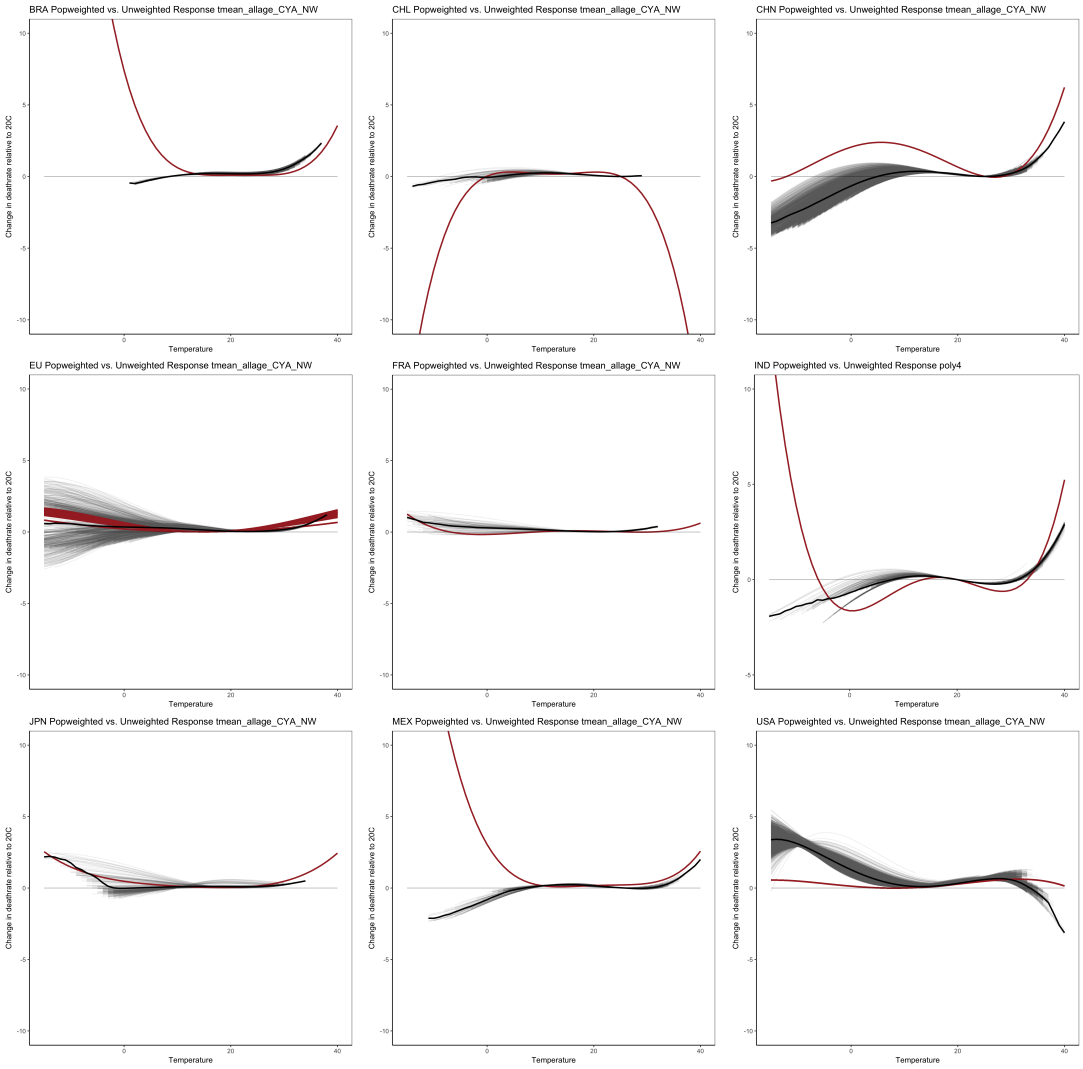

Figure C.7: Leave-one-country-out cross validation

In each panel, we test the ability of our interaction model (results shown in grey) to predict the actual response function for a country omitted from the sample (shown in red). That is, grey lines in all panels show predicted response functions for each impact region in each respective country, where predicted responses are estimated from the interaction model described in Section 4.4, but using a sample *that omits* data from the corresponding country. The solid black line is the unweighted average across all regions, while the red line is the estimated response function using only data from the country of interest. Congruence between red and black lines indicates good performance of our interaction model. Countries, from left to right, top to bottom: Brazil, Chile, China, EU (minus France), France, India, Japan, Mexico, USA.

C.9 Robustness to alternative functional form specifications and alternative historical climate datasets

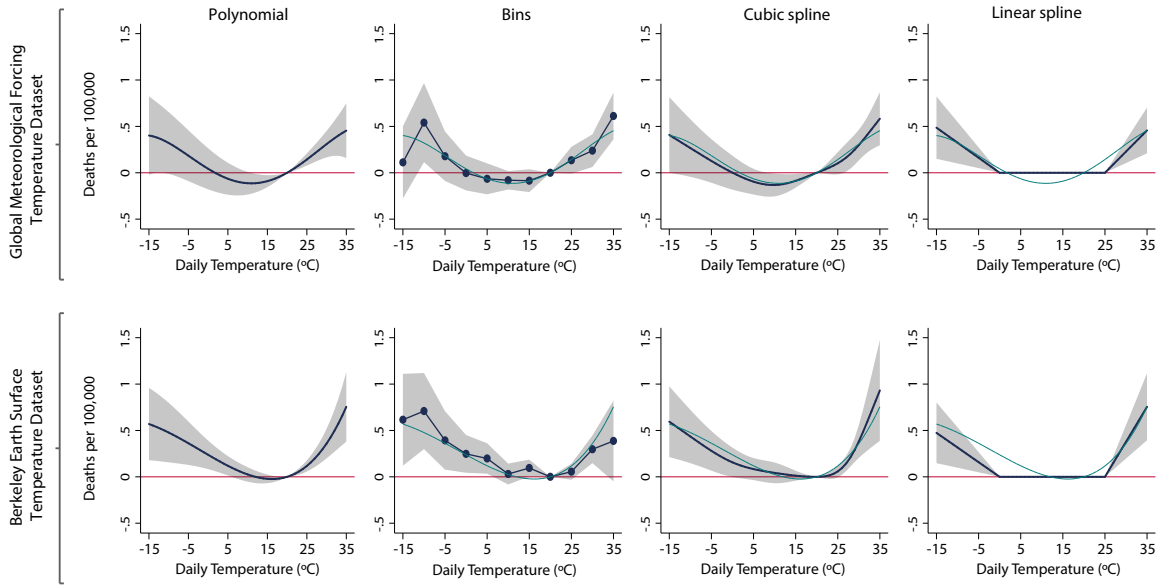

Figure C.8: Functional form robustness

Robustness of the all-age temperature-mortality 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.

C.10 Country-level mortality dose-response function

The regions in our sample represent substantial heterogeneity. To begin to examine this heterogeneity, we look at the variation of the temperature-mortality relationship across countries in our sample.[1] Table C.3 displays these results for the 9 countries or regions in our data. We additionally show only predictions at daily average temperatures actually experienced in each country over our sample period.

Temperature	BRA	CHL	CHN	EUR	FRA	JPN	MEX	USA
35°	0.801*						0.794**	0.473**
	(0.409)						(0.312)	(0.211)
30°	0.115	-1.716**	0.407	0.397	-0.088	0.315*	0.131	0.343***
	(0.117)	(0.856)	(0.684)	(0.312)	(0.762)	(0.167)	(0.145)	(0.094)
25°	0.026	-0.251*	-0.169	0.150	-0.072	0.065	0.018	0.166***
	(0.075)	(0.142)	(0.237)	(0.103)	(0.311)	(0.066)	(0.075)	(0.037)
0°		-0.255	1.113***	0.384	-0.265	0.269*	2.944**	-0.010
		(0.250)	(0.415)	(0.280)	(0.368)	(0.157)	(1.315)	(0.115)
-5°		-1.685**	0.816*	0.655*	-0.178	0.470*	7.761*	0.135
		(0.657)	(0.457)	(0.387)	(0.634)	(0.253)	(4.006)	(0.144)
-10°		-5.588***	0.445	0.937*	0.314	0.855*	16.662*	0.274
		(2.016)	(0.554)	(0.566)	(1.309)	(0.478)	(9.300)	(0.175)
Adj R-squared				.974				
Observations				787329				
Adm2-Age FE				YES				
Cntry-Year-Age FE				YES				

Adm1 clustered standard errors
*** p<0.01, ** p<0.05, * p<0.1

Table C.3: Heterogeneity by country in the mortality-temperature response function.

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. Point estimates indicate the marginal effect of increasing daily average temperature by 1°C, evaluated at each temperature value shown. Country-specific coefficients are generated by interacting all climate variables and fixed effects with country dummies.

C.11 Model selection

We develop a novel model selection criterion, based on our need to predict the sensitivity of mortality to temperature at different points in covariate space. First, we divide our data into

[1] We treat the EU here as a "country" for exposition purposes. 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.

four subsets based on the median values of the two covariates (income and average climate). We turn these subsets into dummy variables; that is, we create indicator variables for hot-rich, hot-cool, cold-rich, and cold-poor. We then define the "true" sensitivities within each of these four subsets by estimating a non-parametric binned regression, with bins that are 5°C or 3°C wide, where all binned temperature terms are interacted with these subset dummies. The model is estimated without regression weights:

$$M_{iat} = \sum_a \sum_k \sum_q \beta_q^{ka} \tilde{T}_{it}^k \times \mathbb{1}\{i \in q\} + FE^q + \varepsilon_{iat}^q \tag{C.3}$$

Where q indicates each subset of the data (e.g. hot-poor). We then predict the response surface using the parametric interaction model described in the main text. We store the predicted values from this model as $\hat{M}_{iat}$. We then create a comparable prediction of the "true" relationship in each subset by regressing predicted values $\hat{M}_{iat}$ on binned daily temperatures, again interacted with subset fixed effects, as above:

$$\hat{M}_{iat} = \sum_a \sum_k \sum_q \delta_q^{ka} \tilde{T}_{it}^k \times \mathbb{1}\{i \in q\} + FE^q + \epsilon_{iat}^q \tag{C.4}$$

Finally, we quantitatively compare these two estimates of the mortality rate by calculating, for every subset q, age group a and bin k:

$$error^{k,q,a} = \hat{\delta}_q^{ka} - \hat{\beta}_q^{ka} \tag{C.5}$$

We calculate $error^{k,q,a}$ for each subset, age group and bin combination. We then report a range of different values, each placing different weight on different bins and subsets.

		Bin 3C			*Bin 5C*		
	Functional form	Unweighted	People-days 2010	People-days 2090	Unweighted	People-days 2010	People-days 2090
	bins	8.697	.985	.45	5.283	1.006	.858
	cspline	7.642	1.116	.448	5.668	1.111	.75
	lspline	7.619	1.154	.568	5.256	1.205	.779
	poly4	6.944	.78	.403	4.867	.809	.705
	poly5	6.594	.909	.354	5.233	.942	.65
Bins >25°C	bins	15.582	.755	.565	4.024	.912	1.188
	cspline	13.152	.787	.509	5.534	.92	1.002
	lspline	14.045	.778	.356	6.143	.984	.837
	poly4	13.272	.619	.413	6.415	.803	.923
	poly5	12.561	.438	.308	7.238	.653	.838
Above median temp.	bins	3.739	.312	.242	3.788	.348	.467
& inc	cspline	3.965	.286	.224	2.946	.338	.384
	lspline	3.602	.271	.207	3.088	.306	.317
	poly4	3.515	.272	.207	2.552	.332	.366
	poly5	2.987	.226	.173	2.654	.273	.326
Age 65+	bins	6.217	2.537	1.322	4.189	2.593	2.537
	cspline	5.255	2.876	1.295	4.545	2.861	2.161
	lspline	5.242	2.966	1.664	4.174	3.099	2.221
	poly4	4.582	1.977	1.173	3.765	2.054	2.028
	poly5	4.105	2.323	1.021	4.129	2.409	1.778
Age 65+	bins	10.693	2.108	1.767	3.283	2.542	3.68
& above median temp.	cspline	8.319	2.193	1.591	4.557	2.556	3.036
	lspline	9.356	2.166	1.096	5.118	2.739	2.438
	poly4	8.716	1.705	1.275	5.32	2.209	2.772
	poly5	7.674	1.176	.95	6.034	1.762	2.362

Table C.4: Model selection across functional forms in temperature

In this table, we show the results from a custom model selection test we developed to ensure our two-factor model of adaptation sufficiently captures heterogeneity in the mortality-temperature response function across different regions of the world. The first column indicates the functional form used to model temperature, and all other columns show the root mean squared error (RMSE) as defined in equation C.5. Each panel shows the RMSE averaged over a distinct portion of the temperature and income support. "Unweighted" columns are unweighted averages, while "people-days" columns take population weighted averages, with population data either from 2010 or 2090. For each panel and each weighting scheme, the lowest RMSE functional form is shown in red.

C.12 Global damage table

VSL Method	**2010**		**2050**		**2098**	
A&G VSL, Global Average	*deaths*	*lifeyears*	*deaths*	*lifeyears*	*deaths*	*lifeyears*
RCP45-SSP3	.1035	.0406	.69	.2571	1.0726	.487
RCP85-SSP3	.006	-.0058	1.1468	.4829	10.5079	4.9282
A&G VSL, Scaled	*deaths*	*lifeyears*	*deaths*	*lifeyears*	*deaths*	*lifeyears*
RCP45-SSP3	.0629	.0204	.3196	.1136	.617	.2626
RCP85-SSP3	-.0063	-.0104	.5315	.2195	7.0837	3.165
EPA VSL, Global Average	*deaths*	*lifeyears*	*deaths*	*lifeyears*	*deaths*	*lifeyears*
RCP45-SSP3	.396	.1554	2.6397	.9835	4.1032	1.8632
RCP85-SSP3	.0231	-.022	4.3871	1.8472	40.1988	18.8533
EPA VSL, Scaled	*deaths*	*lifeyears*	*deaths*	*lifeyears*	*deaths*	*lifeyears*
RCP45-SSP3	.2408	.0782	1.2227	.4345	2.3602	1.0045
RCP85-SSP3	-.0239	-.0399	2.0331	.8398	27.099	12.1079

Table C.5: Mortality-related costs of climate change over time

Values shown are the mortality-related costs of climate change, displayed as shares of global GDP in each year. Values of deaths are either computed using a constant value per lost life, or are calculated using lifeyears, where life expectancy is adjusted for. Results are from a model with fourth order polynomial temperature, under SSP3. All values are the average share of global GDP lost, averaging across all 28 climate models.

Appendix D

Appendix: The implicit global water market

D.1 Supplemental figures

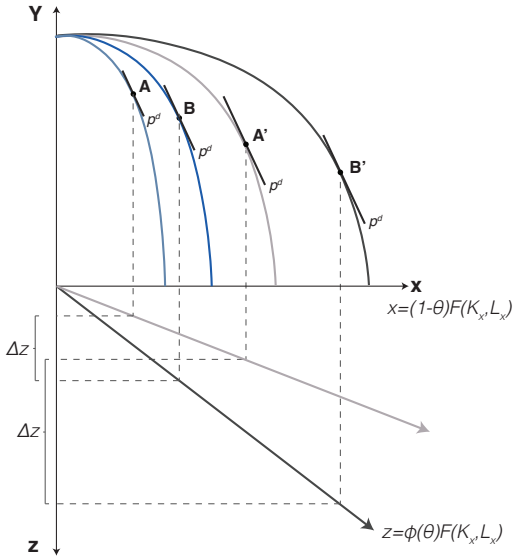

Figure D.1: Interaction between a decrease in cost of irrigation and exogenous agricultural productivity

This figure illustrates heterogeneity in the treatment effect of an irrigation cost shock on total water storage. The blue curves apply to a location with relatively low agricultural productivity (i.e. $F(\cdot)$ is low conditional on inputs K_x and L_x), while the gray curves apply to a location with relatively high agricultural productivity. A decrease in irrigation costs leads to a decrease in factors allocated within agriculture toward mitigating water use, and therefore to a shift outward in the production possibilities frontier, as well as an increase in slope in the $z - x$ dimension. For the location with lower productivity, this decrease in irrigation cost leads to a smaller increase in water use (Δz) than for the location with higher agricultural productivity ($\Delta z'$). Figure adapted from Copeland and Taylor (2005).

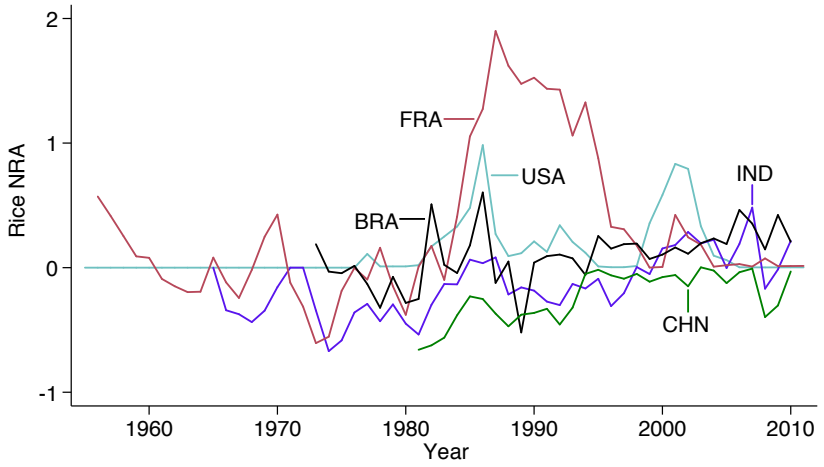

(a) Rice

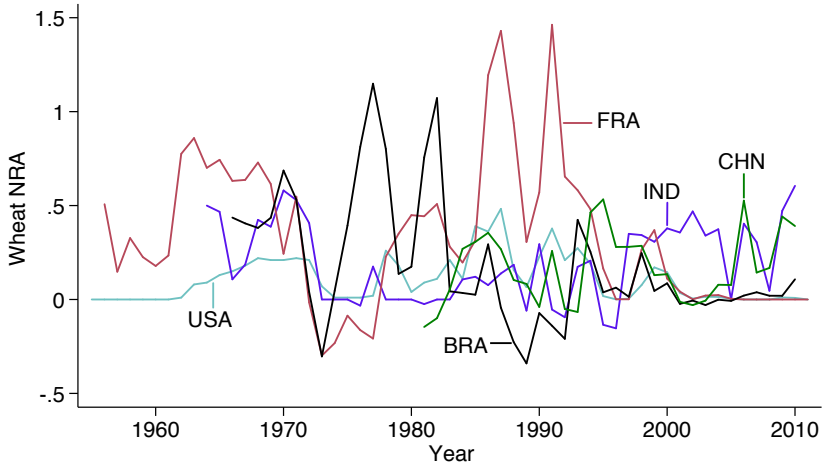

(b) Wheat

Figure D.2: Nominal rates of assistance (NRAs) for rice and wheat in selected countries

An NRA of 0.5 indicates that after accounting for all distortions imposed by governments, domestic producer gross returns are 50% above counterfactual undistorted gross returns. Data from Anderson, Rausser, and Swinnen (2013).

D.2 Supplemental tables

	(1) OLS	(2) OLS	(3) OLS	(4) Lat. FE	(5) Lat. FE	(6) Lat. FE
Trend in temp. (°C/mo)	-0.345**	-0.129	-0.218**	-0.107	0.080*	0.003
	(0.163)	(0.085)	(0.102)	(0.087)	(0.044)	(0.052)
Trend in precip. (mm/mo)	0.070***	0.070***	0.063***	0.056***	0.072***	0.060***
	(0.017)	(0.023)	(0.018)	(0.008)	(0.008)	(0.007)
Trend in lights	0.073	0.098	0.129*	-0.137***	-0.144***	-0.076**
	(0.076)	(0.063)	(0.068)	(0.039)	(0.039)	(0.033)
Crop area fraction (CAF)	-0.026	0.010		-0.001	0.029***	
	(0.016)	(0.023)		(0.010)	(0.011)	
Temp. trend × CAF		-10.866***			-8.931***	
		(2.335)			(0.873)	
Precip. trend × CAF		-0.080			-0.153***	
		(0.091)			(0.054)	
Lights trend × CAF		0.230			0.465**	
		(0.361)			(0.180)	
Temp. trend × CAF rice			-7.436			-6.091
			(11.911)			(5.684)
Temp. trend × CAF maize			-11.757			-7.076
			(11.400)			(7.400)
Temp. trend × CAF wheat			-35.833***			-30.900***
			(7.240)			(3.747)
Temp. trend × CAF soybean			25.355*			17.724***
			(13.963)			(6.103)
Temp. trend × CAF cotton			56.867			30.332
			(44.167)			(26.124)
Precip. trend × CAF rice			-0.492**			-0.587***
			(0.215)			(0.158)
Precip. trend × CAF maize			-1.278**			-1.605***
			(0.563)			(0.399)
Precip. trend × CAF wheat			-0.602			-0.213
			(0.435)			(0.248)
Precip. trend × CAF soybean			1.899***			1.883***
			(0.415)			(0.324)
Precip. trend × CAF cotton			2.228			0.503
			(1.361)			(0.756)
Lights trend × CAF rice			-0.026*			-0.029***
			(0.014)			(0.008)
Lights trend × CAF maize			0.011			0.005
			(0.020)			(0.014)
Lights trend × CAF wheat			-0.020*			-0.009
			(0.012)			(0.006)
Lights trend × CAF soybean			-0.037			-0.028*
			(0.025)			(0.016)
Lights trend × CAF cotton			0.012			0.015
			(0.034)			(0.019)
Observations	12,050	12,050	12,057	12,050	12,050	12,057
R-squared	0.027	0.058	0.098	0.441	0.461	0.490
Latitude FE	NO	NO	NO	YES	YES	YES

Table D.1: Cross-sectional "trends-on-trends" regression detects climatic and agricultural footprints in total water storage

Impacts of climate, night lights and crops on total water storage estimated in a cross-sectional trends-on-trends model. The regression includes a cross-section of estimate trends for 12,050 grid cells over the period 2003 to 2014. Grid cells above 60° and below -50°C latitude are omitted. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. All regressions are precision weighted, using the variance of the trend in total water storage. All crop-specific CAF variables are included; for clarity, these coefficients are not shown. Standard errors in parentheses are clustered at the water basin level. *** p<0.01, ** p<0.05, * p<0.1

	(1) Basin FE	(2) Basin FE	(3) Grid FE	(4) Grid FE
NRA rice × CAF rice	-5.886***	-6.516***	-6.122***	-5.513**
	(1.076)	(1.751)	(1.385)	(2.080)
NRA maize × CAF maize	2.972	4.400*	2.871	4.513
	(1.942)	(2.528)	(1.978)	(2.933)
NRA wheat × CAF wheat	-3.979***	-4.907***	-6.424***	-8.318***
	(1.432)	(1.596)	(2.358)	(2.578)
NRA soybean × CAF soybean	-6.825	-19.690*	1.227	-8.009
	(7.658)	(11.610)	(8.448)	(14.112)
NRA cotton × CAF cotton	0.977	0.317	-0.748	-2.434
	(2.400)	(1.690)	(3.446)	(2.720)
NRA rice	0.889**	6.921	1.108**	26.777***
	(0.409)	(5.685)	(0.457)	(8.935)
NRA maize	-0.940	-5.197	-0.705	6.689
	(0.642)	(4.284)	(0.591)	(4.475)
NRA wheat	0.744*	-5.638***	1.006**	-11.394***
	(0.399)	(2.111)	(0.473)	(2.616)
NRA soybean	-2.452***	-2.994	-3.439***	-22.181***
	(0.656)	(9.413)	(0.619)	(4.502)
NRA cotton	0.351***	-0.616	0.408***	0.968
	(0.123)	(2.383)	(0.135)	(1.870)
Night lights intensity	0.004	0.004	0.047	-0.014
	(0.006)	(0.007)	(0.036)	(0.037)
Night lights × CAF	-0.018	-0.001	-0.397***	-0.141*
	(0.016)	(0.019)	(0.121)	(0.081)
Observations	232,715	232,715	232,658	232,658
R-squared	0.368	0.372	0.371	0.377
Grid FE	NO	NO	YES	YES
Basin FE	YES	YES	NO	NO
Basin trends	YES	NO	YES	NO
Basin Year FE	NO	YES	NO	YES
Basin Month FE	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table D.2: Impact of crop-specific agricultural subsidies on total water storage

Observations are 1° × 1° grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. Because of data limitations for crop-specific nominal rates of assistance (NRA's), crop-specific regressions include grids in Australia, Brazil, China, Columbia, India, the U.S. and Zambia only. All regressions include fifth-order polynomials in temperature and precipitation, and 12 months of lagged precipitation. Standard errors in parentheses are clustered at the water basin level.

	(1) Basin FE	(2) Basin FE	(3) Grid FE	(4) Grid FE	(5) Basin FE	(6) Basin FE	(7) Grid FE	(8) Grid FE
NRA rice	0.564	6.391	0.781*	26.402***	0.889**	6.921	1.108**	26.777***
	(0.420)	(5.666)	(0.466)	(8.992)	(0.409)	(5.685)	(0.457)	(8.935)
NRA maize	-0.926	-4.962	-0.766	6.720	-0.940	-5.197	-0.705	6.689
	(0.635)	(4.192)	(0.573)	(4.503)	(0.642)	(4.284)	(0.591)	(4.475)
NRA wheat	0.429	-6.346***	0.585	-12.205***	0.744*	-5.638***	1.006**	-11.394***
	(0.443)	(2.200)	(0.516)	(2.639)	(0.399)	(2.111)	(0.473)	(2.616)
NRA soybean	-2.319***	-3.531	-3.085***	-22.365***	-2.452***	-2.994	-3.439***	-22.181***
	(0.565)	(9.259)	(0.502)	(4.485)	(0.656)	(9.413)	(0.619)	(4.502)
NRA cotton	0.415***	-0.329	0.450***	1.191	0.351***	-0.616	0.408***	0.968
	(0.131)	(2.354)	(0.138)	(1.884)	(0.123)	(2.383)	(0.135)	(1.870)
NRA rice × CAF rice					-5.886***	-6.516***	-6.122***	-5.513**
					(1.076)	(1.751)	(1.385)	(2.080)
NRA maize × CAF maize					2.972	4.400*	2.871	4.513
					(1.942)	(2.528)	(1.978)	(2.933)
NRA wheat × CAF wheat					-3.979***	-4.907***	-6.424***	-8.318***
					(1.432)	(1.596)	(2.358)	(2.578)
NRA soybean × CAF soybean					-6.825	-19.690*	1.227	-8.009
					(7.658)	(11.610)	(8.448)	(14.112)
NRA cotton × CAF cotton					0.977	0.317	-0.748	-2.434
					(2.400)	(1.690)	(3.446)	(2.720)
Night lights intensity	0.006	0.007	0.055	-0.003	0.004	0.004	0.047	-0.014
	(0.007)	(0.007)	(0.038)	(0.040)	(0.006)	(0.007)	(0.036)	(0.037)
Night lights × CAF	-0.009	0.006	-0.407***	-0.147*	-0.018	-0.001	-0.397***	-0.141*
	(0.019)	(0.021)	(0.120)	(0.082)	(0.016)	(0.019)	(0.121)	(0.081)
Observations	232,715	232,715	232,658	232,658	232,715	232,715	232,658	232,658
R-squared	0.368	0.372	0.370	0.377	0.368	0.372	0.371	0.377
Grid FE	NO	NO	YES	YES	NO	NO	YES	YES
Basin FE	YES	YES	NO	NO	YES	YES	NO	NO
Basin trends	YES	NO	YES	NO	YES	NO	YES	NO
Basin Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Basin Month FE	YES	YES	YES	YES	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table D.3: Impact of crop-specific agricultural subsidies on total water storage: Robustness to specification

Observations are 1° × 1° grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. Because of data limitations for crop-specific NRAs, crop-specific regressions include grids in Australia, Brazil, China, Columbia, India, the U.S. and Zambia only. All regressions include fifth-order polynomials in temperature and precipitation, and 12 months of lagged precipitation. Standard errors in parentheses are clustered at the water basin level.

	(1) Rice	(2) Maize	(3) Wheat	(4) Soybean	(5) Cotton
NRA rice	0.170				
	(0.210)				
NRA rice × CAF rice	-3.791***				
	(1.010)				
NRA maize		-0.305**			
		(0.119)			
NRA maize × CAF maize		0.129			
		(2.170)			
NRA wheat			-0.183		
			(0.137)		
NRA wheat × CAF wheat			-4.248**		
			(1.893)		
NRA soybean				-0.510*	
				(0.305)	
NRA soybean × CAF soybean				-2.230	
				(6.900)	
NRA cotton					-0.144**
					(0.062)
NRA cotton × CAF cotton					-2.028
					(3.814)
Night lights intensity	0.033	0.038*	0.030	0.047	0.033
	(0.026)	(0.021)	(0.019)	(0.034)	(0.027)
Night lights × CAF	-0.219***	-0.153**	-0.132**	-0.226**	-0.297***
	(0.082)	(0.061)	(0.058)	(0.090)	(0.094)
Observations	372,225	618,047	649,046	418,371	346,551
R-squared	0.368	0.400	0.403	0.408	0.363
Basin trends	YES	YES	YES	YES	YES
Basin Year FE	NO	NO	NO	NO	NO
Basin Month FE	YES	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table D.4: Impact of crop-specific agricultural subsidies on total water storage: Crops estimated individually

Observations are 1° × 1° grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. All regressions include fifth-order polynomials in temperature and precipitation, and 12 months of lagged precipitation. Standard errors in parentheses are clustered at the water basin level.

	(1) Robust	(2) Robust	(3) Cluster Lat	(4) Cluster Lat	(5) Cluster Basin	(6) Cluster Basin
NRA rice	1.181***	24.764***	1.181***	24.764***	1.181***	24.764***
	(0.172)	(1.033)	(0.174)	(4.205)	(0.427)	(7.760)
NRA maize	-0.534*	6.991***	-0.534**	6.991**	-0.534	6.991
	(0.276)	(0.942)	(0.252)	(2.667)	(0.484)	(5.135)
NRA wheat	0.723***	-10.060***	0.723***	-10.060***	0.723*	-10.060***
	(0.169)	(0.758)	(0.212)	(2.286)	(0.425)	(2.189)
NRA soybean	-3.439***	-24.865***	-3.439***	-24.865***	-3.439***	-24.865***
	(0.430)	(1.273)	(0.273)	(3.761)	(0.528)	(4.155)
NRA cotton	0.409***	1.829***	0.409***	1.829***	0.409***	1.829
	(0.093)	(0.314)	(0.063)	(0.693)	(0.135)	(1.155)
NRA rice × CAF rice	-6.584***	-5.273***	-6.584***	-5.273***	-6.584***	-5.273**
	(1.422)	(1.780)	(1.035)	(1.123)	(1.563)	(2.038)
NRA maize × CAF maize	-0.313	2.366	-0.313	2.366	-0.313	2.366
	(2.160)	(2.413)	(2.163)	(2.388)	(2.215)	(3.375)
NRA wheat × CAF wheat	-3.327**	-5.223***	-3.327***	-5.223***	-3.327	-5.223*
	(1.685)	(1.857)	(1.265)	(1.574)	(2.359)	(2.637)
NRA soybean × CAF soybean	-0.776	-9.265	-0.776	-9.265	-0.776	-9.265
	(7.269)	(7.910)	(5.969)	(8.277)	(10.046)	(15.433)
NRA cotton × CAF cotton	0.528	-0.982	0.528	-0.982	0.528	-0.982
	(2.086)	(2.158)	(1.528)	(0.942)	(2.876)	(2.079)
Avg monthly temp (°C)	0.034***	0.041***	0.034***	0.041***	0.034	0.041
	(0.008)	(0.008)	(0.012)	(0.013)	(0.058)	(0.062)
Avg monthly temp^2 (°C)	-0.002***	-0.002***	-0.002***	-0.002***	-0.002**	-0.002**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Monthly precip (cm)	-0.141***	-0.141***	-0.141*	-0.141*	-0.141	-0.141
	(0.044)	(0.044)	(0.083)	(0.083)	(0.150)	(0.160)
Monthly precip^2 (cm)	0.000***	0.000***	0.000***	0.000***	0.000*	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Night lights intensity	0.046	-0.025	0.046*	-0.025	0.046	-0.025
	(0.037)	(0.040)	(0.025)	(0.023)	(0.028)	(0.017)
Night lights × CAF	-0.441***	-0.175	-0.441***	-0.175***	-0.441***	-0.175***
	(0.114)	(0.117)	(0.081)	(0.055)	(0.115)	(0.057)
Observations	262,347	262,347	262,347	262,347	262,347	262,347
R-squared	0.351	0.358	0.351	0.358	0.351	0.358
Basin trends	YES	NO	YES	NO	YES	NO
Basin Year FE	NO	YES	NO	YES	NO	YES
Basin Month FE	YES	YES	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Table D.5: Impact of crop-specific agricultural subsidies on total water storage: Three sets of standard errors

Observations are 1° × 1° grid cells, excluding latitudes above 60° and below -50°C. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction.

	(1) OLS	(2) Lat. FE	(3) Continent FE	(4) Basin FE
NRA rice × CAF rice	-510.837***	-276.070***	-438.763***	-232.244***
	(102.503)	(48.104)	(110.404)	(86.744)
NRA maize × CAF maize	-89.918	-74.022***	-124.748**	-96.046**
	(60.234)	(22.499)	(54.619)	(45.797)
NRA wheat × CAF wheat	-22.812	-6.457	-19.102	-25.054
	(20.353)	(4.606)	(18.573)	(19.448)
NRA rice	3.262	17.517***	8.325	-7.981
	(4.632)	(3.564)	(6.241)	(6.451)
NRA maize	29.343***	20.998***	13.574	19.370**
	(6.324)	(2.416)	(11.013)	(8.325)
NRA wheat	11.140**	10.499***	29.707**	14.822*
	(4.658)	(1.413)	(12.637)	(8.006)
Night lights intensity	0.159	-0.222	0.283	0.933
	(0.619)	(0.222)	(0.597)	(0.733)
Night lights × CAF	0.317	-0.959***	-0.031	-2.671**
	(0.957)	(0.352)	(1.152)	(1.333)
Observations	3,882	3,882	3,882	3,882
R-squared	0.252	0.710	0.282	0.457
Lat FE	NO	YES	NO	NO
Continent FE	NO	NO	YES	NO
Basin FE	NO	NO	NO	YES

*** p<0.01, ** p<0.05, * p<0.1

Table D.6: Impact of long-run crop-specific changes in agricultural subsidies on long-run changes in total water storage

Results are from a long differences specification over the period 2003 to 2011. There was not sufficient variation to estimate soybean and cotton. "NRA" indicates Nominal Rate of Assistance, the wedge between domestic prices inclusive of policy interventions and the counterfactual free trade price. "CAF" indicates grid cell level cropped area fraction. Standard errors are clustered at the basin level.

Bibliography

Acemoglu, D., S. Johnson, and J. A. Robinson. 2001. "The colonial origins of comparative development: An empirical investigation." *The American Economic Review* 91 (5):1369–1401.

Acemoglu, Daron, Simon Johnson, and James A Robinson. 2012. "The colonial origins of comparative development: An empirical investigation: Reply." *The American Economic Review* 102 (6):3077–3110.

Adamopoulos, Tasso and Diego Restuccia. 2014. "The size distribution of farms and international productivity differences." *The American Economic Review* 104 (6):1667–1697.

Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson. 2017. "Nonparametric counterfactual predictions in neoclassical models of international trade." *The American Economic Review* 107 (3):633–689.

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham. 2014. "The light and the heat: Productivity co-benefits of energy-saving technology." *National Bureau of Economic Research Working Paper 24314* URL `http://economics.yale.edu/sites/default/files/adhvaryu.pdf`.

Agence-France Presse. 2016. "Thousands of farmer suicides prompt India to set up 1.3bn crop insurance scheme." *The Guardian* .

Alderman, Harold, John Hoddinott, and Bill Kinsey. 2006. "Long term consequences of early childhood malnutrition." *Oxford Economic Papers* 58 (3):450–474.

Allan, Tony. 2011. *Virtual water: tackling the threat to our planet's most precious resource*. IB Tauris.

Almond, Douglas and Janet Currie. 2011. "Killing me softly: The fetal origins hypothesis." *The Journal of Economic Perspectives* 25 (3):153.

Amundson, JL, TL Mader, RJ Rasby, and QS Hu. 2005. "Temperature and temperature–humidity index effects on pregnancy rate in beef cattle." In *Proc. 17th International Congress on Biometeorology, Deutscher Wetterdienst, Offenbach, Germany*.

Anderson, Kym. 2009. *Distortions to agricultural incentives: A global perspective, 1955-2007*. World Bank Publications.

Anderson, Kym, Gordon Rausser, and Johan Swinnen. 2013. "Political economy of public policies: Insights from distortions to agricultural and food markets." *Journal of Economic Literature* 51 (2):423–477.

Anderson, R Warren, Noel D Johnson, and Mark Koyama. 2013. "From the persecuting to the protective state? Jewish expulsions and weather shocks from 1100 to 1800." *SSRN Working Paper* URL `http://ssrn.com/abstract=2212323`.

André, G, B Engel, PBM Berentsen, Th V Vellinga, and AGJM Oude Lansink. 2011. "Quantifying the effect of heat stress on daily milk yield and monitoring dynamic changes using an adaptive dynamic model." *Journal of Dairy Science* 94 (9):4502–4513.

Andres, Antonio Rodriguez. 2005. "Income inequality, unemployment, and suicide: A panel data analysis of 15 European countries." *Applied Economics* 37 (4):439–451.

Ang, Beng Wah and FQ Zhang. 2000. "A survey of index decomposition analysis in energy and environmental studies." *Energy* 25 (12):1149–1176.

Annan, Francis and Wolfram Schlenker. 2015. "Federal crop insurance and the disincentive to adapt to extreme heat." *American Economic Review, Papers and Proceedings* 105 (5):262–66.

Anttila-Hughes, Jesse K. and Solomon M. Hsiang. 2012. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." *SSRN Working Paper* URL `http://papers.ssrn.com/abstract_id=2220501`.

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. 2014. "Cooling China: The weather dependence of air conditioner adoption." *Frontiers of Economics in China* 9 (1):70–84.

———. 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 24397* .

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.

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* .

Auffhammer, Maximilian and Erin T. Mansur. 2014. "Measuring climatic impacts on energy consumption: A review of the empirical literature." *Energy Economics* 46:522 – 530. URL `http://www.sciencedirect.com/science/article/pii/S0140988314001017`.

Auffhammer, Maximilian, V Ramanathan, and Jeffrey R Vincent. 2006. "Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in India." *Proceedings of the National Academy of Sciences* 103 (52):19668–19672.

———. 2012. "Climate change, the monsoon, and rice yield in India." *Climatic Change* 111 (2):411–424.

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

Ayres, Andrew B and Kyle C Meng. 2017. "The Economic Value of Secure Water: Landowner Returns to Defining Groundwater Property Rights." *Working Paper* .

Bai, Y and J.K Kung. 2010. "Climate shocks and Sino-nomadic conflict." *Review of Economics and Statistics* (0).

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, Olivier Deschenes, and Melanie Guldi. 2015. "Maybe next month? Temperature shocks, climate change, and dynamic adjustments in birth rates." Tech. rep., National Bureau of Economic Research Working Paper 21681.

Barreca, Alan I. 2010. "The long-term economic impact of in utero and postnatal exposure to malaria." *Journal of Human Resources* 45 (4):865–892.

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

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

Barreca, Alan I and Jay P Shimshack. 2012. "Absolute humidity, temperature, and influenza mortality: 30 years of county-level evidence from the United States." *American Journal of Epidemiology* 176 (suppl 7):S114–S122.

Barrios, S, L Bertinelli, and E Strobl. 2010. "Trends in rainfall and economic growth in Africa: A neglected cause of the African growth tragedy." *Review of Economics and Statistics* 92 (2):350–366.

Bastos, Paulo, Matias Busso, and Sebastian Miller. 2013. "Adapting to climate change: Long-term effects of drought on local labor markets." *IDB Working Paper No. IDB-WP-466* URL `http://idbdocs.iadb.org/wsdocs/getdocument.aspx?docnum=38335273`.

Basu, Rupa and Jonathan M Samet. 2002. "Relation between elevated ambient temperature and mortality: a review of the epidemiologic evidence." *Epidemiologic Reviews* 24 (2):190–202.

Basyan, Ceren, F. Gonzalez, M. Burke, S. Hsiang, and E. Miguel. 2014. "Economic and non-economic factors in violence: Evidence from organized crime, suicides, and climate in Mexico." *Working paper* .

Baylis, Patrick. 2015. "Temperature and temperament: Evidence from a billion tweets." *Working Paper* URL `https://ei.haas.berkeley.edu/research/papers/WP265.pdf`.

Becker, Gary S and Richard A Posner. 2004. "Suicide: An economic approach." *University of Chicago* .

Besley, Timothy and Robin Burgess. 2002. "The political dconomy of government responsiveness: Theory and evidence from India." *The Quarterly Journal of Economics* 117 (4):1415–1451.

Bhatt, Samir, Peter W Gething, Oliver J Brady, Jane P Messina, Andrew W Farlow, Catherine L Moyes, John M Drake, John S Brownstein, Anne G Hoen, Osman Sankoh et al. 2013. "The global distribution and burden of dengue." *Nature* 496 (7446):504–507.

Blakeslee, David S and Ram Fishman. 2017. "Weather shocks, agriculture, and crime: evidence from India." *Journal of Human Resources* :0715–7234R1.

Bohra-Mishra, Pratikshya, Michael Oppenheimer, and Solomon M. Hsiang. 2014. "Nonlinear permanent migration responses to climatic variations but minimal response to disasters." *Proceedings of the National Academy of Sciences* 111 (27):9780–9785.

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`.

Bryant, JR, N López-Villalobos, JE Pryce, CW Holmes, and DL Johnson. 2007. "Quantifying the effect of thermal environment on production traits in three breeds of dairy cattle in New Zealand." *New Zealand Journal of Agricultural Research* 50 (3):327–338.

Bryla-Tressler, D et al. 2011. "Weather index insurance for agriculture: Guidance for development practitioners." Tech. rep., World Bank.

Buckley, B.M, K.J Anchukaitis, D Penny, R Fletcher, E.R Cook, M Sano, L.C Nam, A Wichienkeeo, T.T Minh, and T.M Hong. 2010. "Climate as a contributing factor in the demise of Angkor, Cambodia." *Proceedings of the National Academy of Sciences* 107 (15):6748.

Büntgen, U., W. Tegel, K. Nicolussi, M. McCormick, D. Frank, V. Trouet, J.O. Kaplan, F. Herzig, K.U. Heussner, H. Wanner et al. 2011. "2500 years of European climate variability and human susceptibility." *Science* 331 (6017):578.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2014. "The unequal effects of weather and climate change: Evidence from mortality in India." *Working Paper* URL `http://www.lse.ac.uk/economics/people/facultyPersonalPages/facultyFiles/RobinBurgess/UnequalEffectsOfWeatherAndClimateChange140514.pdf`.

Burke, M, M Craxton, CD Kolstad, C Onda, H Allcott, E Baker, L Barrage, R Carson, K Gillingham, J 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.

Burke, Marshall, Erick Gong, and Kelly Jones. 2015. "Income shocks and HIV in Africa." *The Economic Journal* 125 (585):1157–1189.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel. 2015a. "Climate and conflict." *Annual Review of Economics* 7 (1):577–617.

———. 2015b. "Global non-linear effect of temperature on economic production." *Nature* 527 (7577):235–239.

Burke, M.B., E. Miguel, S. Satyanath, J.A. Dykema, and D.B. Lobell. 2009. "Warming increases the risk of civil war in Africa." *Proceedings of the National Academy of Sciences* 106 (49):20670.

Burke, Paul J. 2012. "Economic growth and political survival." *The B.E. Journal of Macroeconomics* 12 (1).

Burke, Paul J and Andrew Leigh. 2010. "Do output contractions trigger democratic change?" *American Economic Journal: Macroeconomics* 2 (4):124–157.

Cai, Ruohong, Shuaizhang Feng, Mariola Pytliková, and Michael Oppenheimer. 2014. "Climate variability and international migration: The importance of the agricultural linkage." *IZA Discussion Paper* URL `http://ftp.iza.org/dp8183.pdf`.

Carleton, T, SM Hsiang, and M Burke. 2016. "Conflict in a changing climate." *The European Physical Journal Special Topics* 225 (3):489–511.

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

Cattaneo, Cristina and Giovanni Peri. 2016. "The migration response to increasing temperatures." *Journal of Development Economics* 122:127–146.

Chaney, Eric. 2013. "Revolt on the nile: Economic shocks, religion, and political power." *Econometrica* 81 (5):2033–2053.

Chapagain, A and A Hoekstra. 2008. *Globalization of Water: Sharing the Planet's Freshwater Resources*. Oxford, UK: Blackwell Publishing.

Cherniwchan, Jevan. 2017. "Trade liberalization and the environment: Evidence from NAFTA and US manufacturing." *Journal of International Economics* 105:130–149.

Cherniwchan, Jevan, Brian R Copeland, and M Scott Taylor. 2017. "Trade and the environment: New methods, measurements, and results." *Annual Review of Economics* 9:59–85.

Chichilnisky, Graciela. 1994. "North-south trade and the global environment." *The American Economic Review* :851–874.

Chichilnisky, Graciela and Geoffrey Heal. 1998. "Economic returns from the biosphere." *Nature* 391 (6668):629.

Chong, Howard and David Sunding. 2006. "Water markets and trading." *Annual Review of Environment and Resources* 31 (1):239–264.

Cian, E de, IS Wing et al. 2014. "Climate change impacts on energy demand." *CMCC Research Paper* (RP0240). URL `http://www.cmcc.it/wp-content/uploads/2015/02/rp0240-cip-12-20141.pdf`.

Colacito, Riccardo, Bridget Hoffmann, and Toan Phan. 2014. "Temperatures and Growth: a Panel Analysis of the US." *Working Paper* URL `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2546456`.

Colmer, Jonathan. 2016. "Weather, labour reallocation, and industrial production: Evidence from India." *Working Paper* URL `https://drive.google.com/file/d/0B-BakBtoHwF8UjNtU3ZONmdjOTA/view`.

Colón-González, Felipe J, Carlo Fezzi, Iain R Lake, and Paul R Hunter. 2013. "The effects of weather and climate change on dengue." *PLoS Neglected Tropical Diseases* 7 (11):e2503.

Copeland, Brian R. and M. Scott Taylor. 2005. "Free trade and global warming: a trade theory view of the Kyoto protocol." *Journal of Environmental Economics and Management* 49 (2):205–234.

Costanza, Robert, Ralph d'Arge, Rudolf De Groot, Stephen Farber, Monica Grasso, Bruce Hannon, Karin Limburg, Shahid Naeem, Robert V O'neill, Jose Paruelo et al. 1997. "The value of the world's ecosystem services and natural capital." *Nature* 387 (6630):253.

Costinot, Arnaud, Dave Donaldson, and Cory Smith. 2016. "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world." *Journal of Political Economy* 124 (1):205–248.

Craig, MH, RW Snow, and D Le Sueur. 1999. "A climate-based distribution model of malaria transmission in sub-Saharan Africa." *Parasitology Today* 15 (3):105–111.

Cullen, H, P Demenocal, S Hemming, G Hemming, F Brown, T Guilderson, and F Sirocko. 2000. "Climate change and the collapse of the Akkadian empire: Evidence from the deep sea." *Geology* 28 (4):379.

Culp, Peter W, Robert J Glennon, and Gary Libecap. 2014. *Shopping for water: How the market can mitigate water shortages in the American West*. Springer.

Currie, Janet and Maya Rossin-Slater. 2013. "Weathering the storm: Hurricanes and birth outcomes." *Journal of Health Economics* 32 (3):487–503.

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.

Dalin, Carole, Yoshihide Wada, Thomas Kastner, and Michael J Puma. 2017. "Groundwater depletion embedded in international food trade." *Nature* 543 (7647):700–704.

d'Amour, Christopher Bren, Leonie Wenz, Matthias Kalkuhl, Jan Christoph Steckel, and Felix Creutzig. 2016. "Teleconnected food supply shocks." *Environmental Research Letters* 11 (3):035007.

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.

De Montesquieu, Charles Baron. 2011. *The Spirit of Laws*. Cosimo, Inc.

Debaere, P and T Li. 2016. "The effects of water markets: Evidence from the Rio Grande." In *AGU Fall Meeting Abstracts*.

Deisenhammer, EA, G Kemmler, and P Parson. 2003. "Association of meteorological factors with suicide." *Acta Psychiatrica Scandinavica* 108 (6):455–459.

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* :66–95.

———. 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* .

Department for Environment, Food and Rural Affairs. 2005. "The Social Cost of Carbon Review." Tech. rep., UK Department for Environment, Food and Rural Affairs.

Deryugina, Tatyana. 2013. "How do people update? The effects of local weather fluctuations on beliefs about global warming." *Climatic Change* 118 (2):397–416.

———. 2017. "The fiscal cost of hurricanes: disaster aid versus social insurance." *American Economic Journal: Economic Policy* 9 (3):168–98.

Deryugina, Tatyana and Solomon Hsiang. 2017. "The marginal product of climate." *National Bureau of Economic Research Working Paper 24072* .

Deschenes, Olivier. 2014. "Temperature, human health, and adaptation: A review of the empirical literature." *Energy Economics* 46:606–619.

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 Jonathan Guryan. 2009. "Climate change and birth weight." *American Economic Review, Papers and Proceedings* 99:211–217.

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

Deshpande, RS. 2002. "Suicide by farmers in Karnataka: Agrarian distress and possible alleviatory steps." *Economic and Political Weekly* :2601–2610.

Desmet, Klaus and Esteban Rossi-Hansberg. 2015. "On the spatial economic impact of global warming." *Journal of Urban Economics* 88:16–37.

Diamond, Jared. 2005. *Collapse: How societies choose to fail or succeed*. Penguin.

Diaz, Delavane B. 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." *SSRN Working Paper* .

Donna, Javier and José-Antonio Espín-Sánchez. 2016. "The illiquidity of water markets: Efficient institutions for water allocation in southeastern Spain." *SSRN Working Paper* .

Drabo, Alassane and Linguère Mbaye. 2011. "Climate change, natural disasters and migration: An empirical analysis in developing countries." *IZA Discussion Paper* URL `http://ftp.iza.org/dp5927.pdf`.

Duflo, Esther and Rohini Pande. 2007. "Dams." *The Quarterly Journal of Economics* :601–646.

Durkheim, Emile. 1951. "Suicide: A study in sociology (JA Spaulding & G. Simpson, trans.)." *Glencoe, IL: Free Press.(Original work published 1897)* .

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

Eyer, Jonathan and Casey Wichman. 2014. "The effect of water supply shocks on the electricity generation mix: Implications for climate change." *Working Paper* (14):1459–1469.

Famiglietti, JS, M Lo, SL Ho, J Bethune, KJ Anderson, TH Syed, SC Swenson, CR De Linage, and M Rodell. 2011. "Satellites measure recent rates of groundwater depletion in California's Central Valley." *Geophysical Research Letters* 38 (3).

Feng, S., M. Oppenheimer, and W. Schlenker. 2012. "Climate Change, Crop Yields, and Internal Migration in the United States." *National Bureau of Economic Research Working Paper 17734* .

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.

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

Feyrer, James and Bruce Sacerdote. 2009. "Colonialism and modern income: islands as natural experiments." *The Review of Economics and Statistics* 91 (2):245–262.

Fischer, Günther, HT Van Velthuizen, MM Shah, and Freddy O Nachtergaele. 2002. "Global agro-ecological assessment for agriculture in the 21st century: Methodology and results." *IIASA* .

Fishman, Ram. 2016. "More uneven distributions overturn benefits of higher precipitation for crop yields." *Environmental Research Letters* 11 (2):024004.

Fishman, Ram, Jason Russ, and Paul Carrillo. 2015. "Long-term impacts of high temperatures on economic productivity." *Working Paper* URL `https://www.gwu.edu/~iiep/assets/docs/papers/2015WP/FishmanRussCarrillo_October2015.pdf`.

Fishman, Ram Mukul. 2011. "Climate change, rainfall variability, and adaptation through irrigation: Evidence from Indian agriculture." *Job Market Paper* .

Froom, Paul, Yeheskial Caine, Igal Shochat, and Joseph Ribak. 1993. "Heat Stress and Helicopter Pilot Errors." *Journal of Occupational and Environmental Medicine* 35:720–724.

Gallup, J. L., J. D. Sachs, and A. D. Mellinger. 1999. "Geography and economic development." *Consulting Assistance on Economic Reform II, Discussion Papers* (39).

Garrett, Thomas A and Russell S Sobel. 2003. "The political economy of FEMA disaster payments." *Economic Inquiry* 41 (3):496–509. URL `http://drjoesaviak.com/uploads/EmergencyManagementPolitics_DisastersStudyGarrettSobel.pdf`.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez De Silanes, and Andrei Shleifer. 2014. "Growth in regions." *Journal of Economic Growth* 19 (3):259–309.

Gething, Peter W., David L. Smith, Anand P. Patil, Andrew J. Tatem, Robert W. Snow, and Simon I. Hay. 2010. "Climate change and the global malaria recession." *Nature* 465 (7296):342–345. URL `http://dx.doi.org/10.1038/nature09098`.

Gething, Peter W, Thomas P Van Boeckel, David L Smith, Carlos A Guerra, Anand P Patil, Robert W Snow, and Simon I Hay. 2011. "Modelling the global constraints of temperature on transmission of Plasmodium falciparum and P. vivax." *Parasites & Vectors* 4 (92):4.

Gisser, Micha. 1983. "Groundwater: focusing on the real issue." *Journal of Political Economy* 91 (6):1001–1027.

Gleick, Peter H. 1993. *Water in crisis: a guide to the worlds fresh water resources.* New York New York Oxford University Press 1993.

———. 2016. "Water strategies for the next administration." *Science* 354 (6312):555–556.

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

Graff Zivin, Joshua, Solomon M Hsiang, and Matthew Neidell. 2018. "Temperature and human capital in the short and long Run." *Journal of the Association of Environmental and Resource Economists* 5 (1):77–105.

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.

Grafton, R Quentin, Gary D Libecap, Eric C Edwards, RJ Bob O'Brien, and Clay Landry. 2012. "Comparative assessment of water markets: insights from the Murray–Darling Basin of Australia and the Western USA." *Water Policy* 14 (2):175–193.

Gray, Clark L and Valerie Mueller. 2012. "Natural disasters and population mobility in Bangladesh." *Proceedings of the National Academy of Sciences* 109 (16):6000–6005.

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

Guiteras, Raymond. 2009. "The impact of climate change on Indian agriculture." *Manuscript, Department of Economics, University of Maryland, College Park, Maryland* .

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.

Hagerty, Nick. 2017. "Liquid constrained in California: Estimating the potential gains from water markets." *Working paper* .

Hajat, S, KL Ebi, RS Kovats, B Menne, S Edwards, and A Haines. 2005a. "The human health consequences of flooding in Europe: a review." In *Extreme weather events and public health responses*. Springer, 185–196.

Hajat, Shakoor, Ben G Armstrong, Nelson Gouveia, and Paul Wilkinson. 2005b. "Mortality displacement of heat-related deaths: A comparison of Delhi, Sao Paulo, and London." *Epidemiology* 16 (5):613–620.

Hamermesh, Daniel S and Neal M Soss. 1974. "An economic theory of suicide." *The Journal of Political Economy* :83–98.

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.

Harari, Mariaflavia and Eliana La Ferrara. 2013. "Conflict, climate and cells: A disaggregated analysis." *Working paper* URL `http://economics.mit.edu/files/10058`.

Hardin, Garrett. 1968. "The tragedy of the commons." *Science* 13 (162):1243–1248.

Haug, G.H, D Günther, L.C Peterson, D.M Sigman, K.A Hughen, and B Aeschlimann. 2003. "Climate and the collapse of Maya civilization." *Science* 299 (5613):1731.

Haukka, Jari, Kirsi Suominen, Timo Partonen, and Jouko Lönnqvist. 2008. "Determinants and outcomes of serious attempted suicide: A nationwide study in Finland, 1996–2003." *American Journal of Epidemiology* 167 (10):1155–1163.

Heal, Geoffrey. 2000. "Valuing ecosystem services." *Ecosystems* 3 (1):24–30.

Heal, Geoffrey and Jisung Park. 2013. "Feeling the heat: Temperature, physiology & the wealth of nations." *National Bureau of Economic Research Working Paper 19725* .

———. 2015. "Goldilocks dconomies? Temperature stress and the direct impacts of climate change." *National Bureau of Economic Research Working Paper 21119* URL `http://www.nber.org/papers/w21119`.

Healy, Andrew and Neil Malhotra. 2009. "Myopic voters and natural disaster policy." *American Political Science Review* 103 (03):387–406.

Henderson, J Vernon, Adam Storeygard, and Uwe Deichmann. 2014. "50 years of urbanization in Africa: Examining the role of climate change." *World Bank Policy Research Working Paper* (6925). URL `http://www-wds.worldbank.org/external/default/WDSContentServer/WDSP/IB/2014/06/16/000158349_20140616110124/Rendered/PDF/WPS6925.pdf`.

Henderson, J Vernon, Adam Storeygard, and David N Weil. 2012. "Measuring economic growth from outer space." *The American Eonomic Review* 102 (2):994–1028.

Herring, Ronald J. 2008. "Whose numbers count? Probing discrepant evidence on transgenic cotton in the Warangal district of India." *International Journal of Multiple Research Approaches* 2 (2):145–159.

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 Working Paper* .

Hidalgo, Daniel F, Suresh Naidu, Simeon Nichter, and Neal Richardson. 2010. "Economic determinants of land invasions." *The Review of Economics and Statistics* 92 (3):505–523.

Hijioka, Y., E. Lin, J.J. Pereira, R.T. Corlett, X. Cui, G.E. Insarov, R.D. Lasco, E. Lindgren, and A. Surjan. 2014. "Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." Cambridge University Press Cambridge, UK, 1327–1370.

Hijmans, Robert J, Susan E Cameron, Juan L Parra, Peter G Jones, and Andy Jarvis. 2005. "Very high resolution interpolated climate surfaces for global land areas." *International Journal of Climatology* 25 (15):1965–1978.

Hoddinott, John and Bill Kinsey. 2001. "Child growth in the time of drought." *Oxford Bulletin of Economics and Statistics* 63 (4):409–436.

Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81:945–960.

Hornbeck, Richard. 2012. "The enduring impact of the American Dust Bowl: Short and long-run adjustments to environmental catastrophe." *The American Economic Review* 102 (4):1477–1507.

Houser, Trevor, Robert Kopp, Solomon M. Hsiang, Michael Delgado, Amir Jina, Kate Larsen, Michael Mastrandrea, Shashank Mohan, Robert Muir-Wood, DJ Rasmussen, James Rising, and Paul Wilson. 2015. *American Climate Prospectus: Economic Risks in the United States*. Columbia University Press.

Howe, Stephen. 2002. *Empire: A very short introduction*. Oxford Paperbacks.

Hsiang, S.M., K.C. Meng, and M.A. Cane. 2011. "Civil conflicts are associated with the global climate." *Nature* 476 (7361):438–441.

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 M. 2010. "Temperatures and Cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences* 107 (35):15367–15372.

Hsiang, Solomon M, Marshall Burke, and Edward Miguel. 2013. "Quantifying the influence of climate on human conflict." *Science* 341 (6151):1235367.

Hsiang, Solomon M and Amir S Jina. 2014. "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones." *National Bureau of Economic Research Working Paper* .

———. 2015. "Geography, depreciation, and growth." *American Economic Review, Papers and Proceedings* 105 (5):252–56.

Hsiang, Solomon M and Kyle C Meng. 2014. "Reconciling disagreement over climate–conflict results in Africa." *Proceedings of the National Academy of Sciences* 111 (6):2100–2103.

———. 2015. "Tropical economics." *American Economic Review, Papers and Proceedings* 105 (5):257–261.

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 Socal Cost of Carbon. 2010. "Social Cost of Carbon for Regulatory Impact Analysis - Under Executive Order 12866." 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.

Iyer, L and P Topalova. 2014. "Poverty and crime: Evidence from rainfall and trade shocks in India." *Working paper* URL `http://www.hbs.edu/faculty/Publication%20Files/14-067_45092fee-b164-4662-894b-5d28471fa69b.pdf`.

Jacob, B, L Lefgren, and E Moretti. 2007. "The dynamics of criminal behavior: Evidence from weather shocks." *Journal of Human Resources* 42 (3):489–527.

Jacob, Thomas, John Wahr, W Tad Pfeffer, and Sean Swenson. 2012. "Recent contributions of glaciers and ice caps to sea level rise." *Nature* 482 (7386):514–518.

Jaglom, Wendy S, James R McFarland, Michelle F Colley, Charlotte B Mack, Boddu Venkatesh, Rawlings L Miller, Juanita Haydel, Peter A Schultz, Bill Perkins, Joseph H Casola et al. 2014. "Assessment of projected temperature impacts from climate change on the US electric power sector using the Integrated Planning Model®." *Energy Policy* 73:524–539.

Jessoe, Katrina, Dale T Manning, J Edward Taylor et al. 2014. "Climate Change and Labor Markets in Rural Mexico: Evidence from Annual Fluctuations in Weather." *Working paper* URL `http://ageconsearch.umn.edu/bitstream/170556/2/WeatherLaborMexico_AAEA.pdf`.

Jia, Ruixue. 2014. "Weather shocks, sweet potatoes and peasant revolts in historical China." *The Economic Journal* 124 (575):92–118.

Jones, Benjamin and Benjamin Olken. 2010. "Climate shocks and exports." *American Economic Review, Papers and Proceedings* 100:454–459.

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

Kalnay, Eugenia, Masao Kanamitsu, Robert Kistler, William Collins, Dennis Deaven, Lev Gandin, Mark Iredell, Suranjana Saha, Glenn White, John Woollen et al. 1996. "The NCEP/NCAR 40-year reanalysis project." *Bulletin of the American Meteorological Society* 77 (3):437–471.

Kaluza, Pablo, Andrea Kölzsch, Michael T Gastner, and Bernd Blasius. 2010. "The complex network of global cargo ship movements." *Journal of the Royal Society Interface* 7 (48):1093–1103.

Kazianga, Harounan and Christopher Udry. 2006. "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso." *Journal of Development Economics* 79 (2):413–446.

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

Kenrick, D. T and S. W Macfarlane. 1986. "Ambient temperature and horn honking: A field study of the heat/aggression relationship." *Environment and Behavior* 18 (2):179–191.

Key, Nigel, Stacy Sneeringer, and David Marquardt. 2014. "Climate change, heat stress, and US dairy production." *USDA-ERS Economic Research Report* (175).

Kim, Nam Kyu. 2016. "Revisiting economic shocks and coups." *Journal of Conflict Resolution* 60 (1):3–31.

Kleemans, Marieke. 2014. "Migration Choice under Risk and Liquidity Constraints." *University of California, Berkeley* URL `http://www.economics.illinois.edu/seminars/documents/Kleemans.pdf`.

Kolb, Peter, Christine Gockel, and Lioba Werth. 2012. "The effects of temperature on service employees' customer orientation: an experimental approach." *Ergonomics* 55 (6):621–635.

Kopp, Robert E, Solomon M Hsiang, and Michael Oppenheimer. 2013. "Empirically calibrating damage functions and considering stochasticity when integrated assessment models are used as decision tools." *Impacts World 2013, International Conference on Climate Change Effects* .

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* .

Kovats, R Sari, Shakoor Hajat, and Paul Wilkinson. 2004. "Contrasting patterns of mortality and hospital admissions during hot weather and heat waves in Greater London, UK." *Occupational and Environmental Medicine* 61 (11):893–898.

Kudamatsu, Masayuki, Torsten Persson, and David Strömberg. 2012. "Weather and infant mortality in Africa." *CEPR Discussion Paper No. DP9222* .

Kung, James Kai-sing and Chicheng Ma. 2012. "Can Cultural Norms Reduce Conflicts? Confucianism and Peasant Rebellions in Qing China." *Working paper* URL `http://ahec2012.org/papers/S6B-2_Kai-singKung_Ma.pdf`.

Kunreuther, Howard C and Erwann O Michel-Kerjan. 2009. *At war with the weather: managing large-scale risks in a new era of catastrophes*. MIT Press.

Larrick, Richard P., Thomas A. Timmerman, Andrew M. Carton, and Jason Abrevaya. 2011. "Temper, Temperature, and Temptation: Heat-Related Retaliation in Baseball." *Psychological Science* 22:423–428.

Lee, Henry Desmond Pritchard et al. 1952. *Meteorologica*, vol. 397. Loeb Classical Library.

Leonard, Bryan and Gary D Libecap. 2016. "Economic analysis of property rights: First possession of water in the American West." *Working paper* .

Levinson, Arik. 2009. "Technology, international trade, and pollution from US manufacturing." *The American Economic Review* 99 (5):2177–92. URL `http://www.aeaweb.org/articles?id=10.1257/aer.99.5.2177`.

Linnerud, Kristin, Torben K Mideksa, and Gunnar S Eskeland. 2011. "The impact of climate change on nuclear power supply." *The Energy Journal* 32 (0):1.

Lobell, David B and Marshall B Burke. 2008. "Why are agricultural impacts of climate change so uncertain? The importance of temperature relative to precipitation." *Environmental Research Letters* 3 (3):034007.

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.

Lobell, David B, Graeme L Hammer, Greg McLean, Carlos Messina, Michael J Roberts, and Wolfram Schlenker. 2013. "The critical role of extreme heat for maize production in the United States." *Nature Climate Change* 3 (5):497–501.

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.

Lobell, David B, Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate trends and global crop production since 1980." *Science* 333 (6042):616–620.

Lomborg, Bjørn. 2004. *Global crises, global solutions*. Cambridge University Press.

Long, Di, Bridget R. Scanlon, Laurent Longuevergne, Alexander Y. Sun, D. Nelun Fernando, and Himanshu Save. 2013. "GRACE satellite monitoring of large depletion in water storage in response to the 2011 drought in Texas." *Geophysical Research Letters* 40 (13):3395–3401.

Lövheim, Hugo. 2012. "A new three-dimensional model for emotions and monoamine neurotransmitters." *Medical Hypotheses* 78 (2):341–348.

Luthcke, Scott B, TJ Sabaka, BD Loomis, AA Arendt, JJ McCarthy, and J Camp. 2013. "Antarctica, Greenland and Gulf of Alaska land-ice evolution from an iterated GRACE global mascon solution." *Journal of Glaciology* 59 (216):613–631.

Maccini, Sharon and Dean Yang. 2009. "Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall." *The American Economic Review* 99 (3):1006.

Mackworth, Norman H. 1946. "Effects of Heat on Wireless Telegraphy Operators Hearing and Recording Morse Messages." *British Journal of Industrial Medicine* 3:143–158.

Mann, J John, Alan Apter, Jose Bertolote, Annette Beautrais, Dianne Currier, Ann Haas, Ulrich Hegerl, Jouko Lonnqvist, Kevin Malone, Andrej Marusic et al. 2005. "Suicide prevention strategies: A systematic review." *Journal of the American Medical Association* 294 (16):2064–2074.

Marchiori, Luca, Jean-François Maystadt, and Ingmar Schumacher. 2012. "The impact of weather anomalies on migration in sub-Saharan Africa." *Journal of Environmental Economics and Management* 63 (3):355–374.

Martens, Willem JM. 1998. "Climate change, thermal stress and mortality changes." *Social Science & Medicine* 46 (3):331–344.

Martin, Leslie A. 2012. "Energy efficiency gains from trade: Greenhouse gas emissions and India's manufacturing firms." *Department of Agricultural and Resource Economics, University of California Berkley* .

Maystadt, J and O Ecker. 2014. "Extreme Weather and Civil War in Somalia: Does Drought Fuel Conflict in Somalia through Livestock Price Shocks?" *American Journal of Agricultural Economics* .

McCord, GC. 2016. "Malaria ecology and climate change." *The European Physical Journal Special Topics* 225 (3):459–470.

McDermott, Grant R and Øivind A Nilsen. 2014. "Electricity prices, river temperatures, and cooling water scarcity." *Land Economics* 90 (1):131–148.

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/`.

Mekonnen, M. M. and A. Y. Hoekstra. 2011. "The green, blue and grey water footprint of crops and derived crop products." *Hydrology and Earth System Sciences* 15 (5):1577–1600.

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.

Metcalf, Gilbert E and James H Stock. 2017. "Integrated assessment models and the social cost of carbon: a review and assessment of US experience." *Review of Environmental Economics and Policy* 11 (1):80–99.

Miguel, E. 2005. "Poverty and witch killing." *Review of Economic Studies* 72 (4):1153–1172.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti. 2004. "Economic shocks and civil conflict: An instrumental variables approach." *Journal of Political Economy* 112 (4):725–753.

Milly, Paul CD, Julio Betancourt, Malin Falkenmark, Robert M Hirsch, Zbigniew W Kundzewicz, Dennis P Lettenmaier, and Ronald J Stouffer. 2008. "Stationarity is dead: Whither water management?" *Science* 319 (5863):573–574.

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`.

Mohanty, Bibhuti B. 2005. "'We are Like the Living Dead': Farmer Suicides in Maharashtra, Western India." *Journal of Peasant Studies* 32 (2):243–276.

Monfreda, Chad, Navin Ramankutty, and Jonathan A Foley. 2008. "Farming the planet: 2. Geographic distribution of crop areas, yields, physiological types, and net primary production in the year 2000." *Global Biogeochemical Cycles* 22 (1).

Moore, Frances C and Delavane B Diaz. 2015. "Temperature impacts on economic growth warrant stringent mitigation policy." *Nature Climate Change* 5 (2):127.

Moore, Frances C and David B Lobell. 2015. "The fingerprint of climate trends on European crop yields." *Proceedings of the National Academy of Sciences* 112 (9):2670–2675.

Muñoz, J Ricardo and David J Sailor. 1998. "A modelling methodology for assessing the impact of climate variability and climatic change on hydroelectric generation." *Energy Conversion and Management* 39 (14):1459–1469.

Munshi, Kaivan. 2003. "Networks in the modern economy: Mexican migrants in the US labor market." *The Quarterly Journal of Economics* :549–599.

Nawrotzki, Raphael J, Lori M Hunter, Daniel M Runfola, and Fernando Riosmena. 2015. "Climate change as a migration driver from rural and urban Mexico." *Environmental Research Letters* 10 (11):114023.

Nordhaus, W. D. 2006. "Geography and macroeconomics: New data and new findings." *Proceedings of the National Academy of Sciences* 103 (10).

Nordhaus, William D. 1993. "Optimal greenhouse-gas reductions and tax policy in the DICE model." *The American Economic Review* :313–317.

———. 1994. *Managing the global commons: The economics of climate change*, vol. 31. MIT press Cambridge, MA.

O'Bannon, C., J. Carr, D. A. Seekell, and P. D'Odorico. 2014. "Globalization of agricultural pollution due to international trade." *Hydrology and Earth System Sciences* 18 (2):503–510.

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

———. 2011b. "Responding to climatic challenges: lessons from US agricultural development." In *The Economics of Climate Change: Adaptations Past and Present*. University of Chicago Press, 169–194.

Olmstead, Sheila M. 2010. "The economics of managing scarce water resources." *Review of Environmental Economics and Policy* 4 (2):179–198.

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.

O'Neill, Marie S, Shakoor Hajat, Antonella Zanobetti, Matiana Ramirez-Aguilar, and Joel Schwartz. 2005. "Impact of control for air pollution and respiratory epidemics on the estimated associations of temperature and daily mortality." *International Journal of Biometeorology* 50 (2):121–129.

Patel, Raj. 2007. *Stuffed and starved: Markets, power and the hidden battle for the world food system*. Black Inc.

Patel, Vikram, Chinthanie Ramasundarahettige, Lakshmi Vijayakumar, JS Thakur, Vendhan Gajalakshmi, Gopalkrishna Gururaj, Wilson Suraweera, and Prabhat Jha. 2012. "Suicide mortality in India: A nationally representative survey." *The Lancet* 379 (9834):2343–2351.

Peduzzi, Pascal, Andréa De Bono, and Christian Herold. 2015. "Making Development Sustainable: The Future of Disaster Risk Management, Global Assessment Report on Disaster Risk Reduction." URL `http://www.preventionweb.net/english/hyogo/gar/2015/en/gar-pdf/GAR2015_EN.pdf`.

Piao, Shilong, Philippe Ciais, Yao Huang, Zehao Shen, Shushi Peng, Junsheng Li, Liping Zhou, Hongyan Liu, Yuecun Ma, Yihui Ding et al. 2010. "The impacts of climate change on water resources and agriculture in China." *Nature* 467 (7311):43–51.

Pindyck, Robert S. 2013. "Climate change policy: What do the models tell us?" *NBER Working paper 19244* URL `http://www.nber.org/papers/w19244`.

Ralston, Laura. 2015. "Conflict and Climate: a Micro-level Analysis." *CEGA Working paper* :1–59URL `http://cega.berkeley.edu/assets/miscellaneous_files/19_-ABCA-Conflict_and_Climate_Ralston.pdf`.

Ranson, Matthew. 2014. "Crime, weather, and climate change." *Journal of Environmental Economics and Management* 67 (3):274–302.

Rao, VM and DV Gopalappa. 2004. "Agricultural Growth and Farmer Distress: Tentative Perspectives from Karnataka." *Economic and Political Weekly* :5591–5598.

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`.

Ray, Russell S, Andrea E Corcoran, Rachael D Brust, Jun Chul Kim, George B Richerson, Eugene Nattie, and Susan M Dymecki. 2011. "Impaired respiratory and body temperature control upon acute serotonergic neuron inhibition." *Science* 333 (6042):637–642.

Rehkopf, David H and Stephen L Buka. 2006. "The association between suicide and the socio-economic characteristics of geographical areas: A systematic review." *Psychological Medicine* 36 (02):145–157.

Revesz, Richard L, Peter H Howard, Kenneth Arrow, Lawrence H Goulder, Robert E Kopp, Michael A Livermore, Michael Oppenheimer, and Thomas Sterner. 2014. "Global warming: Improve economic models of climate change." *Nature* 508 (7495):173–175.

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.

Richey, Alexandra S., Brian F. Thomas, Min-Hui Lo, John T. Reager, James S. Famiglietti, Katalyn Voss, Sean Swenson, and Matthew Rodell. 2015. "Quantifying renewable groundwater stress with GRACE." *Water Resources Research* 51 (7):5217–5238.

Roberts, Michael J and Wolfram Schlenker. 2013. "Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate." *The American Economic Review* 103 (6):2265–2295.

Rockström, Johan, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin III, Eric F Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber et al. 2009. "A safe operating space for humanity." *Nature* 461 (7263):472.

Rodell, Matthew, Jianli Chen, Hiroko Kato, James S. Famiglietti, Joe Nigro, and Clark R. Wilson. 2006. "Estimating groundwater storage changes in the Mississippi River basin (USA) using GRACE." *Hydrogeology Journal* 15 (1):159–166.

Rodell, Matthew, Isabella Velicogna, and James S. Famiglietti. 2009. "Satellite-based estimates of groundwater depletion in India." *Nature* 460 (7258):999–1002.

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.

Sailor, David J and AA Pavlova. 2003. "Air conditioning market saturation and long-term response of residential cooling energy demand to climate change." *Energy* 28 (9):941–951.

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* .

Samuelson, Paul A. 1953. "Prices of factors and good in general equilibrium." *The Review of Economic Studies* :1–20.

Sarma, EAS. 2004. "Is rural economy breaking down? Farmers' suicides in Andhra Pradesh." *Economic and Political Weekly* :3087–3089.

Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher. 2005. "Will US agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach." *American Economic Review* :395–406.

Schlenker, Wolfram, W. Michael Hanemann, and Snthony C. Fisher. 2007. "Water availability, degree days and the potential impact of climate change on irrigated agriculture in California." *Climatic Change* 81:19–38.

Schlenker, Wolfram and David B Lobell. 2010. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5 (1):014010.

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.

Schoengold, Karina and David Zilberman. 2007. "The economics of water, irrigation, and development." *Handbook of Agricultural Economics* 3:2933–2977.

Sekhri, Sheetal and Adam Storeygard. 2012. "Dowry deaths: Consumption smoothing in response to climate variability in India." *Working paper* URL `http://www.virginia.edu/economics/RePEc/vir/virpap/papers/virpap407.pdf`.

Sengupta, Somini. 2006. "On India's farms, a plague of suicide." *New York Times* .

Seo, Dongju, Christopher J Patrick, and Patrick J Kennealy. 2008. "Role of serotonin and dopamine system interactions in the neurobiology of impulsive aggression and its comorbidity with other clinical disorders." *Aggression and Violent Behavior* 13 (5):383–395.

Seppanen, O, William J Fisk, and QH Lei. 2006a. "Room temperature and productivity in office work." *Lawrence Berkeley National Laboratory* .

Seppanen, Olli, William J Fisk, and QH Lei. 2006b. "Effect of temperature on task performance in office environment." *Lawrence Berkeley National Laboratory* .

Shaun Morrison, Christopher Madden, Kazuhiro Nakamura. 2008. "Central Control of Thermogenesis in Mammals." *Experimental Physiology* 93:773–797.

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.

Shrader, Jeffrey. 2016. "Expectations and adaptation to environmental risks." *Working paper* URL `http://acsweb.ucsd.edu/~jgshrade/papers/forecasts_and_adaptation.pdf`.

Small, Jennifer, Scott J. Goetz, and Simon I. Hay. 2003. "Climatic suitability for malaria transmission in Africa, 1911–1995." *Proceedings of the National Academy of Sciences* 100 (26):15341–15345. URL `http://www.pnas.org/content/100/26/15341.abstract`.

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

Solow, Robert Merton. 1991. "Sustainability: An economist's perspective." URL `http://isites.harvard.edu/fs/docs/icb.topic203569.files/Solow.Sustainability_An_Economists_Perspective._1993.pdf`.

Somanathan, E, R Somanathan, A Sudarshan, and M Tewari. 2015. "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing." *Working paper* URL `http://www.isid.ac.in/~pu/dispapers/dp14-10.pdf`.

Stack, Steven. 2000. "Suicide: A 15-year review of the sociological literature part I: Cultural and economic factors." *Suicide and Life-Threatening Behavior* 30 (2):145–162.

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

Strobl, Eric. 2011. "The economic growth impact of hurricanes: evidence from US coastal counties." *Review of Economics and Statistics* 93 (2):575–589.

Sutch, Richard. 2011. "The impact of the 1936 corn belt drought on American farmers' adoption of hybrid corn." In *The economics of climate change: Adaptations past and present*. University of Chicago Press, 195–223.

Tapley, Byron D., Srinivas Bettadpur, John C. Ries, Paul F. Thompson, and Michael M. Watkins. 2004. "GRACE Measurements of Mass Variability in the Earth System." *Science* 305 (5683):503–505.

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

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

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

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

Tol, Richard SJ. 2002. "Estimates of the damage costs of climate change. Part 1: Benchmark estimates." *Environmental and Resource Economics* 21 (1):47–73.

Tol, R.S.J and S Wagner. 2010. "Climate change and violent conflict in Europe over the last millennium." *Climatic Change* .

Townsend, Robert M. 1995. "Consumption insurance: An evaluation of risk-bearing systems in low-income economies." *The Journal of Economic Perspectives* 9 (3):83–102.

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

Wahr, John M, Steven R Jayne, and Frank O Bryan. 2002. "A method of inferring changes in deep ocean currents from satellite measurements of time-variable gravity." *Journal of Geophysical Research: Oceans* 107 (C12).

Waldhoff, Stephanie T, David Anthoff, Steven Rose, and Richard SJ Tol. 2011. "The marginal damage costs of different greenhouse gases: An application of FUND." *Economics Discussion Paper* (2011-43).

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." *Proc. Natl. Acad. Sci.* .

Wenz, Leonie and Anders Levermann. 2016. "Enhanced economic connectivity to foster heat stress–related losses." *Science advances* 2 (6):e1501026.

Werner, Antweiler, Copeland Brian, and Taylor Scott. 2001. "Is free trade good for the environment?" *American Economic Review* 91 (4):877–908.

Wilde, Joshua, Bénédicte Apouey, Toni Jung et al. 2014. "Heat waves at conception and later life outcomes." *Unpublished Manuscript* URL `http://economics.usf.edu/PDF/Wilde_Apouey_Jung_HeatWavesConception.pdf`.

Willmott, Cort and Kenji Matsuura. 2014. "Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 2013)." URL `http://climate.geog.udel.edu/~climate/`.

Wineman, Ayala and Brian Mulenga. 2014. "Sensitivity of field crops to climate shocks in Zambia." *Working paper* 18. URL `http://fsg.afre.msu.edu/climate_change/Crop_yield_and_weather_in_Zambia_August_16,_2014.pdf`.

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.

Wooldridge, Jeffrey M. 2010. *Econometric analysis of cross section and panel data*. MIT press.

World Health Organization. 2014. "A global brief on vector-borne diseases." Tech. rep., World Health Organization. URL `http://apps.who.int/iris/bitstream/10665/111008/1/WHO_DCO_WHD_2014.1_eng.pdf`.

Yancheva, Gergana, Norbert R Nowaczyk, Jens Mingram, Peter Dulski, Georg Schettler, Jörg F W Negendank, Jiaqi Liu, Daniel M Sigman, Larry C Peterson, and Gerald H Haug. 2007. "Influence of the intertropical convergence zone on the East Asian monsoon." *Nature* 445 (7123):74.

Zaveri, Esha, Danielle S. Grogan, Karen Fisher-Vanden, Steve Frolking, Richard B. Lammers, Douglas H. Wrenn, Alexander Prusevich, and Robert E. Nicholas. 2016. "Invisible water, visible impact: groundwater use and Indian agriculture under climate change." *Environmental Research Letters* 11 (8):084005.

Zhang, D, C Jim, G Lin, Y He, J Wang, and H Lee. 2006. "Climate change, wars and dynastic cycles in China over the last millennium." *Climatic Change* .

Zhang, David D., Peter Brecke, Harry F. Lee, Yuan-Qing He, , and Jane Zhang. 2007. "Global climate change, war and population decline in recent human history." *Proceedings of the National Academy of Sciences* 104 (49):19214–19219.

Zhang, Peng, Junjie Zhang, Olivier Deschenes, and Kyle Meng. 2016. "Temperature Effects on Productivity and Factor Reallocation: Evidence from a Half Million Chinese Manufacturing Plants." *Working paper* URL `http://www.econ.ucsb.edu/about_us/events/seminar_papers/Zhang.pdf`.

Zhou, Guofa, Noboru Minakawa, Andrew K Githeko, and Guiyun Yan. 2004. "Association between climate variability and malaria epidemics in the East African highlands." *Proceedings of the National Academy of Sciences of the United States of America* 101 (8):2375–2380.

Ziebarth, Nicolas R, Maike Schmitt, and Martin Karlsson. 2014. "The short-term population health effects of weather and pollution." *Working paper* URL `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2377611`.