

Technical appendix: Climate impacts

Solomon M. Hsiang^{a,b,*}, Amir Jina^{a,c}, James Rising^c

^a*University of California, Berkeley, CA*

^b*National Bureau of Economic Research, Cambridge, MA*

^c*Columbia University, New York, NY*

1. Introduction

Recent years have seen an explosion of empirical research leading to an unprecedented advance in our knowledge of the interaction of society and the environment. As the amount of research continues to grow, it becomes more important that studies be used to collectively answer societally important questions. The approach of this report towards environmental assessment is an effort to holistically integrate this research. Our vision was to create a system to not only incorporate recent discoveries, but also to be updated with new research and new findings as they become available. We present below a flexible, open-source, and adaptive system to combine our best estimates of environmental impacts, allowing us to constantly learn the broad societal effects from an evolving body of research. The current approach is not limited to climate, however, as it can be extended to project many different types of impacts (e.g., from policy changes). The tools we have designed can become a central hub enabling researchers to collaborate on a larger body of socially important research.

We identify and employ a meta-analytical approach (described in section 2) that draws on Bayesian methods commonly used in medical research and previously implemented in *Hsiang et al.* (2013b). Using these techniques, we design an open-source tool which can update aggregated dose-response functions in real time as new research becomes available (fig. B2). With the method of meta-analysis in place, we then identify a group of rigorous studies across a number of climate-impacted sectors. Combining the dose-response functions from individual studies (detailed in section 3), we generate a series of aggregate response functions for each sector. Finally, to understand the impact of climate upon each sector, we take the product of our response functions and the downscaled physical climate projections described in Technical Appendix A, giving us partial equilibrium impacts out to the end of the current century (described in section 4). These impacts are then used as input to computationally model the general equilibrium effects of changes in climate.

2. Meta-Analysis Approach

The empirical impact functions are treated as conditional distributions, conditioned on weather variables such as mean temperature and precipitation. This representation facilitates meta-analysis, and also captures the range of uncertainty in the empirical estimates. Each distribution is evaluated at a given quantile only when it is applied to data, as described in section 4.

2.1. Hierarchical Bayesian Modeling

The impact estimates that combine results from more than one study apply a Bayesian hierarchical model structure, as described by *Gelman et al.* (2013). This approach simultaneously estimates a distribution of possible underlying effect sizes, as well as a degree of “partial-pooling”. If the individual study estimates are consistent with a single underlying effect, their estimates are pooled to accurately estimate the effect. However, if the study estimates are inconsistent with each other, the underlying effect is estimated to be only loosely informed by each study, which is considered to have its own idiosyncratic effect.

Consider a collection of impact functions, $f_i(\beta_i|T)$, for $i \in \{1, \dots, N\}$ indexing independently published results. Here, $f_i(\beta_i|T)$ is a probability distribution for β_i conditioned on a weather variable T . We wish

*Corresponding author

Email address: shsiang@berkeley.edu (Solomon M. Hsiang)

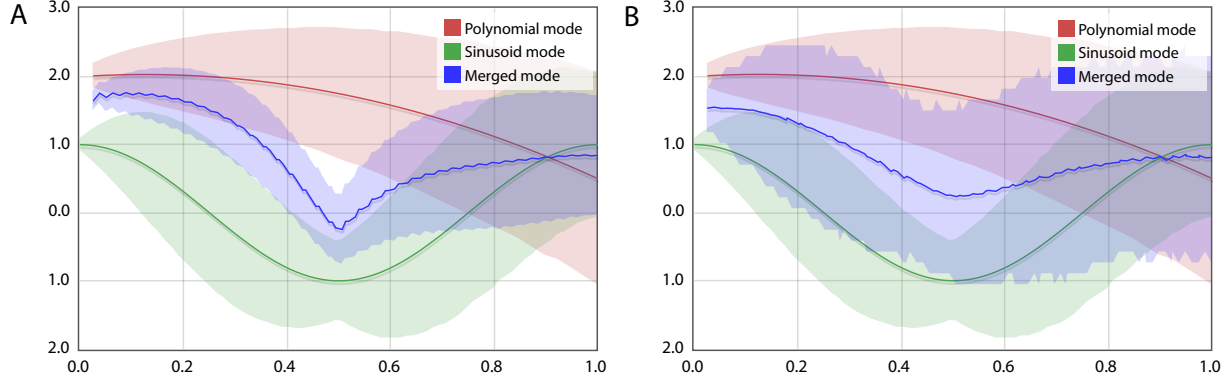


Figure B1: **A)** Pooled and **B)** Bayesian hierarchical estimates for a constructed polynomial and sinusoidal response function. The pooled distribution is calculated as $p(\hat{\beta}|x) = p(\cap_{i=1}^N \hat{\beta} = \beta_i | x) \propto \prod_{i=1}^N p(\beta_i | x)$, where $p(\beta_i | x)$ is the conditional distribution for either the polynomial or sinusoidal response function at a given value of x . The 95% confidence interval of the pooled result does not overlap with the individual estimates when they are far from each other. The confidence intervals on $\hat{\beta}$ are wider, reflecting the uncertainty in resolving the two estimates.

to combine these estimates into a single conditional distribution, $g(\hat{\beta}|T)$, where $\hat{\beta}$ is called the “hyper-parameter”. We treat each value of T independently, so we will write these functions as $f_i(\beta_i)$ and $g(\hat{\beta})$.

Under hierarchical merging, the conditional parameter distributions are required to be Gaussian distributions, and below, Gaussian parameter estimates errors are used in all applicable response functions. The governing equations are,

$$\begin{aligned}\theta_i &\sim \mathcal{N}(\hat{\beta}, \tau^2) \\ \beta_i &\sim \mathcal{N}(\theta_i, \sigma_i^2)\end{aligned}$$

where β_i is a measured parameter, corresponding to a true (unobserved) parameter θ_i which characterizes the response for study i . σ_i^2 is the standard error of β_i . We are interested in $\hat{\beta}$, the underlying hyper-parameter, and τ^2 , the variance between models. We apply non-informative priors to $\hat{\beta}$ and τ . That is, $p(\hat{\beta}) \propto 1$ and $p(\tau) \propto 1$. The values of β_i and σ_i^2 are provided by the published studies, and the rest of the parameters are simultaneously estimated.

An analytic solution exists for how to generate draws from the posterior distribution of this hierarchical model, and is described in chapter 5 of *Gelman et al.* (2013). We approximate the posterior by producing draws and constructing a histogram for each conditional distribution.

Figure B1 shows an example of how pooled and Bayesian hierarchical results differ for a combination of two simple impact functions.

2.2. Distributed Meta-Analysis System

To support the management of empirical results, the meta-analysis combination process, and the application of these results to data, we constructed a new tool called DMAS, the Distributed Meta-Analysis System¹. This system currently operates on a dedicated server housed at the University of California, Berkeley. The heart of DMAS is a database of results that can be easily recombined into many different meta-analyses. This database is designed to be expanded in a decentralized manner, by “crowd-sourcing” from scientists working independently to detail their empirical findings². By combining the efforts of many researchers throughout various academic disciplines in an Internet research community, the capacity of scientists to maintain up-to-date empirical relationships increases drastically. Below, we detail the numerous features which are available to promote academic exchange on DMAS. Unlike many crowd-sourcing projects, the necessary vetting of this information is made much easier by connecting each estimate with published literature. This connection to academic journals further supports the construction of comprehensive meta-analyses.

¹Available online at <http://dmas.berkeley.edu/>

²A similar process of decentralized collection of results has begun for drug discovery (*Lessl et al.*, 2011).

2.2.1. Model Types

The DMAS library results are conditional distributions, representing one or more parameter estimates typically in a dose-response curve. The following representations are used for impact functions in this report:

Discrete-Discrete Probability Models: The discrete-discrete probability model represents either a sampled approximation to a continuous probability density function, $f(y|x)$, at discrete values of $y \in \{y_i\}$ and $x \in \{x_i\}$, or a probability mass function of the same form. This is most appropriate when the collection of response outcomes is limited or categorical. Both the dependent and independent variables may be either categorical or sampled at a collection of numerical levels. For continuous functions, the sampling of the dependent variable, $\{y_i\}$, and the independent variable, $\{x_i\}$, can be uneven.³ This model can be treated as a matrix $P = (p_{ij} = f(y_j|x_i))$. Discrete-discrete probability models are the ultimate form for any Bayesian hierarchical meta-analysis result, after draws from the posterior distribution are organized into histograms.

Spline Models: The spline model represents a continuous conditional probability function, using a spline to denote the log of its values.

$$f(y|x) = \begin{cases} e^{a_0+b_0y+c_0y^2} & \text{for } y_0 \leq y < y_1 \\ e^{a_1+b_1y+c_1y^2} & \text{for } y_1 \leq y < y_2 \\ \dots & \dots \end{cases}$$

Distinct splines are described at distinct values of the conditioning variable $x \in \{x_i\}$, which may be categorical or numerical and may be sampled unevenly. The lowest value of y_0 for each x may be $-\infty$, and the highest value of y_1 may be ∞ . Spline models are used for most impact functions, since they provide arbitrary resolution on the shape of the conditional distribution curve.

Bin Models: A bin model represents a model defined across continuous spans, where the distribution is constant over each span. It is a combination of information describing the width of each bin and an underlying categorical model of one of the other model types describing each bin’s probability distribution. Bin models are used for degree-day impacts as in *Schlenker and Roberts (2009)*, with an underlying spline curve representation for each bin.

2.2.2. Importing Models

Each model type has a file format specification, and parameter estimates can be added by uploading files in these formats. We also provide a variety of simplified ways to specify models. This includes a spreadsheet-style entry for discrete-discrete probability models, a GUI model generator for uniform, Gaussian, and polynomial models, and a “feature interpreter” which allows spline models to be described in terms of any collection of their mean, variance, standard deviation, skew, mode, or arbitrary confidence intervals.

2.2.3. Additional Features

Finally, DMAS includes a wide range of features that can help support the continued evolution and use of these results. The models can be visualized, compared, and combined with weights. They can be selectively included or excluded from a meta-analysis combination, either manually or using arbitrary population or study characteristics. Arbitrary functions of results can be computed, applied to data, and output to for external use. Finally, the entire database can be quickly searched using tags, parameter definitions, and study-specific meta-data.

Some of the greatest advantages of DMAS are its collaborative aspects. Scientists can curate collections of results for meta-analyses, with a moderation system for others to submit additions. They can also create crowd-sourcing templates, which ask for study results in a customized form, to quickly collect information from many researchers. In addition, each result and each collection can act as a discussion board, inviting authorized users to debate the choices used.

We hope that this tool can act as a platform to promote the meta-analysis process and make results available to both modelers and a wider audience.

³For calculating a CDF, we assume that each y_i represents a histogram-style bin, while under interpolation we use linear interpolation (see section 8).

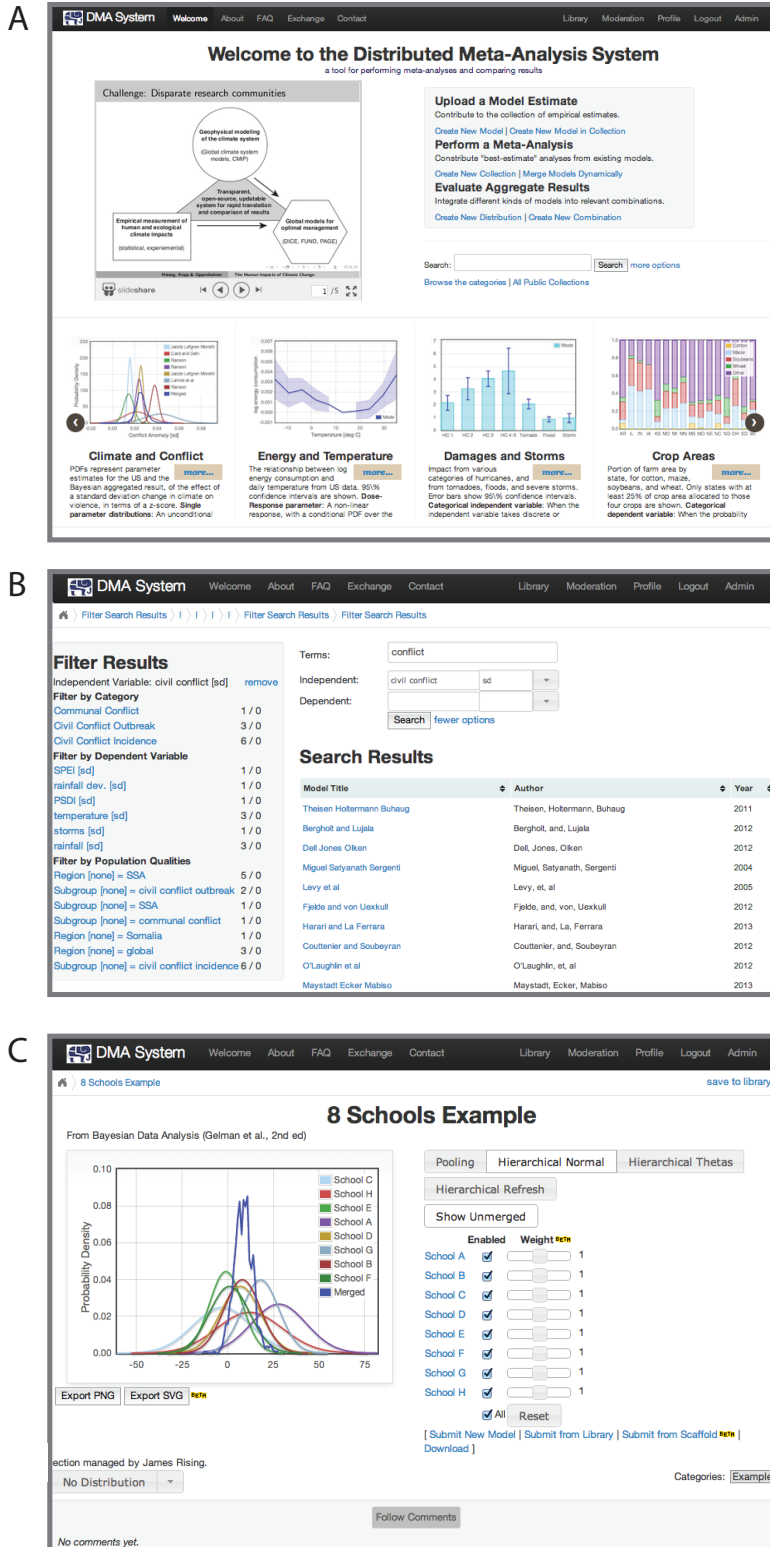


Figure B2: **A:** Front page of the DMAS system, showing a gallery of existing results. **B:** Search screen, allowing results to be filtered by population and study characteristics. **C:** A sample result, the 8-schools example from *Gelman et al.* (2013), showing options for selecting and weighting results and identifying the meta-analysis combination technique.

3. Micro-founding impact functions

We develop empirical, micro-founded impact functions for a number of sectors seen to be economically important. These include agriculture, crime, health, and labor. Within each sector, we draw on statistical studies that robustly account for a number of potential confounding factors when trying to identify the impacts of climate. For the current analysis, we make no claim to having performed an exhaustive quantitative meta-analysis from the reviewed papers. Numerous high-quality and insightful studies are omitted from sectors, though many studies were used to confirm the validity of the selected papers. However, we have designed our approach to be *inclusive in the long-run* by building an open-source system for meta-analysis and collaboration. Incorporating each study took considerable effort, often requiring new data, efforts on the part of the original authors and ourselves to rerun analyses, and extensive discussions to ensure an accurate interpretation of results. In this process, we are indebted to each of the authors listed below. Our final selection required studies to meet the following criteria:

1. **Nationally representative.** We required that studies be conducted at national level or be drawn from a representative random sample of the entire US. This was of particular relevance to health sector studies. For example, many that we considered performed detailed time-series analysis of single or multiple cities (e.g., *Curriero et al.*, 2002; *Anderson and Bell*, 2009). While these were high-quality studies, inclusion would have required either a weighting scheme based on city populations or an assumption of national generalizability.
2. **Analyze recent time-periods in US history.** As we are concerned with potential effects of adaptation, we preferred studies that identified effects as close to the present as possible.
3. **Robust to unobserved factors that differ across spatial units** (jurisdictions, counties, or states). We placed an emphasis on studies that were able to control for unobservable differences between spatial units of analysis with the inclusion of fixed effects. This required the use of longitudinal or panel data, as cross-sectional comparisons between could suffer from omitted variable bias.
4. **Identify responses to high-frequency climatic variables** (days or weeks). The importance of using high-frequency data to estimate climate impacts is demonstrated by all papers included, building on early work by *Deschenes and Greenstone* (2007), and in one case finding large effects by considering sub-daily temperature responses (*Schlenker and Roberts*, 2009).
5. **Identify responses to the full distribution of temperature and rainfall measures.** Many studies looked at single climatic events, or parts of the temperature or rainfall distribution (e.g., heat-waves in *Anderson and Bell*, 2011). As we are modeling annual impacts, we chose only those studies that included the full distribution of realised climate outcomes, and ensured the validity of results by comparison to numerous studies looking at single phenomena or sub-populations.
6. **Account for seasonal patterns and trends in the outcomes.** Cyclicity and seasonality of responses to climate forcings is a source of major concern, so we selected only those studies that robustly accounted for seasonal patterns and time trends in their analysis.
7. **Ecologically valid.** We required studies to be valid for real-life circumstances and levels of exposure, which led us to prefer studies that were quasi-experimental in design, using observational data. For example, in the case of labor, numerous laboratory studies exist on the intensive margin effects of temperature upon productivity (e.g., *Seppanen et al.*, 2006). As these raised a question of ecological validity when applied to the labor sector, we chose to not include them.

Many of the impacts of climate change will unfold over years, but distinguishing between the role of climate change and the role of social, technological, and economic evolution is very difficult over any long time horizon. Our criteria for selecting studies requires that long-term trends are accounted for and are not reflected in the measured impact response functions. As a result, the impacts that we measure are from weather “shocks”, short-term changes in temperature and precipitation which are not captured by long-term trends. This approach has both strengths and weaknesses. Its key strength is that it clearly identifies the

impacts of weather as distinct from longer-term changes. However, it may miss many of the long-term impacts of climate change that do not take the form of increases in the size, frequency, and duration of weather shocks.

We identify a number of studies using panel data to isolate the variation within the relevant spatial unit, while controlling for unobservable difference between units. Estimates from each of the studies were combined, as detailed in section 2. We have been conservative in our choice of studies for the current analysis, using only studies which we think most credibly identify the impact of climate upon specific outcomes in each sector. However, our approach allows for future studies to be incorporated, introducing new findings, and modifying the current results. The following is a complete list of empirical response functions used in this study, with detailed discussion of each of the studies beneath (shown in fig. B4):

Agriculture	Maize yields vs. temperature (East)
	Maize yields vs. temperature (West)
	Maize yields vs. precipitation (East)
	Maize yields vs. precipitation (West)
	Wheat yield vs. temperature
	Soybean yields vs. temperature (East)
	Soybean yields vs. temperature (West)
	Soybean yields vs. precipitation (East)
	Soybean yields vs. precipitation (West)
	Cotton yields vs. temperature
	Cotton yields vs. precipitation
	Maize yields vs. 100ppm CO ₂ increase
	Wheat yields vs. 100ppm CO ₂ increase
	Soybean yields vs. 100ppm CO ₂ increase
	Cotton yields vs. 100ppm CO ₂ increase
Crime	Violent crime vs. temperature
	Violent crime vs. precipitation
	Property crime vs. temperature
	Property crime vs. precipitation
Health	Mortality vs. temperature (all age)
	Mortality vs. temperature (younger than 1 year)
	Mortality vs. temperature (1 - 44 years)
	Mortality vs. temperature (45 - 64 years)
	Mortality vs. temperature (65 years and up)
Labor	Hours worked in high-risk industries vs. temperature
	Hours worked in low-risk industries vs. temperature

3.1. Agriculture

Schlenker and Roberts (2009)

Outcome data:	Yields for maize, soybeans, and cotton from US Department of Agriculture National Agricultural Statistical Service.
Climate data:	PRISM temperature and rainfall, temporally downscaled to daily resolution.
Sample period:	1950-2009
Sample unit:	County-years, for counties with recorded yields of maize, soybeans, or cotton
Methodology:	Piecewise linear response of $\log(\text{yield})$ to cumulative temperature (degree days) and polynomial response to precipitation (seasonal total), controlling for county fixed effects and state-specific quadratic trends. Piecewise linear models are specific to each crop type, with thresholds that capture the beneficial effects of temperatures below a certain point, and the deleterious effects above.
Result:	Modified version of <i>Schlenker and Roberts</i> , 2009, SI Appendix, p. 9, fig. A3; and p. 20, fig. 10
Impact function:	We contacted the authors of the study to select a preferred response function from the multiple methods they had employed, selecting a piecewise-linear specification using degree days for temperature and seasonal total precipitation. We obtained impact functions for each of the three crops studied, for both temperature and precipitation. The authors note the distinct difference in response between counties to the east and west of the 100 th meridian for maize and soybeans, so we obtained separate response functions in for these regions. On December 19 th , 2013, we were sent a complete list of response functions that were updated span the time period up to and including 2011 (as presented in <i>Berry et al.</i> , 2012).

Hsiang, Lobell, Roberts, and Schlenker (2013a)

Outcome data:	USDA-NASS
Climate data:	University of Delaware monthly temperature and precipitation
Sample period:	1950-2007
Sample unit:	County-year
Methodology:	Non-linear response of $\log(\text{yield})$ to crop-specific seasonal average temperature and precipitation, controlling for county and year fixed effects.
Result:	<i>Hsiang et al.</i> , 2013a, p. 19.
Impact function:	We use the response of wheat to seasonal average temperature presented in the paper. Results were obtained from the author. Calorie weighted averages were taken between maize and wheat in order to combine results, as detailed in section 4.

McGrath and Lobell (2013)

Outcome data:	Yield from 1960-2004 from FAOStat.
Climate data:	Keeling CO ₂ concentrations and country average P/PET.
Methodology:	Process model that develops the response of different crops to carbon dioxide concentrations and growing season P/PET from empirical studies. This is then used to estimate the changes to historical yields under a 100ppm increase in CO ₂ .
Result:	<i>McGrath and Lobell</i> , 2013, p. 5, fig. 4 (obtained US result from authors).
Impact function:	We contacted the authors and received estimates of the CO ₂ fertilization relationship with yields of different crops on January 17 th , 2014, specifically for the US. Data were for 8 different crop types. We used an average of all types for cotton estimates.

3.1.1. Storage

In addition to the above impacts on yields, we observe that farmers store crops for sale in the future, and so the overall impact of climate on supply of crops may appear smoother than if there were no storage. For our projections, we also make use of Fisher et al. (2012, Appendix p.xi, table A4) to estimate crop




Figure B3: Predicted consumption of maize, modeled as a moving average of production. Predicted values compare well to observed consumption, and allow us to project the smoothed consumption values out to the end of the century.

consumption as a moving average process of crop production. We estimated the following equation for crop c ,

$$\ln(\text{consumption})_{c,t} = \sum_{l=0}^L [\beta_{c,l} \times \ln(\text{production})_{c,t-l}] + \theta_c t + \gamma_c t^2 + \epsilon_{c,t}$$

where $L = 2$ and $L = 3$ for soybeans, and we account for linear and quadratic time-trends. Results of this process are shown in fig. B3. We project the smoothing of future crops with a time-series structure that incorporates these empirical results on storage. Weights for each crop are constructed from the lagged coefficients, β_l , presented in section 4.1.

3.2. Crime

Jacob, Lefgren, and Moretti (2007)

Outcome data:	FBI National Incident Based Reporting System
Climate data:	Weekly temperature and precipitation from the NCDC GHCN-Daily database.
Sample period:	1995-2001
Sample unit:	Jurisdiction-weeks
Methodology:	Linear response of $\log(\text{crime_rate})$ to average temperature and precipitation, controlling for jurisdiction-by-year and month fixed effects, as well as jurisdiction-specific 4 th order polynomials in day of year.
Result:	Modified version of <i>Jacob et al.</i> , 2007, p. 508-509, table 2.

Impact function: We obtained data and replication files from the authors and generated coefficients for a month-long exposure window, to account for displacement of crime, as noted in the text. The climate variables are at weekly resolution, and in order to make this comparable to *Ranson* (2014) we reran the analysis using maximum temperatures and then scaled the coefficients in *Jacob et al.* (2007). We did this by first dividing the coefficient for the monthly exposure by 7, to get a daily response, and further by 4 to account for the lagged climate variables. This resulted in the marginal effect on crime of a 1°F increase in daily temperature. Taking a reference point of zero response at a temperature of 65°F (to coincide with the central point of the reference bin of (*Ranson*, 2014)) we derived a linear response of violent crimes and property crimes to temperature and precipitation.

Ranson (2014)

Outcome data: FBI Universal Crime Reporting Data.
 Climate data: Daily temperature and precipitation from the NCDC GHCN-Daily database.
 Sample period: 1960-2009
 Sample unit: County-months
 Methodology: Non-linear response of $\log(\text{crime_rate})$ to maximum temperature and precipitation, controlling for county-by-year and state-by-month fixed effects. Temperature is transformed into number of days within 10°F bins, with the 60-69°F bin as a reference point.
 Result: *Ranson*, 2014, p. 9, fig. 4
 Impact function: We contacted the author and received updated estimates of the percentage change for each of 8 different classes of crimes on March 12th, 2014. To derive response functions, we grouped these into violent crimes (murder, rape, aggravated assault, and simple assault) and property crimes (robbery, burglary, larceny, and vehicle theft), and combined results within each class of crimes.

3.3. Health

Deschenes and Greenstone (2011)

Outcome data: National Center for Health Statistics Compressed Mortality Files.
 Climate data: Daily temperature and precipitation from NCDC
 Sample period: 1968-2002
 Sample unit: County-years
 Methodology: Non-linear response of mortality to temperature, controlling for county-by-age-group and state-by-year-by-age-group fixed effects. Temperature is transformed into number of days in an year-long window within 10°F bins, with the 50-59°F bin as a reference point.
 Result: Modified version of *Deschênes and Greenstone*, 2011, p. 9, fig. 2
 Impact function: We contacted the authors and received estimates on November 5th, 2013. To make the study comparable to Barreca et al. (2013), the main analysis was rerun with $\log(\text{mortality})$ as an outcome.

Barreca, Clay, Deschenes, Greenstone, and Shapiro (2013)

Outcome data: Mortality from the Mortality Statistics of the US (pre-1959) and the Multiple Cause of Death files (post-1959).
 Climate data: Daily temperature and precipitation from the NCDC GHCN-Daily database.
 Sample period: 1929-2004
 Sample unit: State-months

Methodology:	Non-linear response of $\log(\text{mortality})$ to temperature, controlling for state-by-month and year-month fixed effects, and state-by-month-specific quadratic time trends. Temperature is transformed into number of days in a two-month window within 10°F bins, with the 60-69°F bin as a reference point.
Result:	Modified version of <i>Barreca et al.</i> , 2013, p. 37, table 3, panel B
Impact function:	We contacted the authors and received estimates on 5 th November, 2013. The preferred specification, to account for forward displacement, was to use monthly mortality with a 2-month exposure window to temperature. We used the estimated response from 1960-2004. To make this response comparable to the response of Deschenes and Greenstone (2011), the analysis was rerun with the reference point changed to the 50-59°F bin. To scale the coefficients, we divided each coefficient value by a factor of six. We also obtained age-specific response functions for ages 0-1, 1-44, 45-64, and 65+.

3.4. Labor

Graff Zivin and Neidell (2014)

Outcome data:	Hours worked from the American Time Use Survey.
Climate data:	Daily temperature, precipitation, and humidity from NCDC.
Sample period:	2003-2006
Sample unit:	Person-days
Methodology:	Seemingly-unrelated regression allowing for correlated errors between time spent working, or indoor and outdoor leisure. Non-linear response to maximum temperatures controlling for county, year-by-month, and day of week fixed effects, as well as individual level controls. Temperature is transformed into number of days within 5°F bins, with the 76-80°F bin as a reference point.
Result:	High-risk: <i>Graff Zivin and Neidell</i> , 2014, p. 15, fig. 3; Low-risk: <i>Graff Zivin and Neidell</i> , 2014, p. 16, fig. 4
Impact function:	We contacted the authors prior to publication and received full estimates for high-risk and low-risk labor responses to temperature on December 18 th , 2013.

4. Application of Impact Functions

We apply two approaches for sampling the conditional distributions for each impact function: a Monte Carlo approach and a constant quantile approach. The Monte Carlo approach captures the full range of uncertainty in impact functions estimates, under the assumption that each impact function is independent. We randomly select quantiles for each of the 26 empirical distributions. The constant quantile approach applies the same quantile across all distributions. In particular, we use a low quantile ($p = 0.33333$), median quantile ($p = 0.5$) and high quantile ($p = .66667$). The ordinality of the quantiles is chosen so that these describe, in essence, low, median, and high impact scenarios. High quantiles correspond to greater losses in yield and labor productivity, and greater increases in crime and mortality, within the range of statistical uncertainty.

Under either approach, the same quantile is used across the entire range of the conditioning variable. By evaluating each impact function at a quantile, we generate a single-dimensional, deterministic function which is used in the evaluation of the impact for each Monte Carlo or constant quantile run.

The impact results for crime, labor productivity, and mortality are all estimated by binning weather values. In these cases, we construct a continuous impact curve by linearly interpolating between the midpoints of these bins.

Impact results were produced in “result sets” that consisted of an impact value for each year in each county in the US. Result sets are grouped into “batches”, with a result set for each RCP (2.6, 4.5, 6.0, and 8.5), each GCM model or model surrogate (from 28 to 44 models in each RCP), and each of 10 weather realization, for a total of 1440 result sets per batch. We produced 25 batches of Monte Carlo results. These results are later combined and weighted as described in section 7. Twenty-six conditional distributions were used

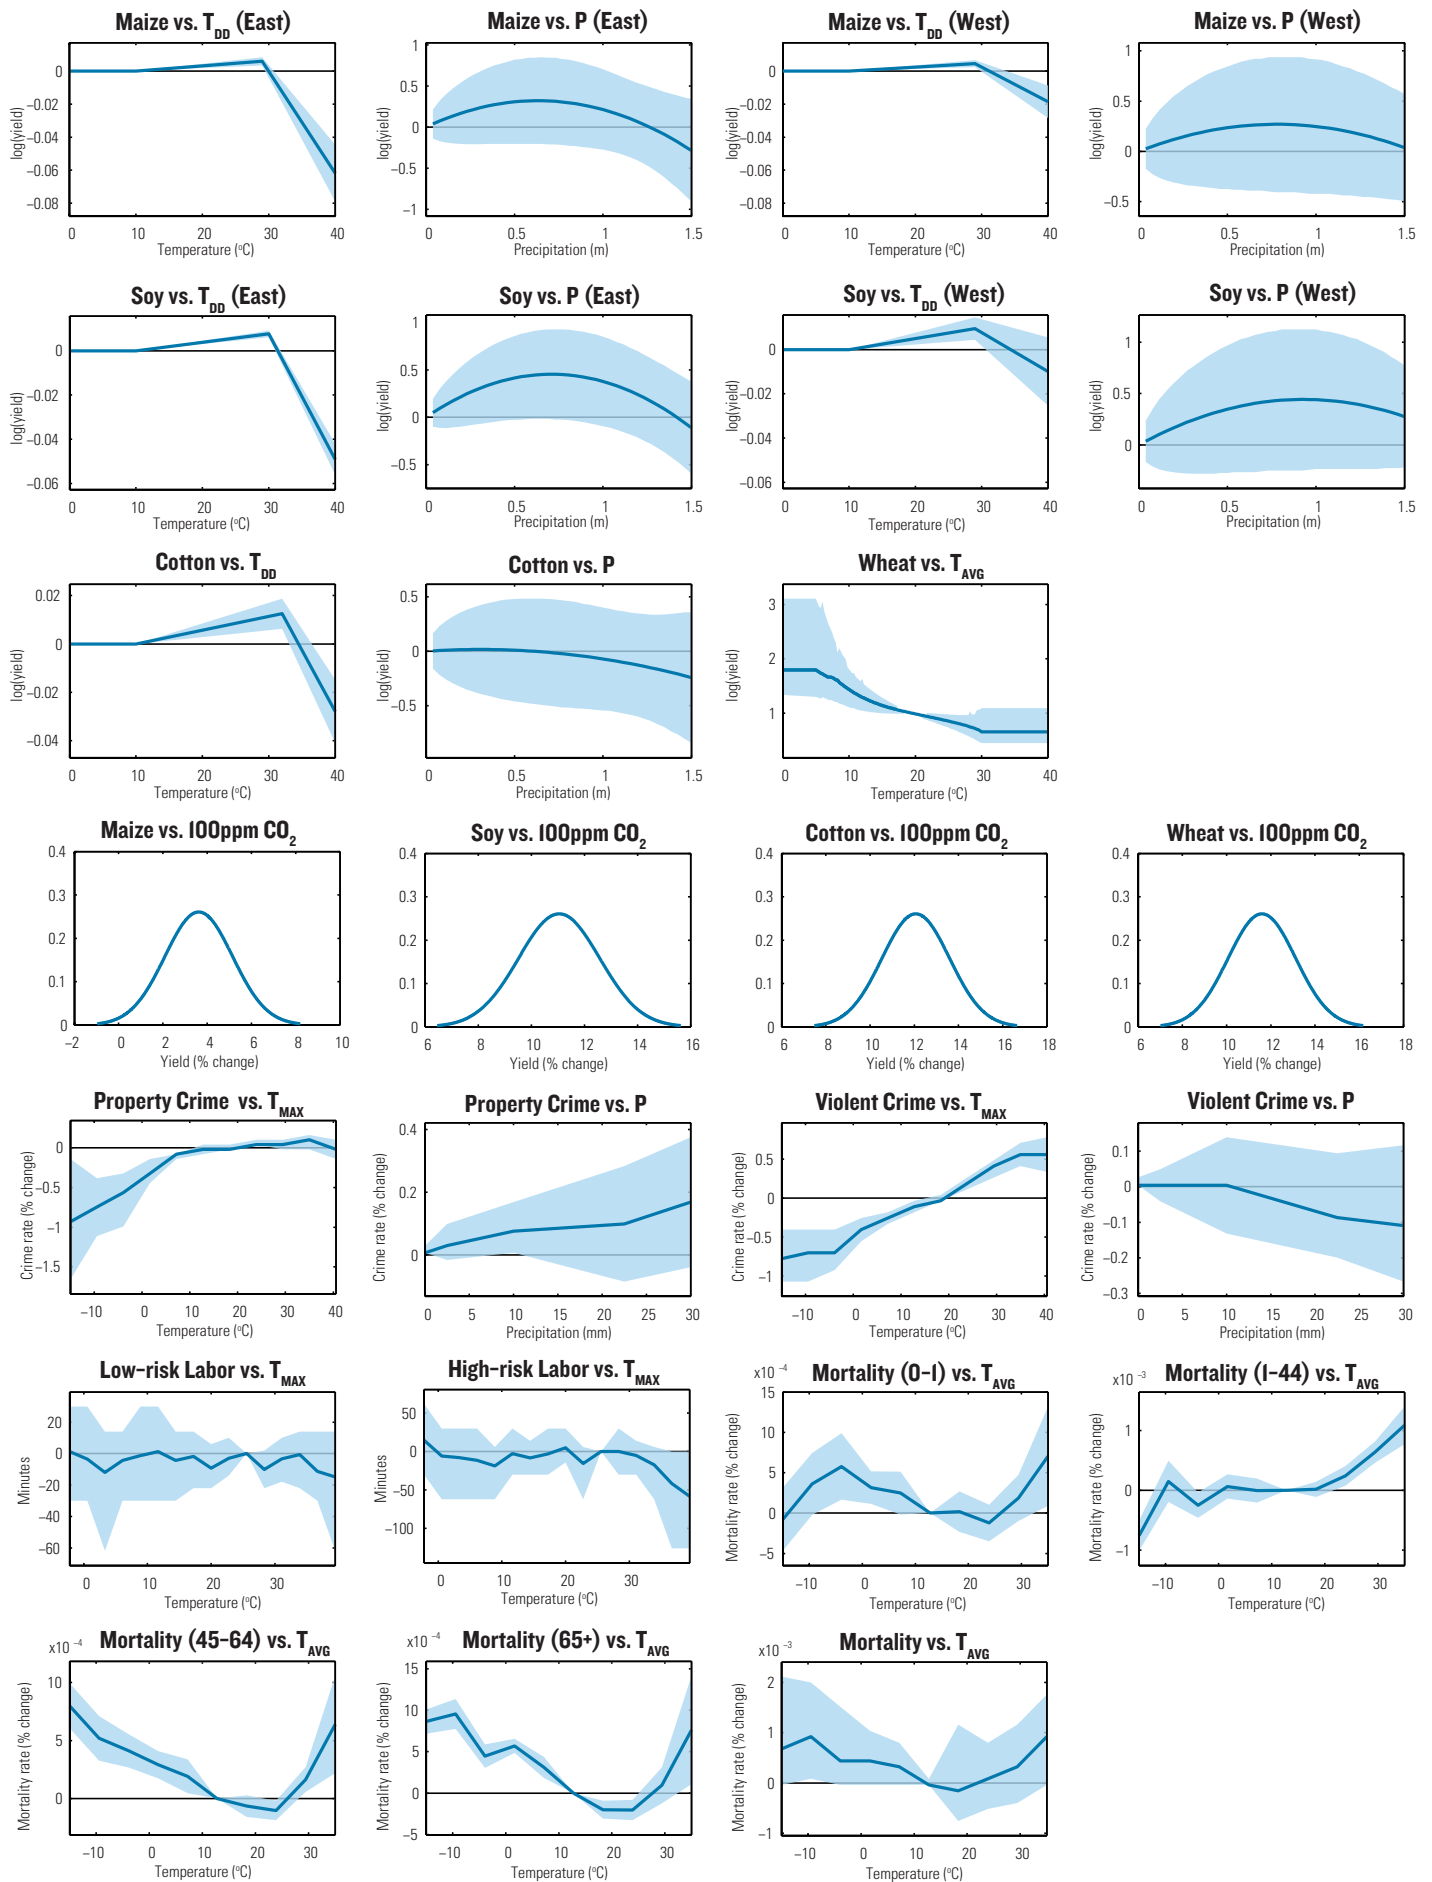


Figure B4: All 26 dose-response functions used in our analysis. 95% confidence intervals are shaded.

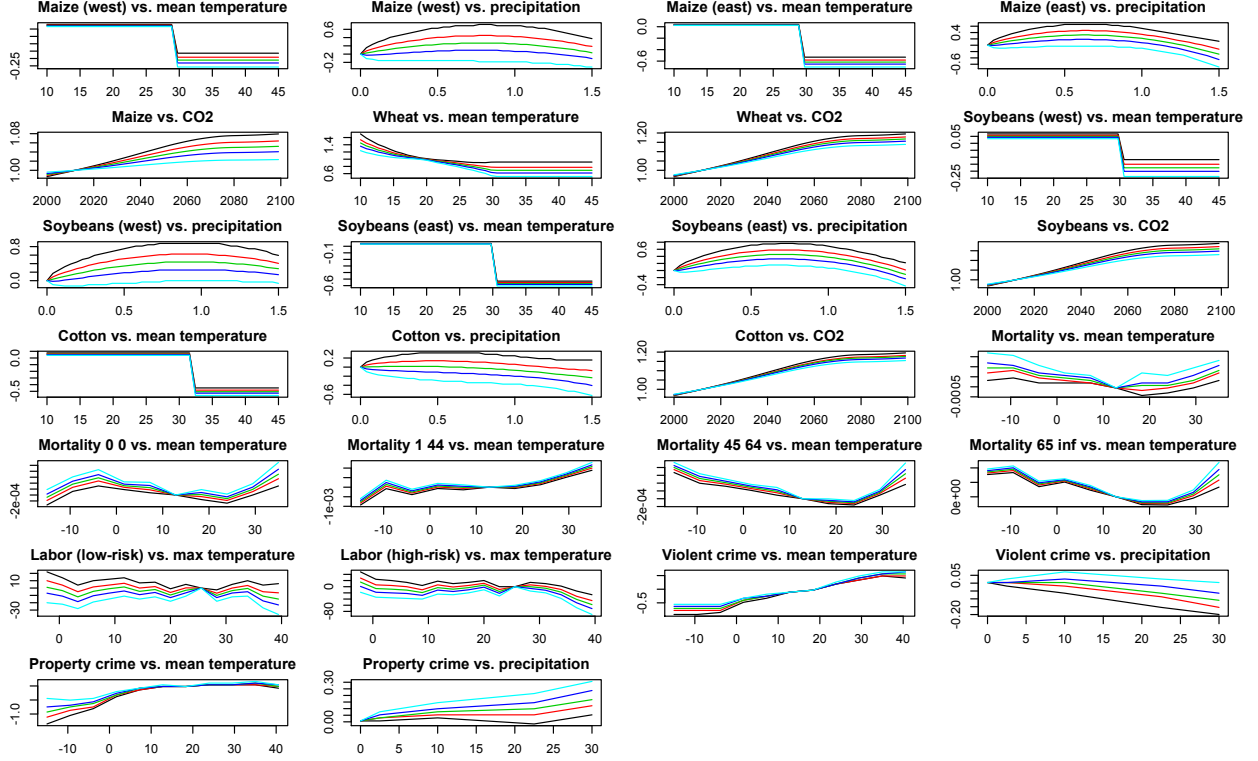


Figure B5: Quantiles of the response functions, showing the impact functions that would be used in evaluation for quantiles of $p \in \{.1, .3, .5, .7, .9\}$. CO₂ fertilization graphs show evolution for RCP 4.5.

to generate 15 impact results. These calculations are described below. Throughout, impacts are ultimately reported relative to 2012, the baseline year for the CGE model. The definitions of each impact below do not reflect this.

Below, we use the notation T_{AVG} for mean daily temperature; T_{MIN} and T_{MAX} for minimum and maximum daily temperature, respectively; and P for precipitation. All weather variables are provided at the county-level and on a daily basis. Below, where counties are indexed by j , weather variables are implicitly also indexed by j . We index days in the year by d and years by t .

4.1. Agricultural Yields and Production

Percent changes in agriculture production, relative to 2012, were generated using fixed, county-specific growing seasons. The growing season, denoted $S(j)$ for county j , is determined using the centroid of the county applied to the planting and harvesting dates in *Sacks et al. (2010)*. For maize and wheat, for which *Sacks et al. (2010)* provides two calendars (two croppings for maize, and summer and winter wheat), the calendar that represented the greatest portion of land area in each county was used.

Relative changes in yield are calculated based on seasonal temperatures and precipitation as follows:

Wheat Wheat uses a seasonal average temperature response function:

$$Y_{jt} = f \left(\frac{1}{N(S(j))} \sum_{d \in S(j)} T_{AVG,d} \right)$$

where $f(\cdot)$ is calculated by *Hsiang et al. (2013a)*, as a function of average mean daily temperature over the growing season, and $N(S(j))$ is the number of days in the growing season for county j . This functional form was only used for wheat, since a degree-day representation was unavailable.

Cotton Cotton uses a single degree-day function:

$$Y_{jt} = e^{f(0.01 \sum_{d \in S(j)} DD_{low}(T_{MAX,d}, T_{MIN,d}), 0.01 \sum_{d \in S(j)} DD_{high}(T_{MAX,d}, T_{MIN,d})) + g(1 \times 10^{-3} \sum_{d \in S(j)} P_d)}$$

where DD_{low} and DD_{high} are growing degree days below and above the crop-specific breakpoint specified in *Schlenker and Roberts (2009)*, and calculated as specified there using the minimum and maximum daily temperatures. The functions $f(\cdot)$ and $g(\cdot)$ translate degree days and precipitation, respectively, into yield effects.

Maize and Soybeans Maize and soybeans have two degree-day responses:

$$Y_{jt} = \begin{cases} e^{f_{east}(0.01 \sum_{d \in S(j)} DD_{low}(T_{MAX,d}, T_{MIN,d}), 0.01 \sum_{d \in S(j)} DD_{high}(T_{MAX,d}, T_{MIN,d})) + g_{east}(1 \times 10^{-3} \sum_{d \in S(j)} P_d)} \\ e^{f_{west}(0.01 \sum_{d \in S(j)} DD_{low}(T_{MAX,d}, T_{MIN,d}), 0.01 \sum_{d \in S(j)} DD_{high}(T_{MAX,d}, T_{MIN,d})) + g_{west}(1 \times 10^{-3} \sum_{d \in S(j)} P_d)} \end{cases}$$

Here, $f_{east}(\cdot)$ and $g_{east}(\cdot)$ are used to the east of the 100th meridian, excluding Florida. In this areas, irrigation is less common and the response to increased temperatures is more extreme.

Figure B6 shows distributions of degree days for RCP 6.0, under the MIROC-ESM-CHEM model.

CO₂ fertilization is modeled as a multiplicative factor applied to yields, and estimated as a linear increase for each additional 100 ppm of CO₂:

$$Y'_{jt} = Y_{jt} \left(1 + \frac{C_t - C_{2012}}{100} X \right)$$

where C_t is the CO₂ concentration in year t under a given RCP, and X is the estimated CO₂ fertilization effect, which varies from 3% to 12% depending on the crop, from *McGrath and Lobell (2013)*. *McGrath and Lobell (2013)* does not provide a value for cotton, so the Bayesian combination of all provided crop effects is used for it.

Economic output from the agricultural sector is not synonymous with yield, due to strategic storage. We model output as an autoregressive process of yields, as estimated from USDA data. For grains and cotton, the expression is:

$$I_{jt} = 0.51Y_{jt} + 0.28Y_{j,t-1} + 0.21Y_{j,t-2}$$

A four-year moving average is used for soybeans:

$$I_{jt} = 0.2Y_{jt} + 0.52Y_{j,t-1} + 0.19Y_{j,t-2} + 0.1Y_{j,t-3}$$

4.2. Crime

Both violent and property crime are calculated as,

$$I_{jt} = \left(1 + 0.01 \frac{1}{12} \sum_{d \in y(t)} f(T_{MAX,d}) \right) \left(1 + 0.01 \frac{1}{12} \sum_{d \in y(t)} g(P_d) \right)$$

where $y(t)$ is the set of days in year t , $f(\cdot)$ is a function of daily maximum temperature, and $g(\cdot)$ is a function of daily precipitation. Both $f(\cdot)$ and $g(\cdot)$ are calculated as a Bayesian combination of the effect for each sub-category of crime estimated by *Ranson (2014)* and the average effect for violent or property crime from *Jacob et al. (2007)*.

4.3. Mortality

Both average and age-specific mortalities are calculated as,

$$I_{jt} = \sum_{d \in y(t)} f(T_{AVG,d})$$

The parameters of $f(\cdot)$ are calculated as a Bayesian combination of the results from *Deschênes and Greenstone (2011)* and a corrected form of *Barreca et al. (2013)*. Age-specific mortalities, for newborns, ages 1-44, ages 45-64, and ages 65 and up, are provided by *Barreca et al. (2013)*.

Mortality is reported both as percentage changes, and as differences in the mortality rate. In either case, the pooled and age-specific mortality rates per county are from the *Center for Disease Control and Prevention (2013)*.

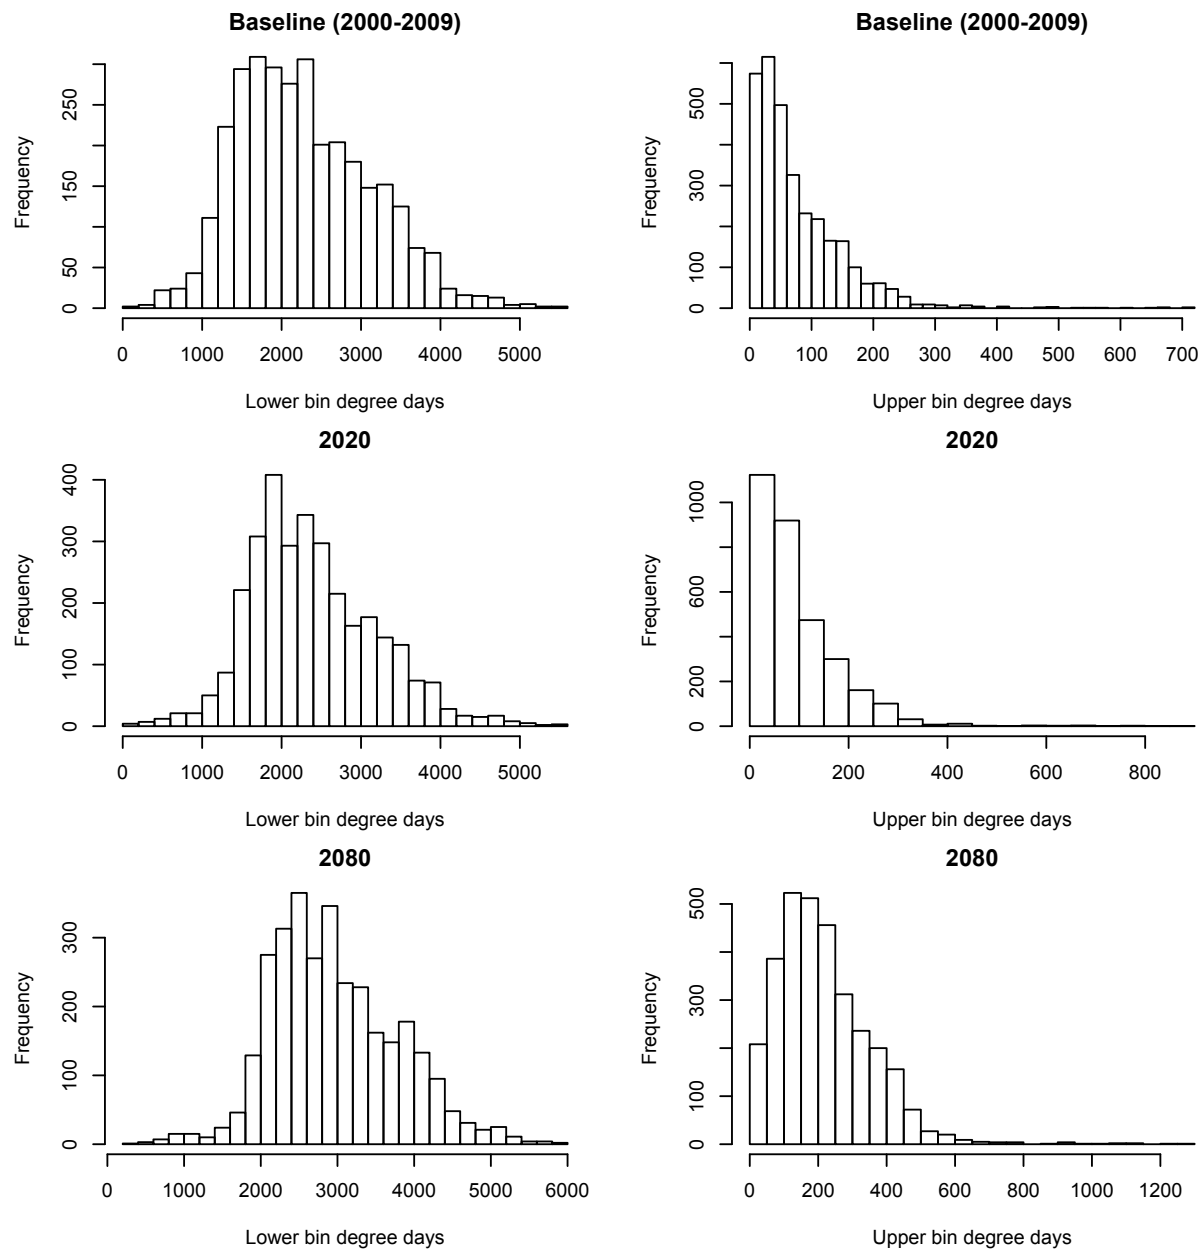


Figure B6: Distributions of the degree days across all US counties in the lower (left) and upper (right) bins of Maize, which has a 29°C bin threshold. The rows represent the average of degree days between 2000 and 2009, and for the years 2020 and 2080.

4.4. Labor Productivity

The structure of the labor productivity calculation is identical for high-risk and low-risk sectors:

$$I_{jt} = \frac{H + \frac{1}{60} \sum_{d \in y(t)} f(T_{MAX,d})}{H}$$

where H is the average number of hours worked per year in the baseline. For high-risk labor, $H = 7.67 \times 365$, and for low-risk labor, $H = 6.92 \times 365$. The parameters of $f(\cdot)$ are provided by *Graff Zivin and Neidell* (2014).

5. Adaptation

The empirical approach used to estimate adaptation applies observed rates of adaptation from recent history to regional differences in response functions to capture the observed capacity for adaptation. Here, we describe the general structure for estimating adaptation. Additional details are added in sections specific to each adapted response, below.

Each impact response curve starts from the same baseline as the unadapted case, in the year 2000. Within each county, it then transitions asymptotically toward a future shape. The evolution of each parameter of the distribution follows,

$$\beta(t) = \beta(\infty) + (\beta(t_{before}) - \beta(\infty))e^{-(t-t_{before})/\tau}$$

where $\beta(\infty)$ is the maximum possible adaptation, which in the cases considered here represents no response to temperature, $\beta(t_{before})$ is the parameter from a historical period, as estimated for a period centered on t_{before} , and τ determines the rate of evolution of the parameter.

Using parameters estimated for a second period closer to the present day than t_{before} , we calculate the rate of adaptation, τ , as,

$$\tau = -\frac{t_{after} - t_{before}}{\log(\beta(t_{after}) - \beta(\infty)) / (\beta(t_{before}) - \beta(\infty))}$$

where $\beta(t_{after})$ is the parameter value during a period centered on t_{after} . This process is represented on the left of figure B9.

We further manipulate τ into an incremental product, $\gamma = e^{-1/\tau}$, so that,

$$\beta(t+1) = \gamma\beta(t) + (1-\gamma)\beta(\infty) \tag{B1}$$

With the exception of maize, a different $\beta(\infty)$ is used for determining the rate of evolution (τ and γ) as is used for estimating the final evolution of the parameter values. The $\beta(\infty)$ used to estimate the actual evolution of the parameters is thought of as the parameter value for a fully adapted region of a particular temperature. As a result, it is a function of average temperature: the response curve for New England, under the current climate, is considered to be fully adapted to a colder average temperature than the response curve for the Southwest. To estimate how $\beta(\infty)$ changes with temperature, we use values from several regions in the US, treated each as fully adapted to its current climate.

Specifically, we construct a linear approximation to the evolution of $\beta(\infty)$ across temperatures, using the parameters and average temperature \bar{T} of various regions. Let this function be denoted $f(\bar{T})$. For the dynamic evolution of the adapted response curves, we take $\beta(\infty) = f(\bar{T}_t)$. At any point in time, \bar{T} is calculated as the average temperature over the previous 15 years. This allows the fully-adapted coefficients to be predicted both between the temperatures of observed regions, and extrapolated beyond them. The process is shown on the right of figure B9.

Each section below describes the evidence used to estimate the temporal evolution of impact responses (the rate of adaptation, τ), and the spatial estimation (how $\beta(\infty)$ varies with \bar{T}).

5.1. Maize in the Eastern U.S.

Increasing the availability of irrigation is a key adaptation to climate change, since additional water for evapotranspiration allows crops to continue to grow at higher temperatures. Fields to the east of the 100th meridian are much less consistently irrigated than those to the west. This helps explain the lower response to killing degree-days in the west ($\beta = -0.21$, compared to -0.62) (*Schlenker and Roberts*, 2009).




Figure B7: Evolution of impacts for FIPS 17161, Rock Island, IL, under RCP 6.0 as modeled by MIROC-ESM-CHEM. In the top panel, the left axis measures linear changes in temperature ($^{\circ}\text{C}$) and CO_2 (100 ppmv). The right axis measures changes in rainfall relative to 2000-2009 annual mean. In the lower panel, both axes show fractional changes in each of the sectors relative to 2000-2009 annual mean.

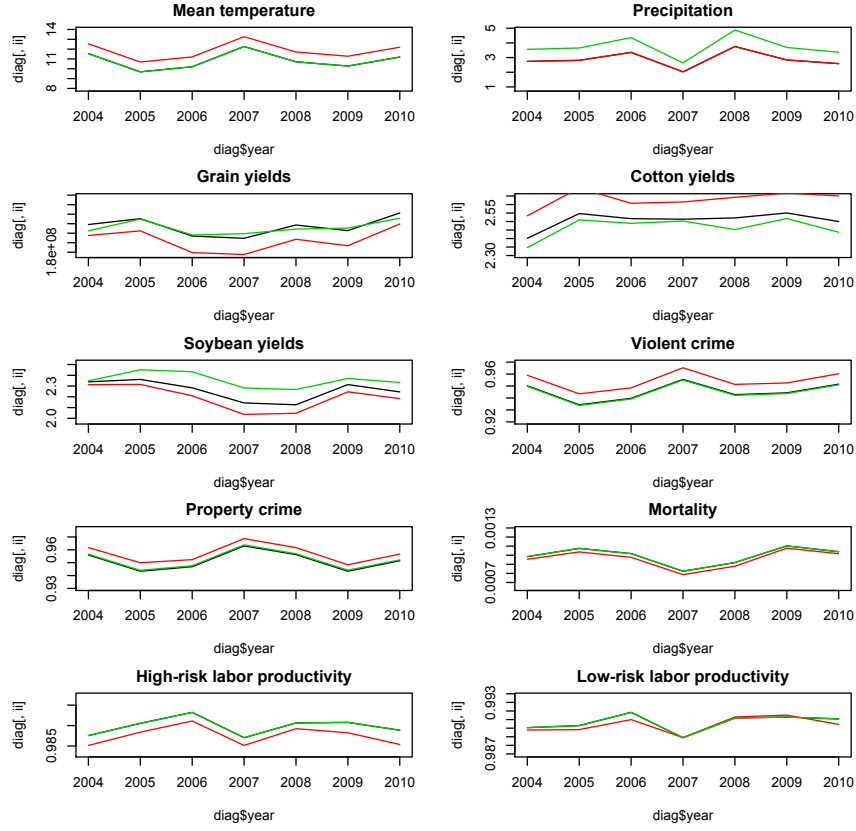


Figure B8: Experiments with FIPS 17161, Rock Island, IL, under RCP 6.0 as modeled by MIROC-ESM-CHEM. The black line denotes a baseline between 2004 and 2010. The red line is applies a 1°C increase on all temperatures. The green line applies a 30% increase on all precipitation.




Figure B9: **Left:** Estimation of temporal evolution of variable values. A historical estimate (at t_0) and a current estimate (at t_1) are used to fit an exponential, characterised by γ . **Right:** Estimation of potential adaptation as a function of average temperature. Regional estimates (in red) are used to produce a linear approximation for how each coefficient varies with average temperature (\bar{T}). This approximation is then evaluated (green) to determine coefficients at future times.

5.1.1. Temporal Evolution

The killing degree-day coefficient has evolved over the course of the last half century, as shown on the left of figure B10. Using Monte Carlo draws from each of these regional parameter distributions independently and fitting exponentials, γ is calculated to be 0.99722 ± 0.00573 . We approximate the distribution of values for γ with a Gaussian, as shown on the right of figure B10. This corresponds to a very slow rate of adaptation, with a time constant of about 360 years (*Burke and Emerick, 2013*).




Figure B10: **Left:** Killing degree-day coefficients computed for 5-year periods between 1955 and 2000. Red lines show the 90% confidence interval on the exponential evolution of the killing degree-day coefficient. **Right:** Values of γ computed in the Monte Carlo, and the corresponding Gaussian approximation in red.

We hold constant the growing degree-days coefficient; only the killing degree-day coefficient adapts. Since the growing degree-day coefficient is higher in the east than in the west, this may result in an optimistic consideration of the trade-offs that result from adaptation.

5.1.2. Spatial Estimation

No extrapolation is performed for maize. Instead, each year the killing degree-day coefficient for all eastern counties is updated according to equation B1.

5.2. Temperature-Related Mortality

5.2.1. Temporal Evolution

For temperature-related mortality, we assume that there are two parallel adaptations occurring: one for cold temperatures and one for hot temperatures. The rate of temporal evolution for these processes is estimated as the mean of the values of γ for each bin for which $T < 65^\circ F$, and for each bin for $T > 65^\circ F$, respectively. Figure B11 shows the coefficient estimates for each bin and the calculation of γ for the median of these coefficient distributions. In the evaluation of actual impacts, Monte Carlo draws are taken from the various distributions, and the value of γ is estimated separately for each Monte Carlo run. The time constant for the cold-related mortality is 31.3 years, and for heat-related deaths is 84.0 years.

5.2.2. Spatial Estimation

Estimates were given for four regions: the Northeast, Midwest, South, and West. These, along with the pooled estimate, were used to construct a linear approximation. See figure B12.

5.3. Violent and Property Crime

The process for estimating adaptation for crime is similar to the process for temperature-related mortality. Pooled results for all forms of property crime and all forms of violent crime were used, as segregated by regions with different maximum temperatures, from *Ranson (2014)*. Since the maximum temperature for the national estimate was very similar to the segregated $65^\circ F$ region, and because the national estimate for the highest temperature bin is outside of the coefficient range of the segregated regions, we used the $65^\circ F$ response function as the baseline for estimating adaptive capacity.

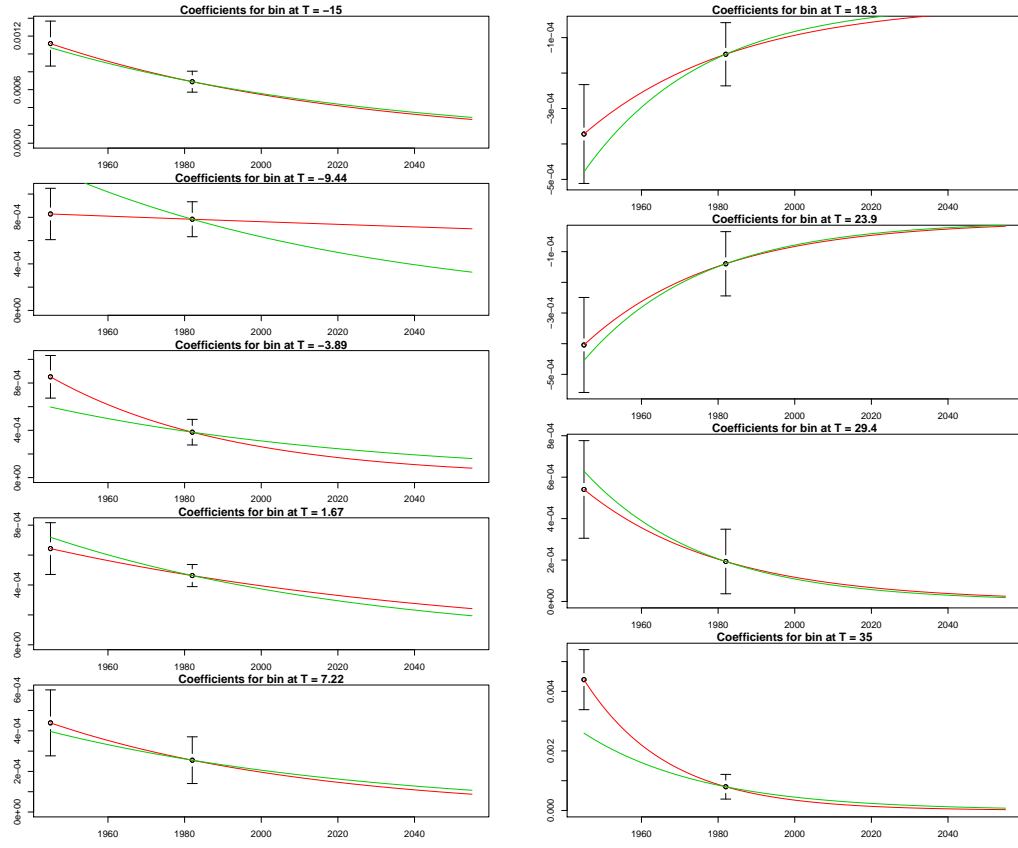


Figure B11: Historical coefficient estimates for cold-related (left) and heat-related (right) mortality. Red shows the time evolution based only on the coefficient described in each panel, while green shows the pooled time evolution for all coefficients in that set.




Figure B12: Linear approximations for cold-related and heat-related mortality β as a function of \bar{T} . Excess mortality for days of a given temperature is clipped at 0. A single coefficient was provided for cold-related mortality ($T < 65^\circ F$).

5.3.1. Temporal Evolution

Estimates for periods centered at 1960 and 2000 were used to estimate the temporal evolution of crime adaptation. Two rates were calculated: one for bins below $65^\circ F$ and one for bins above that value (which also corresponds to the dropped bin in *Ranson* (2014)). These calculations are shown in figure ??.

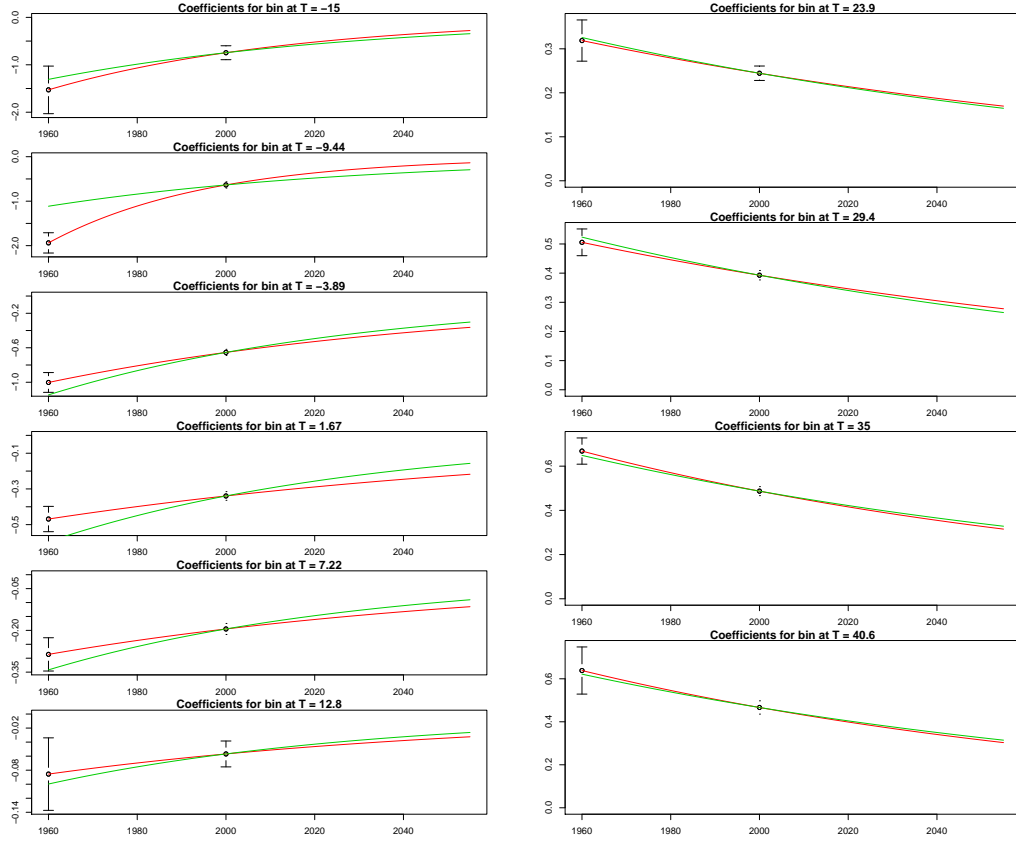


Figure B13: Historical coefficient estimates for violent crime driven by maximum temperatures below (left) and above (right) $65^\circ F$. Red shows the time evolution based only on the single coefficient, while green shows the pooled time evolution for all coefficients in that set.

For violent crime, the time-constant is 71.2 years for the lower range of temperatures and 139 years for the upper range. For property crime, the time-constants are 793 years and 16.4 years, respectively.

5.3.2. Spatial Estimation

Estimates were analyzed from regions with an average maximum temperature of $T = 45^\circ F$, $T = 55^\circ F$, $T = 65^\circ F$, and $T = 75^\circ F$. Figure B15 and B16 shows these estimates by bin.

6. Return Intervals

The number of events exceeding a given threshold is used to communicate return periods of extreme events. We identify the threshold for the 1-in-20 event and count the number of events that exceed this threshold for each 20 year window over the course of the 21st century.

The expected number of events across models for an investment in year t is calculated as a weighted sum,

$$\sum_i w_i e_{it}$$

where w_i is the weight given to model i and e_{it} is the number of events exceeding a threshold experienced in the 20 years following year t . That is, for a threshold x , e_{it} is the number of impacts $y_{i\tau} \geq x$ for




Figure B14: Historical coefficient estimates for property crime driven by maximum temperatures below (left) and above (right) $65^{\circ}F$. Red shows the time evolution based only on the single coefficient, while green shows the pooled time evolution for all coefficients in that set. The value of γ for the bins at $T = 12.7778$, $T = 23.8889$, and $T = 29.4444$ cannot be estimated since the two estimates are on opposite sides of the zero-line.

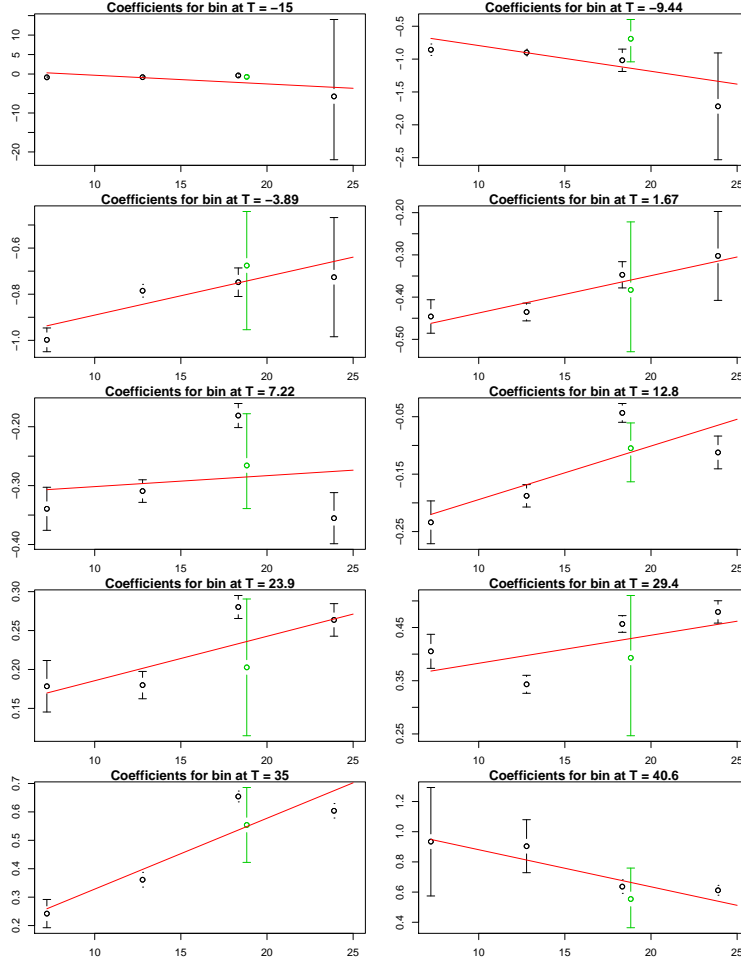


Figure B15: Estimates for the effect of temperatures on violent crime within each temperature bin, for four regions (characterized by $T_{max} = 45$, $T_{max} = 55$, $T_{max} = 65$, and $T_{max} = 75$). The green estimate is the national average estimate, at a national average maximum temperature of $65.88603^{\circ}F$, but was not used for the estimation. The red lines were used for interpolating and extrapolating coefficient values, based on temperature.

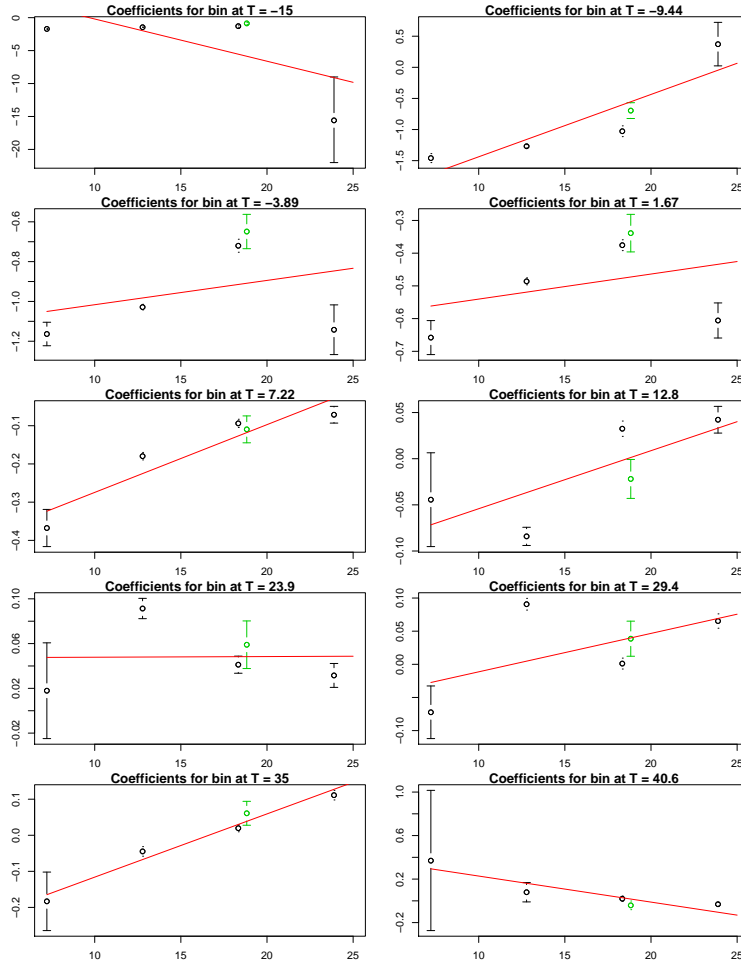


Figure B16: Same as fig. B15 for property crime.

$\tau \in t, \dots, t + 19$. The threshold is chosen such that the expected number of events for the period 1990 - 2009 is approximately 1. These thresholds vary slightly between RCPs, due to the different collections of models used to estimate each RCP. The curves are shifted slightly ($\Delta y < .1$) for display purposes.

7. Impact aggregation

Of the 36000 results generated across all RCPs, models, weather realizations, and Monte Carlo draws of the impact functions, .3% were dropped due to outliers in the estimates of additional mortality in the age 1 - 44 cohort. These outliers were unrealistically high estimates of national averaged mortality, in excess of 100 additional deaths per 100,000. The results from runs that had outlier mortality rates were dropped for all impacts.

County level results are aggregated to the state, NCA-region, and national levels, as weighted sums. This section describes the weighting of results for these aggregations.

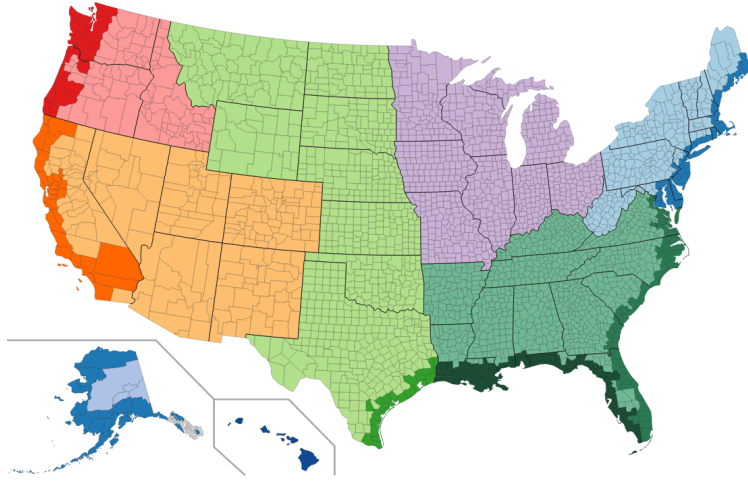


Figure B17: Results were aggregated for passing to the CGE model at both state and NCA region (colors).

7.1. Grain aggregation for CGE analysis

Grains yields (maize and wheat) are both combined within the same county and aggregated to higher scales by calories totals. Within each county, the average impact is,

$$I_i = \frac{I_i^{wheat} A_i^{wheat} C^{wheat} + I_i^{maize} A_i^{maize} C^{maize}}{A_i^{wheat} C^{wheat} + A_i^{maize} C^{maize}}$$

where A_i^{wheat} is the average acres of wheat planted between 2000 - 2005, and A_i^{maize} is the average acres for maize. C^{wheat} and C^{maize} , the calorie density of wheat and maize, are taken to be 1690 calories/kg and 1615 calories/kg, respectively. The weighting of each county result for aggregation to higher scales is

$$W_i = A_i^{wheat} C^{wheat} + A_i^{maize} C^{maize}$$

7.2. Cotton and soybean aggregation for CGE analysis

Cotton and soybean results are aggregated by weighting by the total area planted, as averaged over 2000 - 2005. Since impacts are proportional to changes in yields, and yields are calculated relative to area planted, this is equivalent to weighting counties by production.

7.3. Agricultural aggregation for distributions and maps

All crops are weighted by production (in 1000 MT) for constructing total agriculture impact distributions and maps. Production is calculated from the USDA reported bushels measurement using 56 lbs/bushel and 60 lbs/bushel for soy and wheat, and averaged over 2000 - 2005. The data used here and in the CGE aggregations above come from reproduction data for *Schlenker and Roberts (2009)*.

7.4. Crime aggregation

Counties are weighted by the number of reported property and violent crimes from the Uniform Crime Statistics, averaged over 2000 - 2005, and provided for reproduction of *Ranson (2014)*.

Counties that are not explicitly identified at the county level (of which there are 172) are aggregated using the mean country rates of property and violent crime. Furthermore, since we use the 2010 census for county populations, these rates are scaled by 0.9339, the ratio of the average national population in 2000 - 2005 to the population in 2010, before being scaled by the individual county populations in 2010, to maintain comparability.

7.5. Labor aggregation

Labor employment by county is averaged over 2000 - 2005, as reported by the Bureau of Labor Statistics (*Bureau of Labor Statistics, 2014*). Following *Graff Zivin and Neidell (2014)* high-risk sectors consist of agriculture, forestry, fishing and hunting; mining, quarrying, and oil and gas extraction; utilities; construction; manufacturing; and transportation and warehousing. All others are considered low-risk.

The BLS statistics exclude the counties represented by FIPS codes 02105, 02195, 02198, 02230, and 02275, since these were created after 2005.

7.6. Mortality aggregation

Counties are weighted by 2010 census populations for aggregating mortality. All estimates except for the NCA region estimates used in the text use total populations, irrespective of age-cohort. For the aggregation from sub-NCA regions (in figure B17) to NCA regions used in the text, the census totals for each age are summed into impact cohorts and used for weighting.

7.7. Distributions across result sets

Within each RCP, models are weighted to capture a desired distribution of temperatures, as described in section 1.4 of Technical Appendix I. The weighting process for results for a given weather realization of a given RCP involves constructing an ECDF as follows. Let the calculated value of an impact in a given county and year for model $m \in \{1, \dots, M\}$ be I_m . The CDF of this impact across all models is,

$$F(I) = \frac{1}{\sum_m w_m} \sum_{m \text{ for } I_m \leq I} w_m$$

where w_m is the weight given to model m provided in Tables A3 - A6. Extending this process to multiple weather realizations and multiple batches is done by simply including all available results in the weighted ECDF. Unweighted results are reported in table B10.

8. Linear Extrapolation Assumption

In this section, we consider the consequences of extending impact functions linearly beyond their support. The standard implementation of the impact response functions assumes a flat extrapolation, where the coefficients used in the edge bins remain constant when extrapolated beyond the response function support. In several of the empirical impacts, the response as a function of temperature appears to grow linearly with temperature at the extreme. This is particularly true for mortality and high-risk labor productivity. Agricultural yields are implicitly linearly extrapolated (and estimated as such) since they are in terms of degree-days. Figure B18 displays how quantiles of impact responses change under a linear extrapolation assumption.

The effect of the linear extrapolation assumption is shown in the tails of the distributions and at extreme temperatures (end of century for RCP 8.5). Most impact distributions become wider, although property crime shows narrowing of impacts under linear extrapolation. This is because, while crime generally increases with temperature, it decreases at the very-high end of both temperature and precipitation. Violent crime shows very little change, because of the flat response at high temperatures.

The effects on the tails of the health and labor distributions is large: at the 90th percentile in high-risk labor, the impact increases from -0.5% to -0.7% (a 40% increase), and in mortality, the additional mortality increases from 20 per 100,000 to 24 per 100,000 (a 20% increase).

Table B10: Unweighted impact result percentiles

Percentiles	RCP 8.5					RCP 4.5					RCP 2.6				
	5	17	50	83	95	5	17	50	83	95	5	17	50	83	95
Agricultural Yields															
2080-2099	-42.44	-34.14	-16.49	0.42	13.26	-26.74	-19.37	-4.26	5.09	9.54	-16.64	-12.1	-2.99	1.34	4.57
2040-2059	-18.27	-11.31	-3.96	4.69	8.61	-14.11	-10.76	-1.59	4.83	8.28	-12.9	-9.58	-3.43	1.3	4.28
2020-2039	-9.55	-6.16	-1.92	2.58	7.98	-9.32	-6.86	-1.16	3.71	9.27	-13.61	-8.33	-2.11	1.32	2.77
Agricultural Yields (w/o CO₂fertilization)															
2080-2099	-59.76	-54.04	-40.86	-28.97	-19.24	-35.23	-28.84	-15.04	-6.71	-2.91	-19.12	-14.76	-5.99	-1.84	1.33
2040-2059	-27.49	-21.36	-14.89	-7.26	-3.91	-20.7	-17.56	-8.94	-3.01	0.15	-16.72	-13.49	-7.59	-3.09	-0.29
2020-2039	-13.59	-10.39	-6.36	-2.07	3.09	-12.45	-10.05	-4.54	0.12	5.44	-16.29	-11.1	-5.11	-1.72	-0.42
High Risk Labor															
2080-2099	-2.53	-2.06	-1.44	-0.98	-0.76	-1.18	-0.89	-0.59	-0.3	-0.13	-0.54	-0.42	-0.28	-0.11	0.18
2040-2059	-1.03	-0.77	-0.49	-0.26	-0.15	-0.74	-0.63	-0.38	-0.17	-0.02	-0.47	-0.4	-0.31	-0.13	0.15
2020-2039	-0.5	-0.38	-0.16	-0.02	0.11	-0.44	-0.37	-0.2	-0.03	0.12	-0.44	-0.33	-0.21	-0.03	0.15
Low Risk Labor															
2080-2099	-0.68	-0.46	-0.24	-0.1	0	-0.24	-0.16	-0.08	-0.03	0	-0.12	-0.08	-0.04	0.01	0.05
2040-2059	-0.22	-0.15	-0.08	-0.02	0.02	-0.15	-0.1	-0.06	-0.01	0.03	-0.11	-0.07	-0.04	0.01	0.05
2020-2039	-0.11	-0.07	-0.03	0.01	0.05	-0.1	-0.06	-0.03	0.01	0.05	-0.1	-0.06	-0.02	0.02	0.06
Mortality															
2080-2099	0.75	3.85	9.63	18.04	24.86	-2.53	-2.06	-1.44	-0.98	-0.76	-0.68	-0.46	-0.24	-0.1	0
2040-2059	-3.39	-1.27	2.48	5.89	8.75	-1.03	-0.77	-0.49	-0.26	-0.15	-0.22	-0.15	-0.08	-0.02	0.02
2020-2039	-3.82	-1.83	1.15	4.27	6.13	-0.5	-0.38	-0.16	-0.02	0.11	-0.11	-0.07	-0.03	0.01	0.05
Property Crime															
2080-2099	0.34	0.51	0.74	0.97	1.07	-0.03	0.14	0.51	0.83	1.01	-0.18	-0.01	0.23	0.46	0.61
2040-2059	-0.01	0.13	0.33	0.59	0.73	-0.13	0.03	0.36	0.67	0.85	-0.14	-0.03	0.22	0.44	0.59
2020-2039	-0.29	-0.07	0.1	0.31	0.46	-0.25	-0.09	0.19	0.48	0.63	-0.2	-0.06	0.11	0.32	0.42
Violent Crime															
2080-2099	1.84	2.2	3.05	3.76	4.41	0.21	0.74	1.54	2.31	2.54	-0.27	0.05	0.69	1.27	1.48
2040-2059	0.34	0.7	1.19	1.75	2.41	-0.02	0.42	1.09	1.65	1.84	-0.25	0.26	0.79	1.02	1.52
2020-2039	-0.29	0.04	0.45	0.9	1.34	-0.47	0.07	0.56	0.97	1.3	-0.25	0.01	0.48	0.81	1.06

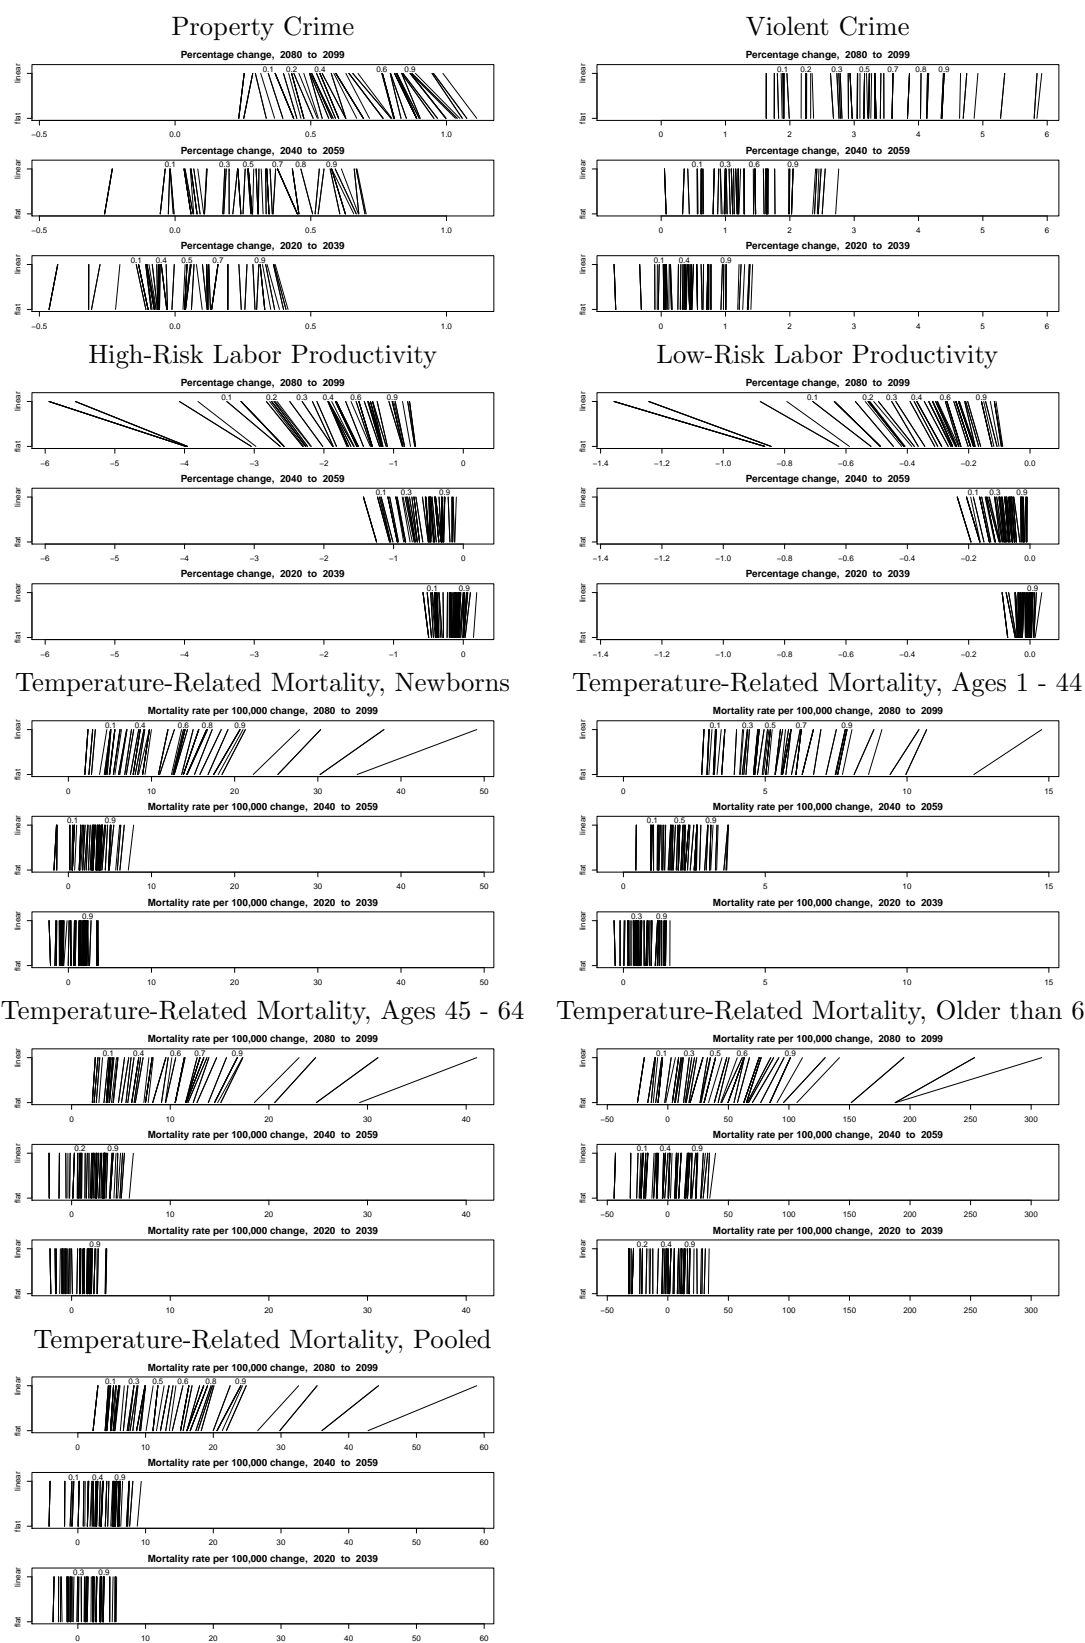


Figure B18: The evolution of impact response quantiles, under flat and linear extrapolation, under RCP 8.5 at a national level. The lines are quantiles (labeled at the top), showing how these shift between flat extrapolation (bottom) and linear extrapolation (top).

Bibliography

- Anderson, B. G., and M. L. Bell (2009), Weather-related mortality: how heat, cold, and heat waves affect mortality in the United States., *Epidemiology (Cambridge, Mass.)*, 20(2), 205–13, doi:10.1097/EDE.0b013e318190ee08.
- Anderson, G. B., and M. L. Bell (2011), Heat waves in the United States: mortality risk during heat waves and effect modification by heat wave characteristics in 43 U.S. communities., *Environmental health perspectives*, 119(2), 210–8, doi:10.1289/ehp.1002313.
- Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro (2013), Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the 20th Century, *NBER working paper*.
- Berry, S. T., M. J. Roberts, and W. Schlenker (2012), Corn Production Shocks in 2012 and Beyond : Implications for Food Price Volatility Model, (September).
- Bureau of Labor Statistics, U. D. o. L. (2014), Occupational employment statistics, <http://data.bls.gov/oes/>, accessed on Mar 20, 2014.
- Burke, M., and K. Emerick (2013), Adaptation to Climate Change : Evidence from US Agriculture.
- Center for Disease Control and Prevention (2013), Compressed Mortality File 1999-2010 on CDC WONDER Online Database, released January 2013, Data are compiled from Compressed Mortality File 1999-2010 Series 20 No. 2P, accessed at <http://wonder.cdc.gov/cmfile.html> on Mar 22, 2014.
- Curriero, F. C., K. S. Heiner, J. M. Samet, S. L. Zeger, L. Strug, and J. a. Patz (2002), Temperature and mortality in 11 cities of the eastern United States., *American journal of epidemiology*, 155(1), 80–7.
- Deschenes, O., and M. Greenstone (2007), The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather, *The American Economic ...*, 91(1), 354–385.
- Deschènes, O., and M. Greenstone (2011), Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US, *American Economic Journal: Applied Economics*, 3(October), 152–185.
- Fisher, A., W. Hanemann, M. J. Roberts, and W. Schlenker (2012), The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment, *American Economic Review*, 102(7), 3749–60.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013), *Bayesian data analysis*, CRC press.
- Graff Zivin, J., and M. Neidell (2014), Temperature and the Allocation of Time: Implications for Climate Change, *Journal of Labor Economics*, 32(1), 1–26, doi:10.1086/671766.
- Hsiang, S., D. Lobell, M. Roberts, and W. Schlenker (2013a), Climate Change and Crop Choice: Evidence from Australia, Brazil, China, Europe, and the United States, *Working paper*.
- Hsiang, S. M., M. Burke, and E. Miguel (2013b), Quantifying the influence of climate on human conflict., *Science (New York, N.Y.)*, 341(6151), 1235,367, doi:10.1126/science.1235367.
- Jacob, B., L. Lefgren, and E. Moretti (2007), The Dynamics of Criminal Behavior Evidence from Weather Shocks, *Journal of Human Resources*, 42(3).
- Lessl, M., J. S. Bryans, D. Richards, and K. Asadullah (2011), Crowd sourcing in drug discovery, *Nature Reviews Drug Discovery*, 10(4), 241–242.
- McGrath, J. M., and D. B. Lobell (2013), Regional disparities in the CO2 fertilization effect and implications for crop yields, *Environmental Research Letters*, 8(1), 014,054, doi:10.1088/1748-9326/8/1/014054.
- Ranson, M. (2014), Crime, weather and climate change, *Journal of Environmental Economics and Management*.
- Sacks, W. J., D. Deryng, J. A. Foley, and N. Ramankutty (2010), Crop planting dates: an analysis of global patterns, *Global Ecology and Biogeography*, 19(5), 607–620.
- Schlenker, W., and M. J. Roberts (2009), Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change., *Proceedings of the National Academy of Sciences of the United States of America*, 106(37), 15,594–8, doi:10.1073/pnas.0906865106.
- Seppanen, O., W. Fisk, and Q. Lei (2006), Effect of temperature on task performance in office environment, *LBNL-60946*.