

Final report

Indian climate early warning system

Tamma Carleton
Michael Greenstone
Solomon Hsiang
Andrew Hultgren
Amir Jina
Robert Kopp
Ashwin Rode

July 2017

When citing this paper, please
use the title and the following
reference number:
E-89342-INC-1

IGC

International
Growth Centre



DIRECTED BY



FUNDED BY



Indian Climate Early Warning System

Tamma Carleton¹, Michael Greenstone², Solomon Hsiang¹, Andrew Hultgren¹,
Amir Jina², Robert Kopp³, and Ashwin Rode²

¹University of California, Berkeley

²University of Chicago

³Rutgers University

July 26, 2017

Abstract

Extreme weather poses numerous risks to India’s economy and society, and future climate change is expected to amplify these risks. However, at present there exist no comprehensive assessments of how either contemporaneous weather shocks or future climate change might affect India’s economic sectors, even at a nationally aggregated level. Furthermore, location-specific information on climate vulnerabilities within India is completely lacking, leaving both government and private actors ill-equipped to make informed decisions. In this project, we develop India-wide, spatially-disaggregated estimates of temperature sensitivities for a range of socio-economically relevant outcomes (i.e. human mortality, labour supply, agricultural yields, crime rates, and conflict incidence). Our approach utilizes outcomes data (on mortality rates, work hours, agricultural yields, conflict incidence, and crime rates) from around the world to flexibly estimate temperature sensitivities as a function of a location’s adaptive capacities, which are proxied by measures of physical and socioeconomic characteristics. We generate high-resolution datasets of these measures and thereby extrapolate locally-relevant temperature sensitivities throughout India.

1 Introduction

Weather and climate shape India’s economy and society. Temperature and precipitation affect such diverse outcomes as human health, labour productivity, agricultural yields, crime, and conflict. While India is often discussed as one of the countries most vulnerable to future climate change impacts, contemporaneous weather shocks also pose clear and present dangers, as exemplified by the 2015 heat wave that claimed over 2,500 Indian lives (Ratnam et al., 2016).

The need for an Indian Climate Early Warning System stems from the fact that climate impacts are not just future eventualities but also current exigencies with immediate policy implications. As in any other policy domain, decision-makers deserve to be equipped with rigorous, data-driven analyses of damages from climatic extremes. However, at present, policy decisions— whether pertaining to long-term mitigation and adaptation or near-term responses to weather shocks— are made based upon incomplete, speculative, and highly aggregated estimates of damages. For instance, damage functions that link climate with human well-being are key inputs into Integrated Assessment Models of climate change. However, there is a general consensus that these damage functions no longer reflect the frontier of knowledge and that there exists a tremendous opportunity for updating them (Pindyck, 2013; Stern, 2013).¹ At the same time, much empirical evidence now exists for climate impacts (Carleton and Hsiang, 2016), but it is rarely, if ever, used collectively to inform either near-term or long-term economy-wide assessments. Policy design and business investments thus can only roughly take into account the role of climate today or in future decades.

In this project, we develop India-wide, spatially-disaggregated estimates of temperature sensitivities for a range of socio-economically relevant outcomes (i.e. human mortality, labour supply, agricultural yields, crime rates, and conflict incidence). These estimates are the product of two major research steps. First, we assemble a large, globally representative collection of historical

¹Moreover, the damage functions used in Integrated Assessment Models only represent aggregates over broad regions of the world.

outcomes data and daily climate data. The comprehensiveness of these data allows us to flexibly estimate plausibly causal temperature responses for each outcome. Second, we account for varying levels of adaptation across locations by directly modeling spatial heterogeneity in the temperature response functions. Specifically, we allow the response functions to vary as a function of long-run climate (e.g., Delhi is likely better adapted to high temperature days than Srinagar) (Barreca et al., 2015), income per capita (Hsiang and Narita, 2012), prevalence of irrigation (Schlenker and Roberts, 2009), and urban-ness (Burgess et al., 2014). These covariates were carefully chosen based on the literature, and new datasets were created to provide high-resolution global coverage. While the heterogeneity analysis is interesting in its own right, it is a key ingredient of the Early Warning System. By inserting local values of the covariates, it is possible to map out locally-relevant temperature-related vulnerabilities for each outcome across India.

The rest of this report is organized as follows: Section 2 describes the data collected on outcomes, climate, and covariates; Section 3 explains the methodological innovations whereby adaptation is modeled through spatial heterogeneity in temperature responses; Section 4 applies our methods to India, with maps of local temperature-related vulnerabilities throughout the country; Section 5 concludes.

2 Data Collection

Our approach makes use of three distinct types of data: a) Outcomes data on mortality rates, work hours, agricultural yields, crime rates, and conflict incidence; b) Historical daily temperature and precipitation data; and c) Data on state-level incomes, historical long-run average temperature, area equipped for irrigation, and population density. Each type of data is described in turn.

2.1 Outcomes Data

Mortality Mortality data represent 41 countries. Combined, our dataset covers mortality outcomes for 56% of the global population. Data are drawn from multiple, often restricted, national and international sources. Eight of these countries (United States, Brazil, India, China, Chile, Japan, Mexico, and France) were collected as individual samples, while the EU data was drawn from a single source which collects data for each member state. In some cases our data are the universe of mortality data in those countries, while in others (namely China and India) data are representative samples as no vital statistics registry system exists. All mortality datasets contain information on deaths per 100,000 population from all causes at an annual frequency. All datasets contain age-specific mortality rates at various resolutions with the exception of India which only provides all-age mortality.

Labor Individual-level work hours data are taken from large-sample time use surveys and labor force surveys that meet the data needs of this project. Specifically, in order to assign temperature exposure, we require that surveys identify the subnational geographic location of individuals and the exact calendar date(s) on which work hours are measured. Datasets meeting this requirement were acquired for 9 countries that represent nearly a third of global population. For the USA, Guatemala, Nicaragua, UK, France, Spain and India, we obtain daily work hours through personal time use diaries. For Brazil and Mexico, we obtain weekly work hours from labor force surveys. The datasets cover varying time spans ranging from from 1974 to 2015. Locations of individuals are identified at the municipality level in Brazil, Mexico, and Nicaragua; at the district, county, and department level in India, USA, and Guatemala respectively, and the region level in France, Spain, and the UK. Combining the separate country datasets yields a consolidated global dataset consisting of over 16 million observations (i.e. person-day or person-week). We restrict our attention to persons in the labor force who are between the ages of 15 and 65.

Agriculture Subnational rice production and area cropped data was collected for 31 countries with an unbalanced panel of production data spanning 1950 – 2015.² Yields were calculated as production divided by area planted where available, or divided by area harvested otherwise. Rice is

²The countries are Argentina, Brazil, Bolivia, Cambodia, Chile, China, Colombia, Dominican Republic, Ecuador, India, Indonesia, Japan, Laos, Malaysia, Myanmar, Nigeria, Philippines, Sri Lanka, Tanzania, Thailand, Vietnam, USA, and the countries of the European Union.

calorically the most important crop in India, accounting for nearly 40% of total caloric production (Cassidy et al., 2013).

Crime and Conflict We have collected data on the number and characteristics of both large-scale civil conflicts and interpersonal violent crimes at various spatial and temporal scales throughout the world. Guided by an extensive literature review, we have obtained the original data behind 21 distinct articles examining the link between climate and crime or conflict.³

The studies from which we have obtained data focus on a wide range of outcome variables, ranging from civil war and ethnic riots to homicide and assault. We therefore categorize these diverse outcomes into two groups: we call larger scale conflicts involving groups of individuals “intergroup conflict”, and smaller scale person-to-person crimes “interpersonal crime”. For intergroup conflict, we have coverage for nearly all countries in the world, although with much richer data in locations where conflict is common and heavily studied (e.g. sub-Saharan Africa). For interpersonal crime, we have subnational data for a smaller subset of countries. The geographic coverage of our data, along with spatial resolution, is shown in Figure 1.

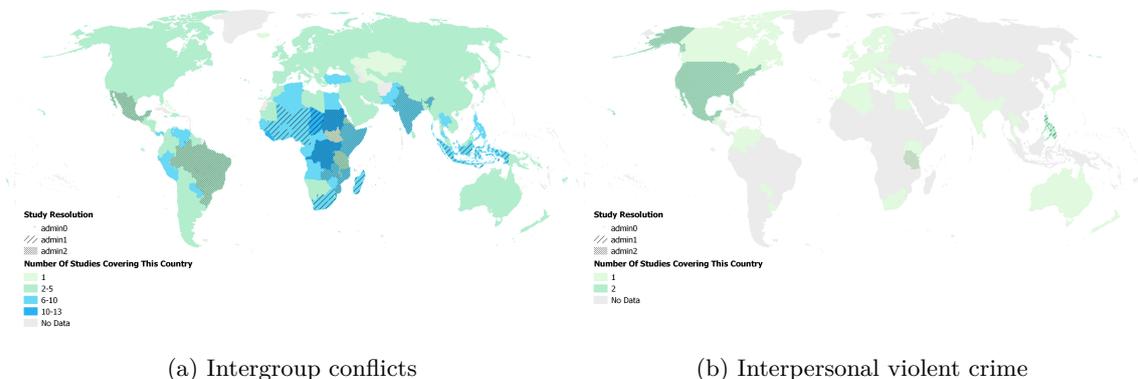


Figure 1: Data coverage and spatial resolution for conflict and crime

Because our data cover a wide range of outcome variables at distinct temporal and spatial scales, we standardize both climate (X) and outcome (Y) variables to account for heterogeneous average levels of conflict outcomes and heterogeneous temperature variances across these different scales. Our standardized variables are defined as:

$$\check{Y}_{it} = \frac{Y_{it}}{\bar{Y}_i} \quad (1)$$

$$\check{X}_{it} = \frac{X_{it} - \bar{X}_i}{\sigma_i(X_{it})} \quad (2)$$

where Y_{it} and X_{it} respectively denote conflict outcome variables and climate variables observed at location i and time period t . The objects \bar{Y}_i and \bar{X}_i respectively denote the averages of Y and X at location i over all periods t , while $\sigma_i(X_{it})$ is the standard deviation of climate variable X at location i over all periods t .

2.2 Climate Data

Historical For all sectors, we use gridded daily maximum, minimum, and/or average temperatures from two sources: The Berkeley Earth Surface Temperature (BEST) dataset, as well as the Global Meteorological Forcing Dataset for land surface modeling (GMFD) dataset. We aggregate grid cells into observations at administrative scales that match our outcome variables by taking a

³Data were sourced from the following 21 articles: Baysan et al. (2015), Bergholt and Lujala (2012), Bohlken and Sergenti (2010), Brückner and Ciccone (2011), Burke et al. (2015), Burke et al. (2009), Burke (2012), Burke and Leigh (2010), Caruso et al. (2016), Couttenier and Soubeyran (2014), Fetzner (2014), Fjelde and von Uexkull (2012), Hendrix and Salehyan (2012), Hidalgo et al. (2010), Kim (2016), Mares and Moffett (2016), Miguel (2005), Miguel et al. (2004), O’Loughlin et al. (2012), Ranson (2014), and Wetherley (2014).

weighted average of the grid cells that fall within a given region. The weights used will vary by sector. In agriculture the weights are cropped area, using the Center for Sustainability and the Global Environment (SAGE) 0.1 degree gridded product, such that regional weather variables are crop-area-weighted exposure measures. For mortality and labor, we use population weights, where gridded population comes from the Gridded Population of the World (GPW) version 4 dataset. For crime and conflict, we use simple area weights.

Each sector uses distinct weather variables to capture appropriate climate exposure. These variables are summarized in Table 1. For agriculture, growing degree days (GDD) were calculated using a sinusoidal fit to daily T_{min} and T_{max} temperatures. This approach estimates a smooth fit to temperatures over time, allowing for sub-hourly measures of crop exposure to extreme temperatures and avoiding underestimating exposure to the extremes of the temperature distribution which would be associated with time-averaged measures. The GDD lower and upper breakpoints were searched for best fit to the data and set at 10°C and 31°C, respectively. "Killing degree days" (KDD) were calculated as growing degree days with a lower bound of 31°C and no upper bound, under the same sinusoidal fit to daily T_{min} and T_{max} temperatures. In mortality and labor, flexible polynomials in daily average (mortality) and daily maximum (labor) temperatures were estimated at the grid cell level before aggregation to the regional level. In conflict, we rest heavily on the expansive existing literature and use the temperature data included in each individual study from which we obtained original data, implying that our temperature variables can include both averages as well as maximum temperatures.⁴

In all sectors, we use gridded daily precipitation data from GMFD, or gridded monthly data from the University of Delaware, calculating a appropriately-weighted regional cumulative measures of total rainfall.

2.3 Covariate Data

Data on the following covariates are used to explain spatial heterogeneity in the temperature responses of the outcomes. Exact details on how the covariates are constructed for each outcome can be found in Table 2.

Temperature Historical long-run average temperature is calculated from the temperature datasets used in the analysis and described above. The average is taken over the years that correspond to the outcome data. This is calculated at the 1st-level administrative division (i.e. ADM1) in order to match the level at which the administrative income data is available.

Sub-national Incomes Sub-national incomes (in constant 2005 dollars adjusted for purchasing power parity) are obtained from [Gennaioli et al. \(2012\)](#), which reports sub-national incomes gathered from administrative data around the world. These data are typically not annual, and are drawn from census data. We match incomes in the cross-section to ADM1 units in our sample based on the sample period for each country.

Area Equipped for Irrigation To measure the cropland area equipped for irrigation (AEI), we used the Global Map of Irrigated Areas from [Siebert et al. \(2013\)](#), which shows the percentage of grid cell area equipped for irrigation around the year 2005. The Global Map of Irrigated Areas is provided at 5 minute grid cells.

Population density Population density is derived from Landsat. We use this high-resolution gridded dataset to create a measure of urban-ness, distinct from population density in an administrative region. As average population density across any of our regions is highly dependent on the value of area in the denominator, this measure will not capture how urban an area is. This is particularly true in cases where large administrative regions comprise unpopulated desert and highly concentrated urban centers. We transform this into a measure of urban-ness by calculating the population-weighted population density, or the population density that *an average person perceives* in an administrative region.

⁴We do not use our own temperature data because of the diversity of contexts and hence appropriate climate measures studied in these different conflict contexts. Because we standardize all analysis in conflict (see Equations 1 and 2), these differences are accounted for.

3 Methodological Innovations

Our method involves 3 steps. We first estimate a dose-response relationship between temperature and each outcome using all global data for that outcome. Second, we explore how this relationship varies across locations. Spatial (i.e. cross-sectional) heterogeneity in the dose-response relationship is empirically modeled as a function of location-specific factors, which, depending on the outcome, may include per-capita GDP, long-run temperature, urban-ness, and percent of cropped area that is irrigated. These covariates are meant to capture dimensions along which adaptation to extreme temperatures might occur. The third step consists of interpolation based on the heterogeneity estimates. By inserting the appropriate values of the covariates, a dose-response relationship, reflective of adaptive capabilities, can be characterized for any location.

3.1 Step 1: Dose-response Function Estimation

The global dose-response function flexibly models the relationship between temperature in degrees Celsius (T) and outcome (Y). Each outcome variable is measured for some unit i at time t , with the set of outcomes and their spatial and temporal resolutions detailed in Table 1. The explanatory variable of interest (i.e. temperature) is modeled through a non-linear function $f(T)$.

The regression specifications take the following general form:

$$Y_{i,t} = f(T_{i,t}) + Controls_{i,t} + FixedEffects + \varepsilon_{it}. \quad (3)$$

The exact form for the function f varies by outcome, as do the control variables and fixed effects. The particulars for each outcome are listed in Table 1. The variable ε_{it} denotes the remaining idiosyncratic error.

With the exception of crime and conflict outcomes, the response of an outcome Y at any given temperature is expressed as the difference between the predicted value of Y at that temperature and the predicted value at some reference temperature T_{ref} , where the exact value of T_{ref} varies by outcome (see Table 1).⁵ For example, the response of Y on a day where the temperature is x degrees Celsius is:

$$Response(x) = f(x) - f(T_{ref}) \quad (4)$$

For the crime and conflict outcomes, the estimation employs a temperature exposure measured in z-scores. Therefore, the function f itself captures a response that is relative to the sample average temperature. Thus, the estimated predicted response of these outcomes at a temperature z-score x is simply $Response(x) = f(x)$.

In the next step, we exploit data on location-specific covariates to model how differential adaptation affects the shape of the response function.

3.2 Step 2: Interpolation Surface Estimation

While an estimate of equation (3) yields a global average response function for a given outcome, it does not capture how the response function varies across locations due to differential adaptation opportunities. In order to address this issue, we develop a model of interactions to characterize heterogeneity in the response. Specifically, we augment equation (3) to include interactions between temperature and other factors (i.e. location-specific covariates) at the first administrative level (i.e. state or province). The factors are as follows, with the exact interactions described in Table 2:

1. *Income* measured as the natural log of per-capita GDP over some representative year(s). See Table 2 for the years used for each outcome.
2. *Long – run Temperature* measured for each outcome as described in Table 2.
3. *Percent cropped area irrigated* (for agricultural yield outcome only).
4. *Population density* (for crime and conflict only), measured as the natural log of population density in the year 2000.

⁵In other words, the response to a temperature of T_{ref} is normalized to zero.

Table 1: Dose-Response Estimation by Sector

| Outcome | Spatial Resolution | Temporal Resolution | Temperature Functional Form | Fixed Effects | Controls | T_{ref} |
|----------------------------|--|--------------------------|--|--|---|-----------|
| Deaths per 100,000 persons | District (Adm2) | Annual | Fourth-order polynomial in daily average temperature (aggregated annually) | District \times age, Country-year \times age | Quadratic precipitation (Country-specific coefficients) | 20°C |
| Minutes Worked | Person-level data, District (Adm2) of residence identified | Daily | Fourth-order polynomial in daily max temperature | District, Day-of-week, State \times mo. of sample | Age, sex, household size | 27°C |
| Yield | District (Adm2) | Annual | Piecewise linear, Sinusoidal intraday fit | District, country-year | Quadratic precipitation | 31°C |
| Violent Crime Incidence | Various (grid cell to country) | Various (daily-annual) | Linear in z-score of original study temperature exposure | Study \times region, study \times region \times month, study \times year | Quadratic precipitation | N/A |
| Conflict Incidence | Various (grid cell to country) | Various (monthly-annual) | Linear in z-score of original study temperature exposure | Study \times region, study \times region \times month, study \times year | Quadratic precipitation | N/A |

“Region” refers to the spatial resolution of the observation (e.g. county, state, country).

We chose these factors because they plausibly capture differential adaptive responses across locations. Prior literature has emphasized the adaptive significance of average climate (Barreca et al., 2015), income per-capita (Hsiang and Narita, 2012), urban-ness (Burgess et al., 2014), and, for agricultural yields, irrigation (Schlenker and Roberts, 2009). For example, higher per-capita GDP entails greater capability to invest in protective measures such as air-conditioning, which may dampen the mortality dose-response relationship. Furthermore, the dose-response relationship may also depend on long-run exposure to extreme temperatures as places with greater previous exposure may differ in their adaptive behaviors. Additionally, irrigation may mitigate the harmful effects of extreme temperatures on agriculture yields.

For each outcome, we parameterize the heterogeneity in the function $f(T)$ due to the factors (see Table 2 for specific parameterizations by outcome), and thereby estimate $f(T)$ as conditional on the factors. Letting \mathbf{z} denote the vector of factors, the conditional function is thus $f(T|\mathbf{z})$. Estimating the object $f(T|\mathbf{z})$ paves the way for the third step of the analysis— predicting dose-response values at a local level.

Table 2: Factor Interactions by Sector

| Outcome | Income | Long-run Temperature | Other Covariate |
|----------------------------|--|---|----------------------------------|
| Deaths per 100,000 persons | $\log(GDP_{pc})$ in sample period ^a | Average annual temperature in sample period | N/A |
| Minutes Worked | $\log(GDP_{pc})$ in representative year ^b | Average annual heating and cooling degree days ^c | N/A |
| Yield | $\log(GDP_{pc})$ in year of observation ^a | Growing season average T_{max} in sample period | Cropland share irrigated |
| Violent Crime Incidence | $\log(GDP_{pc})$ in sample period ^a | Average annual temperature in sample period | $\log(popdens)$ in sample period |
| Conflict Incidence | $\log(GDP_{pc})$ in sample period ^a | Average annual temperature in sample period | $\log(popdens)$ in sample period |

^a Subnational income data are linearly interpolated between the years provided in Gennaioli et al. (2012).

^b For each dataset, the representative year is the year closest to the midpoint of the survey year(s) for which subnational income data are available from Gennaioli et al. (2012). The representative years range from 1995 to 2005.

^c Heating degree days in a year are defined as the cumulative deviations of daily maximum temperatures from a benchmark of 10° C, over all days where the maximum temperature fell below 10° C. Formally, heating degree days in year y are defined as $\sum_{d \in y} |T_{max}^d - 10| * I_{T_{max}^d < 10}$, where T_{max}^d is the maximum temperature on day d and $I_{T_{max}^d < 10}$ is an indicator variable equal to one if $T_{max}^d < 10$. Cooling degree days are defined similarly over days where the daily maximum temperature exceeds 30° C, i.e. cooling degree days in year y are $\sum_{d \in y} |T_{max}^d - 30| * I_{T_{max}^d > 30}$. We use the average value of both variables over 1970-2013. The heating degree days measure is modeled to influence the response for days where the temperature is below 27° C. The cooling degree days measure is modeled to influence the response for days where the temperature is at or above 27° C.

3.3 Step 3: Interpolating Dose-response Relations

Estimating $f(T|\mathbf{z})$ makes it possible to interpolate a response function to any location of the world where the values of the covariates are known. We predict dose-response values at the *impact region* level throughout India, where we have created data to identify the values of each factor. *Impact regions* are a set of boundaries created by us for the purpose of extrapolating temperature sensitivities derived from our estimation. Impact regions are constructed such that they are identical to existing administrative regions or are a union of a small number of administrative regions.⁶ There are 2300 impact regions in India, the boundaries of which are depicted in Figure 2.

⁶We use the Global Administrative Region dataset (Global Administrative Areas, 2012) to delineate boundaries for impact regions, but require fewer than the approximately 295,000 spatial units present in that dataset. We

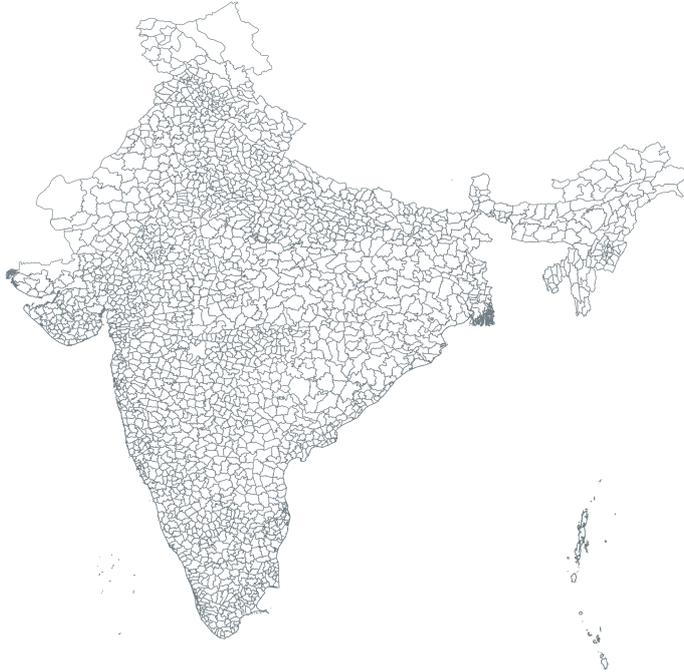


Figure 2: Impact Regions in India

To obtain per-capita income at the impact region level, we allocate national GDP based on the impact region’s mean frequency of nightlights between 1992 and 2013 (Henderson et al., 2012).⁷ All other factors are obtained at the impact region level by aggregating gridded data from the sources described in Section 2.3.

Formally, for outcomes other than crime and conflict, the estimated, predicted response at a temperature x (relative to T_{ref}) in an impact region with covariate vector \mathbf{z}_i is:

$$Response(x|\mathbf{z}_i) = f(x|\mathbf{z}_i) - f(T_{ref}|\mathbf{z}_i). \quad (5)$$

For the crime and conflict outcomes, which are estimated using a z-score measure of temperature exposure, the estimated predicted response at a temperature z-score x in an impact region with covariate vector \mathbf{z}_i is simply $Response(x|\mathbf{z}_i) = f(x|\mathbf{z}_i)$.

In the next section, we map these estimates of predicted responses for each outcome at extreme temperature values.

4 Applications

To demonstrate the application of our methodological innovations, we explore the local impacts of high temperature days on each outcome throughout India. The subsequent maps illustrate impacts that reflect adaptive capabilities (i.e. income, long-run temperature, irrigation, and urban-ness) as of 2010.

therefore agglomerate the spatial units to create a set of 24,378 impact regions globally 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 across regions with respect to population, and internally consistent with respect mean temperature, diurnal temperature range, and mean precipitation.

⁷Specifically, we first estimate a linear relationship between the z-score of GDP per capita and z-score of nightlight frequency for the sample of US counties in 2011, where both these variables are available. We then apply these coefficient estimates to nightlights density measurements to predict GDP per capita at the impact region level in India.

4.1 Mortality

We use the estimated response surface for the world to project a high-resolution temperature sensitivity relationship for all of India. Figure 3 displays the sensitivity to temperature for four different daily temperatures with reference to the sensitivity at 20°C.

To place these estimates in context, we estimate the effect of raising the entirety of India to an average daily temperature of 32°C. While this is an unrealistic temperature increase for the whole of India, it is useful to give some intuition for the size of the effects. We find that a single day at 32°C would raise mortality rates by 0.66 *deaths per 100,000* in the year 2010⁸. Taking a population in that year as 1.23 billion people, this would result in 8124 extra deaths in India for that day.

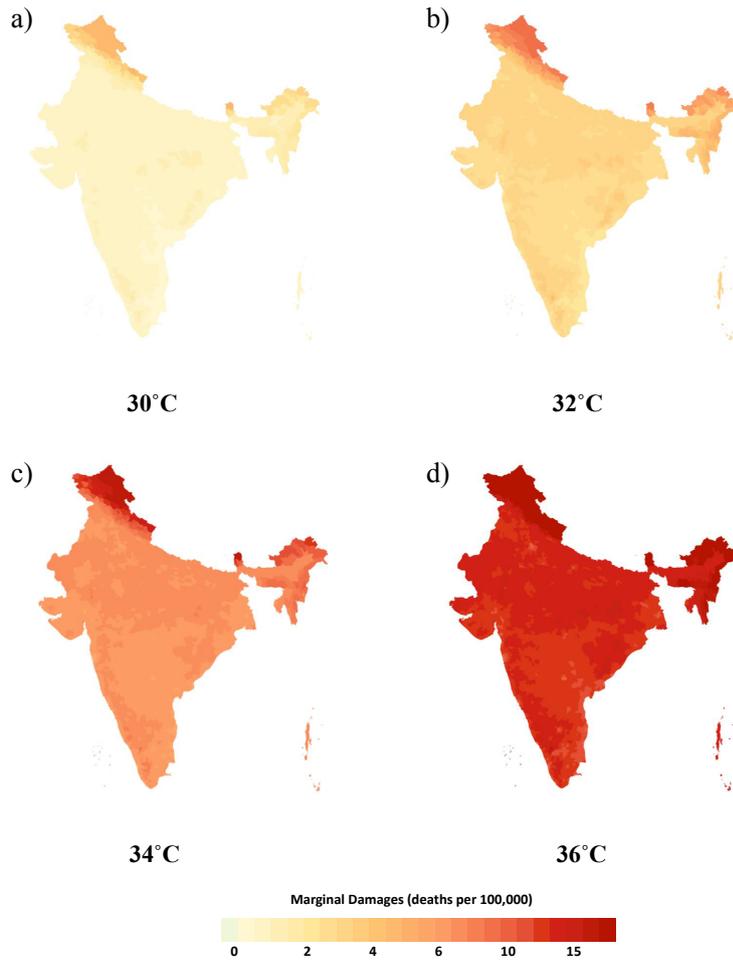


Figure 3: Impact of extreme temperatures on mortality rates in India. Each panel shows the impact of a single day at the stated temperature relative to a day at 20°C. Figure shows result for over 65 year olds.

4.2 Labor

Figure 4 maps a high-resolution temperature sensitivity relationship for work hours. Each panel displays the average lost minutes per worker on a day with a particular maximum temperature (35°C, 37°C, 40°C, and 43°C), relative to a baseline day with maximum temperature of 27°C. The maps reveal sharp increases in lost work time at higher temperatures.

⁸We obtain this number by taking the age and population averages of the mortality rate response for each Indian district.

To place our estimates in perspective, a day where the maximum temperature is 37°C throughout India is associated with 101 million lost work hours nationwide.⁹ This corresponds to over 12.5 million full-time equivalent workers if one assumes an eight-hour workday. Similarly, a 40°C day is associated with a nationwide loss of 189 million lost work hours, or over 23 million lost full-time equivalent workers.

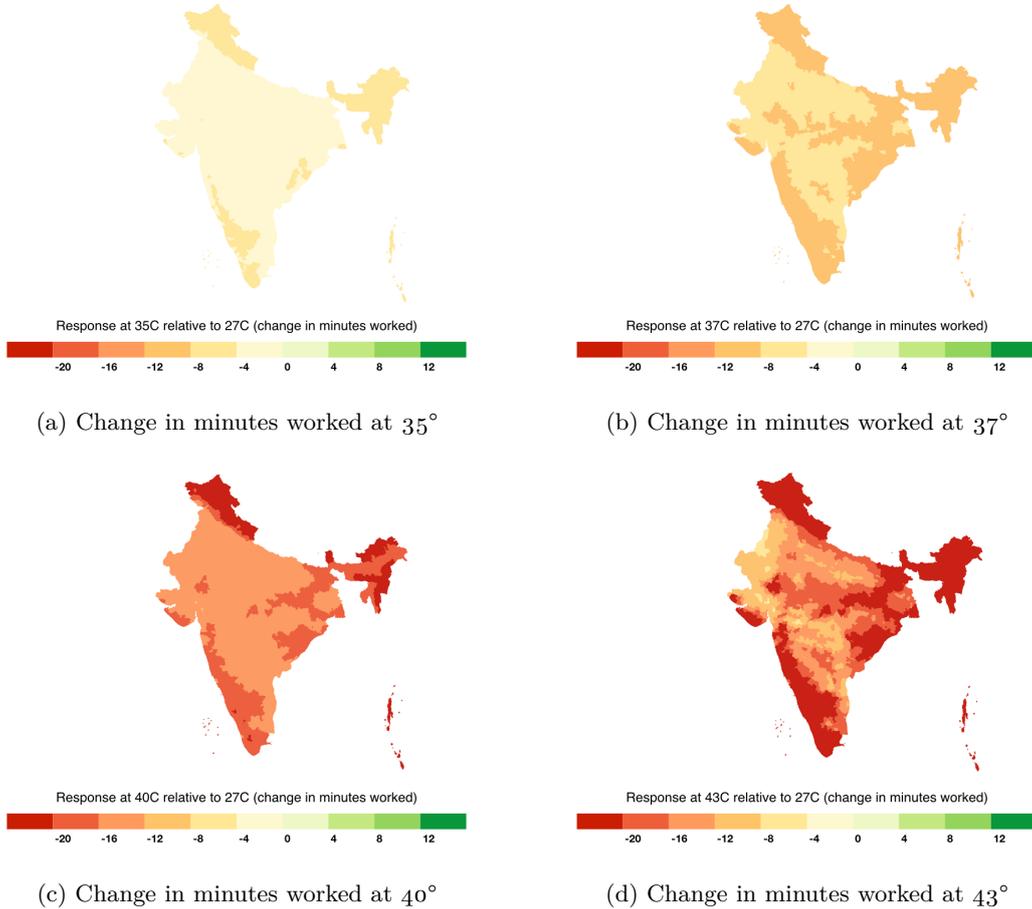


Figure 4: Impact of Extreme Heat on Labor Supply. Each panel depicts the average lost work minutes per worker associated with a given daily maximum temperature, relative to the minutes worked at the reference daily maximum temperature of 27°C .

4.3 Agriculture

In estimating the rice response surface, we find two forms of climate adaptation. Farms that are more frequently exposed to extreme heat, measured via exposure to long-run T_{max} , exhibit a muted negative response of yields to extreme heat shocks. Likewise, high levels of irrigation also mute the adverse effects of extreme heat shocks, presumably through reducing the water-stress effects of extreme heat shocks on crops.

The rice response surface was interpolated for all of India using impact area covariate data, as described previously. The results of this interpolation are displayed in Figure 5. The first chart depicts the yield loss (in log points) associated with 24 hours of constant exposure to a temperature of 35°C , relative to the yield at the estimated optimal rice growing temperature of 31°C . These are the projected responses, allowing for the adaptive effects of climate adaptation and irrigation described above. So, a value of -0.03 would indicate a 3% yield loss relative to the ideal

⁹This is obtained by multiplying the estimated average lost minutes per worker in each impact region by the number of workers in that impact region, and then summing the lost work time across all impact regions. For this calculation, we assume the number of workers in each impact region is equal to the population times the India-wide labor-force participation rate of 57.1 percent in 2009-2010 (International Labour Organization, 2016).

rice temperature, conditional on the long-run climate and current level of irrigation in any given impact region. Regions with a lesser estimated yield loss will have higher levels of irrigation and/or higher levels of long-run T_{max} . The second chart in Figure 5 is identical to the first, except the exposure depicted is for 24 hours of constant temperature at 40°C, instead of 35°C.

It should be noted that these effects, as modeled, are additive. Take, for example, a region which is estimated to exhibit a 6% loss of yields under exposure to 24 hours of 40°C temperatures. If that region is exposed to 72 hours of 40°C temperatures over the course of a growing season, yield losses will be projected to be 18% on average, relative to the counterfactual of optimal growing temperature during those 72 hours.

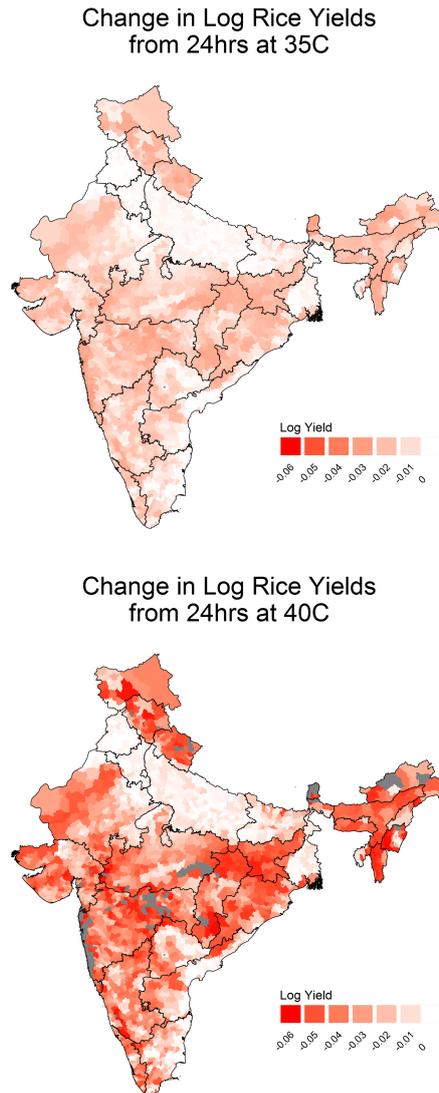


Figure 5: Rice response interpolation for impact regions in India. The first chart depicts the yield loss (in log points) associated with 24 hours of constant exposure to a temperature of 35°C, relative to the yield at the estimated optimal rice growing temperature of 31°C. The second chart depicts the same, except for 24 hours of constant exposure to a temperature of 40°C.

4.4 Crime and Conflict

For interpersonal crime and intergroup conflict independently, we use the interpolation method described above to generate predicted temperature sensitivities throughout India. The results of this interpolation are displayed in Figure 6. In panel (a), the map shows the predicted increase in civil conflict incidence caused by a one standard deviation increase in temperature exposure. In panel (b), the map shows the same value, but for interpersonal violent crime. Note that as above for other sectors, these are the projected responses, allowing for the adaptive effects of climate adaptation and income described above.

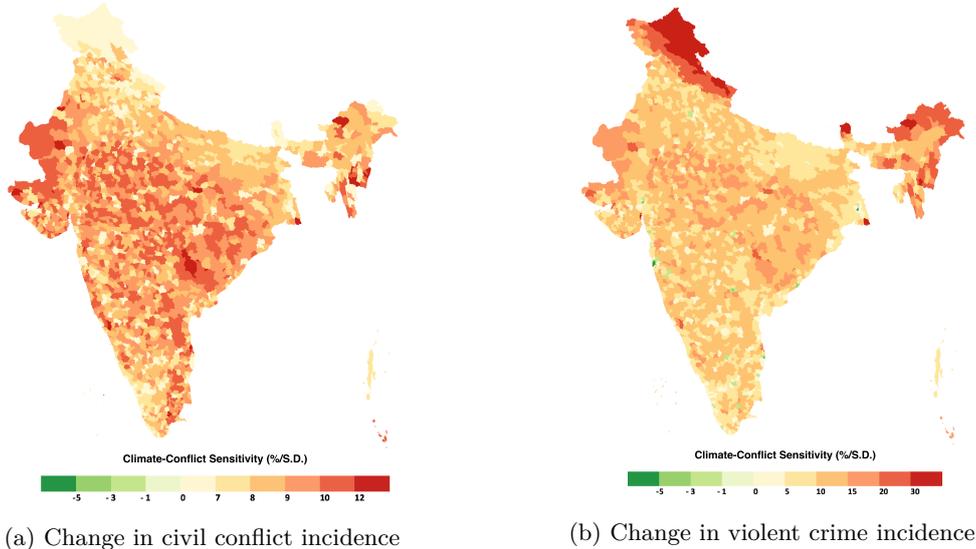


Figure 6: Impact of temperature on crime and conflict. Panel (a) depicts the % change in conflict per standard deviation (σ) increase in temperature, where temperature exposure is defined using a range of variables, from monthly maximum temperature to annual average temperature. Panel (b) shows the same results, but for interpersonal violent crimes.

Across all of India, the average baseline number of annual conflicts is 7,900, according to the Armed Conflict Location and Event Data Project (ACLED). Our estimates shown in panel (a) of Figure 6 predict that a one standard deviation increase in temperature will cause an additional 697 conflicts across India; this corresponds to an average effect of an 8.8% increase in conflict per standard deviation increase in temperature, although this average masks substantial spatial heterogeneity shown in panel (a).

While both interpersonal crimes and larger scale civil conflicts tend to increase with warming temperatures, the magnitudes of these effects and their spatial patterns have been shown to differ substantially (Burke et al., 2015). Consistent with this previous evidence, we find that predicted temperature sensitivities for interpersonal crime exhibit distinct behavior to those estimated for intergroup conflict. The interpolation results for crime are shown in panel (b) of Figure 6. As for intergroup conflict, the map in panel (b) shows the predicted increase in violent crime incidence caused by a one standard deviation increase in temperature exposure. Across all of India, we predict that the average effect of a one standard deviation increase in temperature exposure is 11.6% for violent crime, a significantly higher average value than that for intergroup conflict. Applying this estimate to the case of homicides, we estimate that the average number of homicides per year, currently at 3.5 per 100,000 (Mares and Moffett, 2016), will rise by .411 per 100,000 in response to a one standard deviation increase in temperature. However, the spatial pattern of these deaths is distinct from that shown in panel (a); for interpersonal crime, the highest sensitivity locations are the relatively cool, northern regions. By contrast, for intergroup conflict the highest sensitivity locations tend to be in the central regions of India.

Together, these findings demonstrate how our methodology can harness large amounts of diversely-collected data to generate hyper-local predictions of climate impacts. In particular, these crime and conflict climate impact maps have the potential to be exceptionally valuable to policymakers responsible for preventing and combating outbreaks of violence which may be linked to fluctuations in the climate.

5 Conclusion

While it is widely acknowledged that India is vulnerable to near-term weather shocks as well as long-term climate change, there exist no comprehensive analyses on how these vulnerabilities matter in specific locations and economic sectors. As a result, India's government and private decision-makers have been informationally impaired in anticipating the economic and social repercussions of extreme weather. The methodological innovations of this project are designed to remedy this deficiency and yield unprecedented, hyper-local insights into climate risks for multiple impact sectors— mortality, labour productivity, agricultural yields, crime, and conflict.

Our approach utilizes outcomes data (i.e. data on mortality rates, work hours, agricultural yields, conflict incidence, and crime rates) from around the world to flexibly estimate temperature sensitivities as a function of a location's adaptive capacities, which we model through measures of income, long-run climate, irrigation, and urban-ness. We have developed high-resolution, global datasets of these measures, thus making it possible to extrapolate locally-relevant temperature sensitivities at any given location in India. Our findings suggest that across India, extreme heat exposure leads to substantial loss of human life, work hours, and rice yields, and sharp increases in crime and conflict.

In addition to shedding light on local temperature sensitivities of various outcomes, the methods developed in this project also have the potential to be used as a near-term forecasting tool. Such a tool can equip policymakers with valuable information for better allocating scarce resources to areas of immediate concern.

References

- A. Barreca, K. Clay, O. Deschênes, M. Greenstone, and J. S. Shapiro. Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004. *The American Economic Review*, 105(5):247–251, 2015.
- C. Baysan, M. Burke, F. Gonzalez, S. Hsiang, and E. Miguel. Economic and non-economic factors in violence: Evidence from organized crime, suicides and climate in mexico. *Department of Agricultural and Resource Economics, University of California, Berkeley*, 2015.
- D. Bergholt and P. Lujala. Climate-related natural disasters, economic growth, and armed civil conflict. *Journal of Peace Research*, 49(1):147–162, 2012.
- A. T. Bohlken and E. J. Sergenti. Economic growth and ethnic violence: An empirical investigation of hindu—muslim riots in india. *Journal of Peace research*, 47(5):589–600, 2010.
- M. Brückner and A. Ciccone. Rain and the democratic window of opportunity. *Econometrica*, 79(3):923–947, 2011.
- R. Burgess, O. Deschenes, D. Donaldson, and M. Greenstone. The unequal effects of weather and climate change: Evidence from mortality in india. *Cambridge, United States: Massachusetts Institute of Technology, Department of Economics. Manuscript*, 2014.
- M. Burke, S. M. Hsiang, and E. Miguel. Climate and conflict. *Annu. Rev. Econ.*, 7(1):577–617, 2015.
- M. B. Burke, E. Miguel, S. Satyanath, J. A. Dykema, and D. B. Lobell. Warming increases the risk of civil war in africa. *Proceedings of the national Academy of sciences*, 106(49):20670–20674, 2009.
- P. J. Burke. Economic growth and political survival. *The BE Journal of Macroeconomics*, 12(1), 2012.
- P. J. Burke and A. Leigh. Do output contractions trigger democratic change? *American Economic Journal: Macroeconomics*, 2(4):124–157, 2010.
- T. A. Carleton and S. M. Hsiang. Social and economic impacts of climate. *Science*, 353(6304):aad9837, 2016.
- R. Caruso, I. Petrarca, and R. Ricciuti. Climate change, rice crops, and violence: Evidence from indonesia. *Journal of Peace Research*, 53(1):66–83, 2016.
- E. S. Cassidy, P. C. West, J. S. Gerber, and J. A. Foley. Redefining agricultural yields: from tonnes to people nourished per hectare. *Environmental Research Letters*, 8(3):034015, 2013.
- M. Couttenier and R. Soubeyran. Drought and civil war in sub-saharan africa. *The Economic Journal*, 124(575):201–244, 2014.
- T. Fetzer. Can workfare programs moderate violence? evidence from india. 2014.
- H. Fjelde and N. von Uexkull. Climate triggers: Rainfall anomalies, vulnerability and communal conflict in sub-saharan africa. *Political Geography*, 31(7):444–453, 2012.
- N. Gennaioli, R. La Porta, F. Lopez-de Silanes, and A. Shleifer. Human capital and regional development. *The Quarterly journal of economics*, 128(1):105–164, 2012.
- Global Administrative Areas. Gadm database of global administrative areas, version 2.0. URL: <http://www.gadm.org> [accessed 201-03-24], 2012.
- J. V. Henderson, A. Storeygard, and D. N. Weil. Measuring economic growth from outer space. *The American Economic Review*, 102(2):994–1028, 2012.
- C. S. Hendrix and I. Salehyan. Climate change, rainfall, and social conflict in africa. *Journal of peace research*, 49(1):35–50, 2012.

- F. D. Hidalgo, S. Naidu, S. Nichter, and N. Richardson. Economic determinants of land invasions. *The Review of Economics and Statistics*, 92(3):505–523, 2010.
- S. M. Hsiang and D. Narita. Adaptation to cyclone risk: Evidence from the global cross-section. *Climate Change Economics*, 3(02):1250011, 2012.
- International Labour Organization. India Labour Market Update (July 2016). *ILO Country Office for India*, 2016.
- N. K. Kim. Revisiting economic shocks and coups. *Journal of Conflict Resolution*, 60(1):3–31, 2016.
- D. M. Mares and K. W. Moffett. Climate change and interpersonal violence: a “global” estimate and regional inequities. *Climatic change*, 135(2):297–310, 2016.
- E. Miguel. Poverty and witch killing. *The Review of Economic Studies*, 72(4):1153–1172, 2005.
- E. Miguel, S. Satyanath, and E. Sergenti. Economic shocks and civil conflict: An instrumental variables approach. *Journal of political Economy*, 112(4):725–753, 2004.
- J. O’Loughlin, F. D. Witmer, A. M. Linke, A. Laing, A. Gettelman, and J. Dudhia. Climate variability and conflict risk in east africa, 1990–2009. *Proceedings of the National Academy of Sciences*, 109(45):18344–18349, 2012.
- R. S. Pindyck. Climate change policy: What do the models tell us? *Journal of Economic Literature*, 51(3):860–872, 2013.
- M. Ranson. Crime, weather, and climate change. *Journal of environmental economics and management*, 67(3):274–302, 2014.
- J. Ratnam, S. K. Behera, S. B. Ratna, M. Rajeevan, and T. Yamagata. Anatomy of indian heatwaves. *Scientific reports*, 6, 2016.
- W. Schlenker and M. J. Roberts. Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences*, 106(37):15594–15598, 2009.
- S. Siebert, V. Henrich, K. Frenken, and J. Burke. Update of the digital global map of irrigation areas to version 5. *Rheinische Friedrich-Wilhelms-Universität, Bonn, Germany and Food and Agriculture Organization of the United Nations, Rome, Italy*, 2013.
- N. Stern. The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. *Journal of Economic Literature*, 51(3):838–859, 2013.
- E. Wetherley. Typhoons and temperature impact crime rates: evidence from the philippines. 2014.

The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about our work on our website
www.theigc.org

For media or communications enquiries, please contact
mail@theigc.org

Subscribe to our newsletter and topic updates
www.theigc.org/newsletter

Follow us on Twitter
[@the_igc](https://twitter.com/the_igc)

Contact us
International Growth Centre,
London School of Economic and Political Science,
Houghton Street,
London WC2A 2AE

IGC

**International
Growth Centre**

DIRECTED BY



FUNDED BY



Designed by soapbox.co.uk