Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change

Wolfram Schlenker and Michael J. Roberts

Supporting Information Appendix

Contents

1 Data Appendix

Dependent Variable

Yields for corn, soybeans, and cotton for the years 1950-2005 are reported by the U.S. Department of Agriculture's National Agricultural Statistical Service (USDA-NASS). We include all reported yields, even though some appear spuriously low. These outliers are few. If we drop outliers, results do not change, but the cutoff point for omission becomes somewhat arbitrary, so we include them. The descriptive statistics are:

Table A1: Descriptive Statistics: Crop Yields

Crop	**Mean**	**Minimum**	**Maximum**	**Std. Dev.**
Corn (Bushels/Acre)	75.60	0.04	203.49	35.52
Soybeans (Bushels/Acre)	27.00	1.80	249.39	9.35
Cotton (Pounds/Acre)	451.48	8.00	3433.00	234.86

Yields equal total county-level production divided by harvested acres. For about 80% of the observations, NASS reports planted acres in addition to harvested acres. Because planted acres may not be harvested in a low-yielding year, yield per harvested acre may cause bias. As a sensitivity check, we derived an alternative yield measure by taking total production over total acres planted, with results reported in Section 11 below. The results are very similar to the results using harvested acres. Since the area planted is not reported in all areas and years, our analysis focuses on the larger sample of output per acre harvested that is the standard USDA definition of yield.

In the main paper we report results for corn and soybean yields in counties east of the 100 degree meridian. We do this because counties west of the 100 degree meridian are supported by large and heavily subsidized irrigation systems. For cotton the main paper reports results for all cotton yields, since most cotton (58%) is irrigated and the sample of counties is much smaller. This gives us 105981 observations with corn yields, 82385 observations with soybeans yields, and 31540 observations with cotton yields.

Given the potential importance of irrigation and particularly how irrigation might differentially affect yield outcomes in eastern and wester regions, Section 7 reports separate results for counties east and west of the 100th meridian for all three crops. Results differ substantially for corn and soybeans, but to a lesser degree for cotton, and we hence pool observations from the East and West for the latter.

Past Weather

Earlier statistical studies use average temperature over an entire season or months, which can hide extreme temperatures that occur within a month or even during a fraction of a day. Our fine-scale weather aids identification of nonlinear effects which are diluted when weather outcomes are averaged over time or space. Construction of these data is briefly described here and in more detail in [1].

The basic steps are as follows. We first develop daily predictions of minimum and maximum temperature on a 2.5x2.5mile grid for the entire United States. We then derive the time a crop is exposed to each 1 degree Celsius interval in each grid cell. These predictions are merged with a satellite scan that allows us to select only those grid cells with cropland. We then aggregate the whole distribution of outcomes for all days in the growing season in each county. Since our study emphasizes nonlinearities, it is important to derive the time each grid cell is exposed to each 1 degree Celsius interval before aggregating to obtain the county-level distribution in each growing season. This preserves within-county variation in temperatures in each year.[1]

We obtain geographic specificity of weather outcomes using the Parameter-elevation Regressions on Independent Slopes Model (PRISM), widely regarded as one of the best geographic interpolation procedures (http://www.ocs.orst.edu/prism/). PRISM accounts for elevation and prevailing winds to predict weather outcomes on 2.5x2.5 mile grid across the contiguous United States. The PRISM data, however, are reported on a *monthly* time scale. We therefore combine the advantages of the PRISM model (good spatial interpolation) with better temporal coverage of individual weather stations (daily instead of monthly values). We do this by pairing each of the 259,287 PRISM grid cells that cover agricultural area in a LandSat satellite scan with the closest seven weather stations having a continuous record of daily observations.[2] We then estimate a separate regression for each grid cell, where the dependent variable is the monthly PRISM grid cell estimate and the explanatory variables

[1] An alternative method to approximate the distribution of daily temperatures from the distribution of average monthly temperatures is developed by [2]. This method appears appropriate for predicting the *average* frequency that a certain weather outcome will be realized, but less appropriate in predicting a *specific* frequency of a weather outcome in a particular year. As a result, these methods work well in a forward-looking cross-sectional analysis where the dependent variable is tied to weather expected outcomes rather than realized outcomes (for example, the link between land values and climate), but less well in our analysis where the dependent variable (yield) linked to specific weather outcomes. Obtaining daily values on a small scale requires a spatial interpolation procedure to approximate daily weather outcomes between individual weather stations.

[2] Stations have a continuous record if at least 90 percent of all months have at most 3 missing observations. Missing values are filled in by regressing the daily observations at a station on the seven closest weather stations including half-month fixed effects.

are the monthly averages at each of the seven closest weather stations, plus fixed effects for each month. The R-squares are usually in excess of 0.999. The estimated relationship between *monthly* PRISM grid cell averages and *monthly* averages at each of the seven closest stations is then used to predict *daily* records at each PRISM grid cell from the *daily* records at the seven closest weather stations.

A cross-validation exercise is used to test the accuracy of the daily weather predictions. For this exercise we construct a daily weather record at each PRISM cell that harbors a weather station without using that weather station in the interpolation procedure. We then compare predicted daily outcomes to actual outcomes within all PRISM cells that harbor weather stations. The mean absolute error is 1.36°C for minimum temperature and 1.49°C for maximum temperature. Due to the law of large numbers, our county-level distribution estimates contain much less error because they average errors over all grid cells in each county and all days of the growing season.

The distribution of temperatures within each day is approximated using a sinusoidal curve between predicted minimum and predicted maximum temperatures [3]. In a sensitivity check, reported in Section 13, we instead use a linear interpolation between minimum and maximum temperature. Both methods give similar results. Using the sinusoidal or linear interpolation we then estimate time in each 1°C-degree temperature interval between -5°C and $+50$°C. Finally, we construct the area-weighted average time at each degree over all PRISM grid cells in a county. The agricultural area in each cell was obtained from LandSat satellite images.[3] The weather variables are summed over the six-month period from March through August for corn and soybeans, and the seven-month period April through October for cotton.

Boxplots in Figure A1 show historical temperature distributions and how they vary over time and across counties. Whiskers indicate the minimum and maximum exposure to a each temperature range across all counties. The box marks the 25%-75% range, and the middle line within each box is the median across all counties and years.

Climate Change Scenarios

Climate change predictions are drawn from the Hadley III model.[4] It is one of the models that form the basis for the report by the Intergovernmental Panel on Climate Change (IPCC).

[3] Vince Breneman and Shawn Bucholtz at the Economic Research Service provided us with the agricultural area in each PRISM grid cell. Since we use the LandSat scan of a given year, we are not able to pick up shifts in growing regions.

[4] http://www.metoffice.com/research/hadleycentre/

Figure A1: Descriptive Weather Statistics

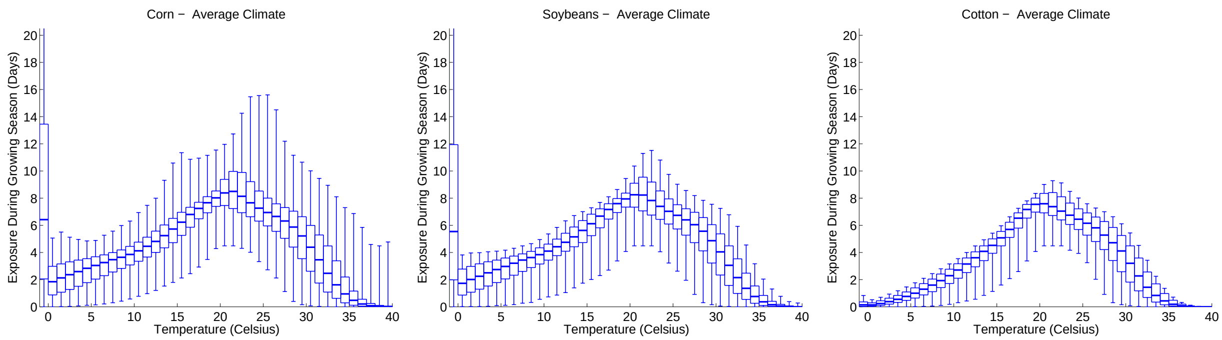

Notes: Graphs show distributions of time each crop was exposed to each 1°C interval during the 184-day growing season March-August for corn and soybeans and the 214-day growing season April-October for cotton. The lowest interval has no lower bound and includes the time temperatures fall below 0°C. The highest interval has no upper bound and includes the time temperatures are above 39°C. For each temperature interval, the range between minimum and maximum time across counties and years is shown by whiskers, the 25%-75% percentile range is outlined by a box, and the median is added as a solid bold line.

We obtain monthly model output for both minimum and maximum temperatures under four major emissions scenarios (A1FI, A2, B1, and B2) for the years 1960-2099. Each emission scenario rests on different assumptions about population growth and use of alternative fuels, among other factors [4]. The B1 scenario assumes CO_2 emissions soon decline and therefore results in the slowest rate of warming over the next century. The A1FI assumes continued use of fossil fuels and the largest increase in CO_2-concentrations and temperatures. The other two scenarios are between these extremes.

Predicted weather under climate change is derived as follows. At each of 216 Hadley grid nodes covering the United States we find the predicted difference in monthly mean temperature for 2020-2049 (medium-term), 2070-2099 (long-term), and historic averages (1960-1989). Next, predicted changes in monthly minimum and maximum temperature at each 2.5x2.5 mile PRISM grid are calculated as the weighted average of the monthly mean change in the four surrounding Hadley grid points, where the weights are proportional to the inverse squared distance and forced to sum to one. In a final step, we add the predicted *absolute* changes in monthly minimum and maximum temperatures at each PRISM grid to observed daily time series from 1960 to 1989. In other words, we shift the entire historical distribution of each month in each grid cell by the predicted mean monthly change of each climate scenario. An analogous approach was used for precipitation, except that we use the *relative* ratio of future predicted rainfall to historic rainfall instead of absolute changes. Each

Figure A2: Climate Change Predictions under the Hadley HCM3 Model

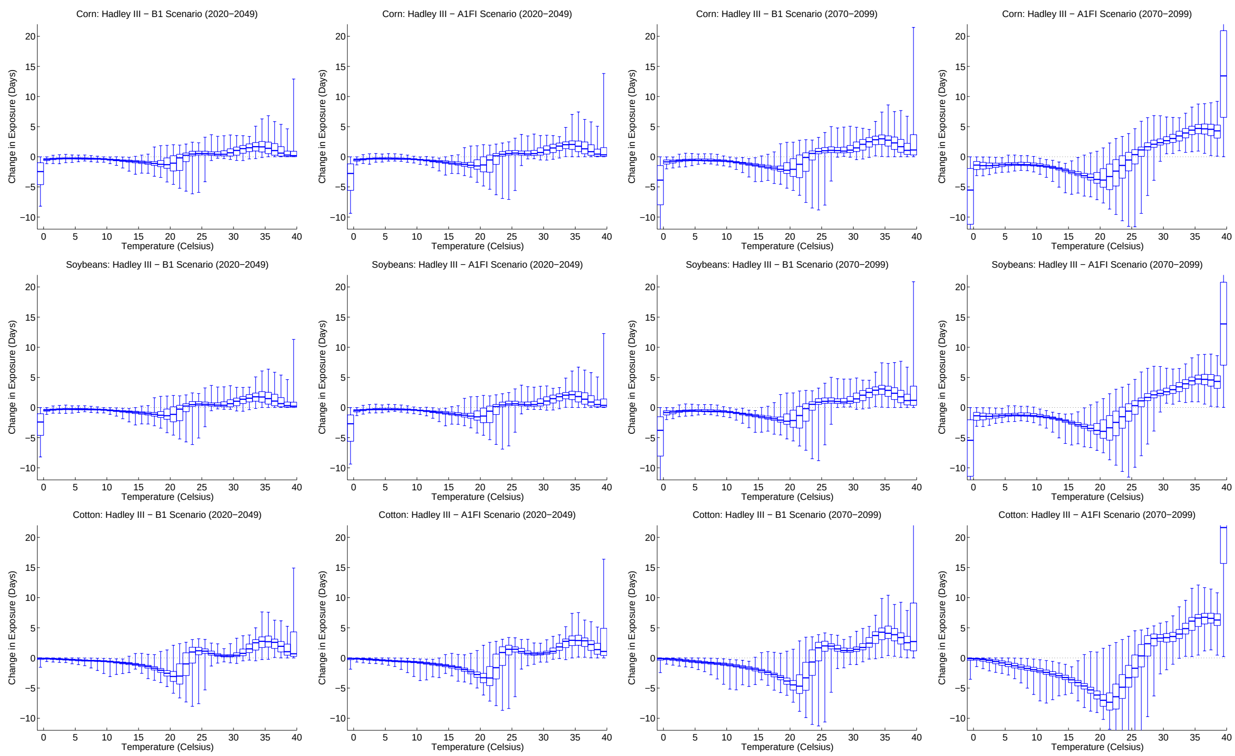

Notes: Graphs show the distribution of predicted temperature changes for each crop and each 1°C interval during the March-through-August 184-day growing season for corn and soybeans and the April-through-October 214-day growing season for cotton. The lowest interval has no lower bound and includes the change in time that temperatures fall below 0°C. The highest temperature interval has no upper bound and includes the change in time that temperatures are above 39°C. For each temperature interval, whiskers show the range of duration changes across counties, the box outlines the 25%-75% percentile of changes, and the solid bold line inside each box shows the median change. The top row displays predicted changes for corn-growing counties, the middle row shows soybeans, and the bottom row shows cotton. The first two columns show slow-warming (B1) and fast-warming (A1FI) predictions for the medium term (2020-2049) and last two columns show slow- and fast-warming predictions for the long-term (2070-2099).

county's weather outcomes are the area-weighted averages of all PRISM grids.

Figure A2 shows the shift in the temperature distribution under the B1 and A1FI scenarios in the medium-term (2020-2049) and long-term (2070-2099) for each of the three crops in our data. Each figure shows a series of box plots, one for each degree Celsius. Each box-plot summarizes the predicted change in the frequency of that specific temperature across all counties growing that crop. Temperatures below 22°C generally become less frequent in corn and soybeans counties and temperatures below 25°C become less frequent in cotton counties. Temperatures above these levels generally become more frequent.

2 Regression Models

The regression models assume temperature effects on yields are cumulative over time and that yield is proportional to total exposure. This implies the temperature effect on yield growth rates are the same and additively substitutable over time. The empirical validity of this assumption is considered in more detail below in Section 10. Specifically, yield growth $g(h)$ depends nonlinearly on heat h such that log yield, y_{it}, in county i and year t is

$$y_{it} = \int_{\underline{h}}^{\overline{h}} g(h)\phi_{it}(h)dh + \mathbf{z}_{it}\boldsymbol{\delta} + c_i + \epsilon_{it} \tag{A1}$$

where $\phi_{it}(h)$ is the time distribution of heat over the growing season in county i and year t. We fix the growing season to months March through August for corn and soybeans and the months April through October for cotton. Observed temperatures during this time period range between the lower bound $\underline{h}$ and the upper bound $\overline{h}$. Other control factors are denoted $\mathbf{z}_{it}$ and include a quadratic in total precipitation as well as a quadratic time trend for each state to capture technological change. A time-invariant county fixed effect c_i is to control for heterogeneity, such as soil type and quality. We allow the error terms ϵ_{it} to be spatially correlated using the non-parametric routine by [5].

Although time separability is partially rooted in agronomy, we implicitly validate this assumption by showing a statistically significant relationship between the cumulative distribution of temperatures and yields. We would not observe this if time separability were not appropriate, because random pairing of various temperatures over a season and between years would not provide clear identification. In Section 10 we consider the assumption of time-additivity of temperature effects in more detail. These alternative models estimate a separate growth function $g(h)$ for various sub-periods of the growing season. In one alternative specification we split the six-month growing season into three two-month intervals and estimate a growth function for each one of them by omitting the remaining four months. In a second specification we jointly estimate a separate growth function for each month of the growing season. In a third specification we test for whether the temperature response function is different July than it is in other months. These considerably more flexible alternative models do little to improve the overall fit of the model and give similar predictions under climate-change scenarios.

A special case of time-separable growth is the concept of growing degree days, typically defined as the sum of truncated degrees between two bounds. For example, bounds of 8°C and 32°C for "beneficial heat" have been suggested by [6]. A day of 9°C hence contributes 1 degree day, a day of 10°C contributes 2 degree days, up to a temperature of 32°C, which contributes 24 degree days. All temperatures above 32°C also contribute 24 degree days. Degree days are then summed over the entire season. These particular bounds have been implemented in a cross-sectional analysis by [7]. Thus, growing degree days are the special case of our model where (using the above bounds as an example)

$$g(h) = \begin{cases} 0 & \text{if } h \leq 8 \\ h - 8 & \text{if } 8 < h < 32 \\ 24 & \text{if } 32 \leq h \end{cases}$$

The appropriate bounds for growing degree days are still debated, partly because earlier studies use a limited number of observations from field experiments to identify them. There is also uncertainty about temperature effects above the upper bound. Degree days above 34°C are sometimes included as a separate variable and speculated to be harmful.

Using data on exposure to each 1-degree Celsius temperature interval, we approximate the above integral with

$$y_{it} = \sum_{h=-5}^{49} g(h + 0.5)[\Phi_{it}(h + 1) - \Phi_{it}(h)] + \mathbf{z}_{it}\boldsymbol{\delta} + c_i + \epsilon_{it} \tag{A2}$$

where $\Phi_{it}(h)$ is the cumulative distribution function of heat in county i and year t.

We consider three specifications for $g(h)$. First, we approximate $g(h)$ using dummy variables for each three-degree temperature interval. We obtain similar results when estimating more flexible models with dummy variables for each one-degree interval. We report results for three-degree intervals in order to make figures easier to interpret. This step function effectively regresses yield on season-total time within each temperature interval. Temperatures above 39°C (102 degrees Fahrenheit) occur less frequently and we therefore lump all time a plant is exposed to a temperature above 39°C into one category. Similarly, we lump all time where temperature is below freezing into the interval $[-1; 0]$. The historical temperature distribution is displayed in Figure A1. The model becomes[5]

$$y_{it} = \sum_{j=0,3,6,9,\ldots}^{39} \gamma_j \underbrace{\left[\Phi_{it}(h+3) - \Phi_{it}(h)\right]}_{x_{it,j}} + \mathbf{z}_{it}\boldsymbol{\delta} + c_i + \epsilon_{it}. \tag{A3}$$

The second specification assumes $g(h)$ is an m-th order Chebychev polynomial of the form $g(h) = \sum_{j=1}^{m} \gamma_j T_j(h)$, where $T_j()$ is the $j-th$ order Chebyshev polynomial. Chebyshev polynomials are a relatively parsimonious approximation for the function $g(h)$, assuming it is smooth. By interchanging the sum we obtain

$$\begin{aligned} y_{it} &= \sum_{h=-1}^{39} \sum_{j=1}^{m} \gamma_j T_j\left(h+0.5\right)\left[\Phi_{it}\left(h+1\right) - \Phi_{it}\left(h\right)\right] + \mathbf{z}_{it}\boldsymbol{\delta} + c_i + \epsilon_{it} \\ &= \sum_{j=1}^{m} \gamma_j \underbrace{\sum_{h=-1}^{39} T_j\left(h+0.5\right)\left[\Phi_{it}\left(h+1\right) - \Phi_{it}\left(h\right)\right]}_{x_{it,j}} + \mathbf{z}_{it}\boldsymbol{\delta} + c_i + \epsilon_{it} \end{aligned} \tag{A4}$$

where $x_{ij,t}$ is the exogenous variable obtained by summing the $j-th$ Chebyshev polynomial evaluated at each temperature interval midpoint, multiplied by the time spent in each temperature interval. Successively higher-order polynomials were estimated until the relationship appeared stable.

The third specification models $g(h)$ as piecewise linear function, which is similar to the concept of growing degree days. In this case growth increases linearly up to an endogenous threshold and then decreases linearly above the threshold. We loop over all possible thresholds, estimate the least-squares segment slopes for each one, and pick the threshold and segment slopes with the best fit.

The predicted change in production when the average weather variables change from the 1960-1989 average $\bar{\mathbf{w}}_{i0}$ to the new values $\bar{\mathbf{w}}_{i1}$ is using the regression coefficients on the weather variables $\boldsymbol{\beta}_w$:[6]

$$\text{impact} = \frac{\sum_{i=1}^{N} a_i e^{\bar{\mathbf{w}}_{i1}\boldsymbol{\beta}_w + \bar{c}_i}}{\sum_{i=1}^{N} a_i e^{\bar{\mathbf{w}}_{i0}\boldsymbol{\beta}_w + \bar{c}_i}} - 1$$

where a_i is the average growing area for counties $i = 1 \ldots N$ and c_i is the county fixed effect.

[5]The omitted category, captured by the model's intercept, is the time temperature is below 0°C.

[6]For some sensitivity checks, we use regression coefficients estimated from one subset of data and use them to make climate-change predictions for another subset of data. For example, in one sensitivity check we use southern counties to estimate $\boldsymbol{\beta}_w$ and then use these coefficient estimates to predict climate-change impacts for all counties east of the 100 degree meridian.

It should be noted that the $\bar{c}_i$ act like another scale factor besides the area a_i. If we omit them, the results differ by at most 1%, so the reweighing has a minor impact on the results.

3 Regression Results and Significance Levels

Regression results are displayed in Figure 1 in the main paper. Here we replicate the analysis and present each of the three model specifications as a separate row in Figure A3. We also

Figure A3: Nonlinear Relation Between Temperature and Yields With Confidence Bands

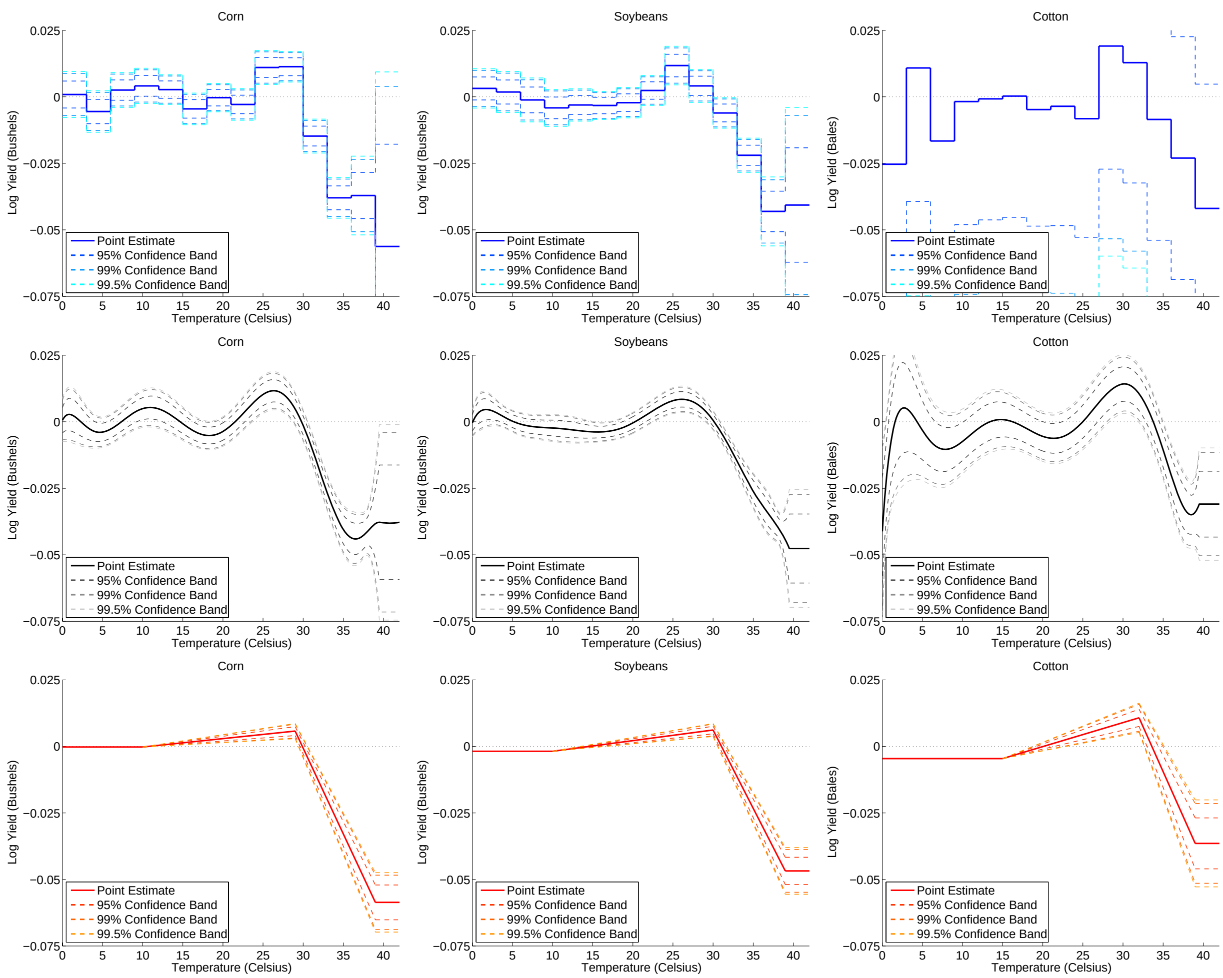

Notes: Graphs display changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95%, 99%, and 99.5% confidence bands, after adjusting for spatial correlation, are added. Curves are centered so the exposure-weighted impact is zero.

include three confidence bands given that our data set is rather large: 95%, 99% and 99.5%. Note that even for the latter, the piecewise linear function shows a statistically significant increase in temperature followed by a statistically significant decrease.

4 Out-of-Sample Prediction Accuracy

Tables A2-A4 report encompassing tests that compare our model's predictions to others in the literature. Comparisons are based on out-of-sample forecast accuracy. Each model is estimated 1000 times, where each replication randomly selects 48 of the 56 years in our full sample. Relative performance is measured according to the accuracy of each model's prediction for the omitted 8 years of the sample (about 14 percent). We sample years instead of observations as year-to-year weather fluctuations are random, but there is considerable spatial correlations across counties within a year.[7] Models include our own three specifications of temperature effects (step function, polynomial (8th order), and piecewise linear), a model with average temperatures for each of four months [8], an approximation of growing-degree days based on monthly average temperatures (Thom's formula) used in [7], and a measure of growing degree days that is calculated using daily mean temperatures used by [9].[8] As a baseline, we also report a model with county fixed effects and time trends but no weather variables. Comparisons are made using the root-mean squared prediction error (RMS) and Welch t-tests against the null hypothesis that the RMS in our 1000 replications is the same for any comparison of two models.

Models that average temperatures over time or space have significantly inferior out-of-sample predictions relative to our new three models. There is little difference in forecasting ability between the step function, polynomial, and piecewise linear model, but large and statistically significant differences between these three models and the other models. For all three crops – corn, soybeans, and cotton – our three new specifications do not give significantly different RMS, yet when paired with any of the other models, the reduction in RMS is statistically significant at the 0.1% level.

[7] We would like to thank a referee for pointing this out.

[8] We use degree days bounds of each cited study, but the ranking of models does not change if instead we were to use bounds suggested by this study.

Table A2: Model Comparison Test for Out-Of-Sample Prediction Accuracy for Corn

	RMS	**Welch Test for Equal Forecasting Accuracy (RMS)**				
		Step Function	**Piecewise Linear**	**Monthly Averages**	**Degree Days (Thom)**	**Degree Days (Daily Mean)**
Polynomial (8th-order)	0.2259	0.10	0.68	13.94	23.58	25.76
Step Function	0.2260		0.58	13.90	23.57	25.75
Piecewise Linear	0.2265			13.49	23.24	25.44
Monthly Averages	0.2387				9.22	12.02
Degree Days 8-32°C, >34°C (Thom)	0.2481					3.18
Degree Days 8-32°C (Daily Mean)	0.2516					
County-Fixed Effects (No Weather)	0.2688					

Notes: Table compares alternative temperature specifications for corn according to out-of sample prediction accuracy. The first column reports the average root mean squared out-of sample prediction error (RMS) from 1000 replications. Rows are sorted from best forecast performance (lowest average RMS) to worst. The last five columns present pair-wise Welch t-tests against the null hypothesis of equal RMS.

Step Function, **Polynomial (8th order)**, and **Piecewise Linear** are the models developed in this paper; **Monthly Averages** uses a quadratic specification in both average temperature and total precipitation for the months January, April, July, and October [8]; **Degree Days 8-32°C, >34°C (Thom)** uses Thom's formula to estimate degree days (which are based on daily data) from monthly average temperature data [7]; **Degree Days 8-32°C (Daily Mean)** first derive the average temperature for each day from daily temperature readings and then construct degree days from this average [9];**County fixed Effects (No weather)** is a baseline model with county fixed effects and state-level trends but no weather measures.

Table A3: Model Comparison Test for Out-Of-Sample Prediction Accuracy for Soybeans

	RMS	Welch Test for Equal Forecasting Accuracy (RMS)				
		Step Function	Piecewise Linear	Monthly Averages	Degree Days (Thom)	Degree Days (Daily Mean)
Polynomial (8th-order)	0.1962	0.78	0.89	22.62	23.86	30.75
Step Function	0.1966		0.10	21.86	23.07	30.04
Piecewise Linear	0.1967			21.96	23.19	30.17
Monthly Averages	0.2096				0.66	9.32
Degree Days 8-32°C, >34°C (Thom)	0.2100					8.86
Degree Days 8-32°C (Daily Mean)	0.2161					
County-Fixed Effects (No Weather)	0.2300					

Notes: Table compares alternative temperature specifications for soybeans according to out-of sample prediction accuracy. The first column reports the average root mean squared out-of sample prediction error (RMS) from 1000 replications. Rows are sorted from best forecast performance (lowest average RMS) to worst. The last five columns present pair-wise Welch t-tests against the null hypothesis of equal RMS.
Step Function, **Polynomial (8th order)**, and **Piecewise Linear** are the models developed in this paper; **Monthly Averages** uses a quadratic specification in both average temperature and total precipitation for the months January, April, July, and October [8]; **Degree Days 8-32°C, >34°C (Thom)** uses Thom's formula to estimate degree days (which are based on daily data) from monthly average temperature data [7]; **Degree Days 8-32°C (Daily Mean)** first derive the average temperature for each day from daily temperature readings and then construct degree days from this average [9]; **County fixed Effects (No weather)** is a baseline model with county fixed effects and state-level trends but no weather measures.

Table A4: Model Comparison Test for Out-Of-Sample Prediction Accuracy for Cotton

	RMS	**Welch Test for Equal Forecasting Accuracy (RMS)**				
		Polynomial (8th-order)	**Step Function**	**Degree Days (Thom)**	**Monthly Averages**	**Degree Days (Daily Mean)**
Piecewise Linear	0.3350	1.41	1.55	3.85	4.41	5.89
Polynomial (8th-order)	0.3379		0.14	2.35	2.87	4.33
Step Function	0.3382			2.20	2.72	4.17
Degree Days 8-32°C, >34°C (Thom)	0.3425				0.50	2.04
Monthly Averages	0.3434					1.57
Degree Days 8-32°C (Daily Mean)	0.3462					
County-Fixed Effects (No Weather)	0.3484					

Notes: Table compares alternative temperature specifications for cotton according to out-of sample prediction accuracy. The first column reports the average root mean squared out-of sample prediction error (RMS) from 1000 replications. Rows are sorted from best forecast performance (lowest average RMS) to worst. The last five columns present pair-wise Welch t-tests against the null hypothesis of equal RMS.

Step Function, **Polynomial (8th order)**, and **Piecewise Linear** are the models developed in this paper; **Monthly Averages** uses a quadratic specification in both average temperature and total precipitation for the months January, April, July, and October [8]; **Degree Days 8-32°C, >34°C (Thom)** uses Thom's formula to estimate degree days (which are based on daily data) from monthly average temperature data [7]; **Degree Days 8-32°C (Daily Mean)** first derive the average temperature for each day from daily temperature readings and then construct degree days from this average [9]; **County fixed Effects (No weather)** is a baseline model with county fixed effects and state-level trends but no weather measures.

5 Nonlinearities and Climate Change Impacts

The last section showed how a model that accounts for nonlinear temperature effects improves model accuracy as measured by the out-of-sample prediction error. This section shows how an account of nonlinear temperature effects influences predictions under projected climate change scenarios.

Figure A4 replicates Figure 2 of the main paper except that it adds the results for a linear model in average growing season temperature. All other controls (a quadratic in precipitation, quadratic time trends by state, and county fixed effects) are the same. The linear temperature term therefore should pick up the marginal effect at the sample mean. Since predicted climate change can lead to very significant non-marginal changes, the linear model can be misleading. The predicted climate change impacts in Figure A4 are much lower and statistically significantly different. Accounting for the nonlinear effect of temperatures has a big effect on predicted climate change impacts.

Figure A4: Predicted Climate Change Impacts on Crop Yields under the Hadley III Climate Model

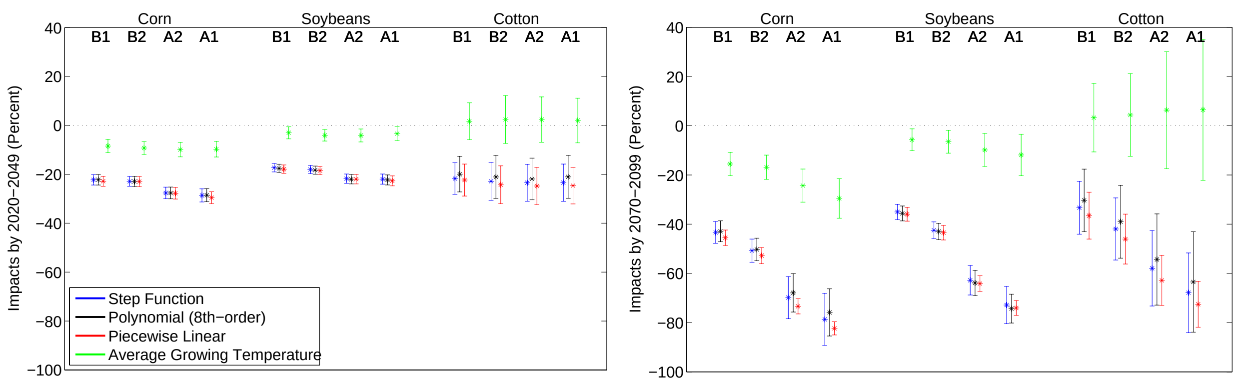

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. The left panel shows predicted impacts in the medium-term (2020-2049) and the right panel shows the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The color corresponds to the regression models in Figures 1 and 2 of the main paper. A model using average temperature over the growing season is added in green.

6 Cross-section versus Time Series

A panel data analysis with county fixed effects is identified from time-series variation in weather, which cannot account for grower adaptation to changes in climate. A clear advantage of the panel-data analysis is that year-to-year variations in weather are plausibly random to farmers, so this approach embodies a viable natural experiment. In contrast, a cross-sectional analysis compares yields and grower management choices across areas with typically different weather (i.e., different climates). While this approach can accounts for a wide array of grower adaptations, much like a hedonic model of land values, a potential problem with cross sectional analysis is confounding from omitted variables associated with location and climate. Unlike year-to-year weather variations, climate variations are less plausibly random. It is informative to consider both methods of identification and compare them.

This comparison, replicated for each of the three crops, is presented in Figure A5. The cross-sectional regression fits each county's average deviation from the nationwide average yield as a function of each county's average temperature distribution. Since the cross-sectional regression is most susceptible to omitted variables biases, we also consider a specification with soil controls as used in [7]. The time-series regression fits nationwide average yield as a function of each year's acre-weighed average temperature distribution. Another

Figure A5: Nonlinear Relation Between Temperature and Yields Using Various Sources of Identification

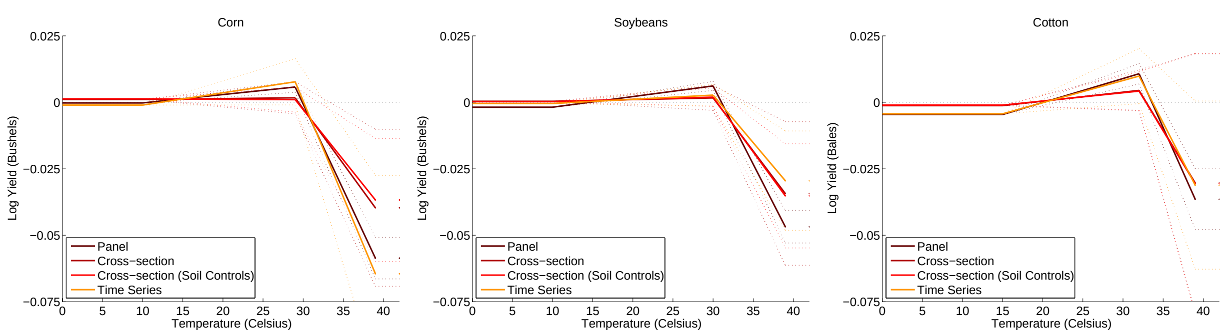

Notes: Graphs show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where within-day temperature variations are counted. Each graph shows results for the full panel, pure cross-section, and pure time-series as the sources of identification. The fitted curves are centered so the exposure-weighted impact is zero. Precipitation and other controls are included in estimation but not reported.

regression uses the full panel. Since there are only 56 observations in the time series, comparisons are made using the piecewise linear model, which has just two degrees of freedom for temperature.

The similarity of the estimated temperature growth functions is notable given the very different variations used to identify them. It seems unlikely that omitted variables could arise in a manner to indicate the same nonlinear relationship in the cross-section as the time series. Holding the locations where crops are grown fixed, this gives a strong indication that growers have historically been unable to easily adapt to warmer temperatures.

The cross-section and time-series models also make similar predictions under climate

Figure A6: Predicted Climate Change Impacts under the Hadley III Climate Model Using Various Sources of Identification

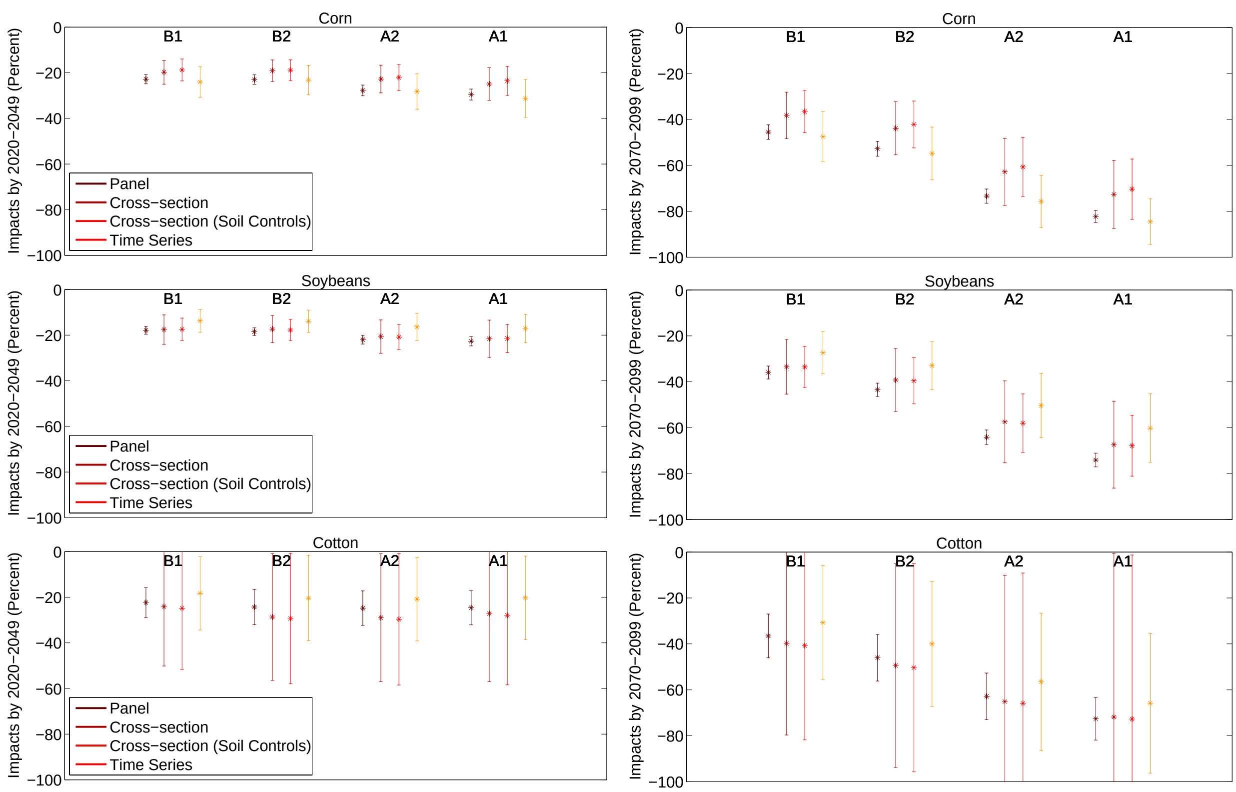

Notes: Graphs show predicted percent changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The line labeled "panel" replicates Figure 2 in the main paper. Lighter colors show predictions from the pure cross-section and pure time-series coefficients.

change. These comparisons are shown in Figure A6, except that standard errors become larger in the case of cotton. Again, this suggests adaptation possibilities have either been too costly or have not been available. Farmers might, of course, change the crops they grow. And future adaptation possibilities may differ from those in the past.

7 Geographic Subregions

7.1 North versus South

To test the sensitivity of the temperature-growth function to location, we divided the sample into three mutually exclusive geographical regions: the most northern (and coolest) states, the most southern (and warmest) states, and those in the middle.

If farmers can adapt to permanently higher temperatures by growing different varieties, one would expect yields in the South would be less sensitive to extreme heat. We replicate

Figure A7: Geographic Subsets

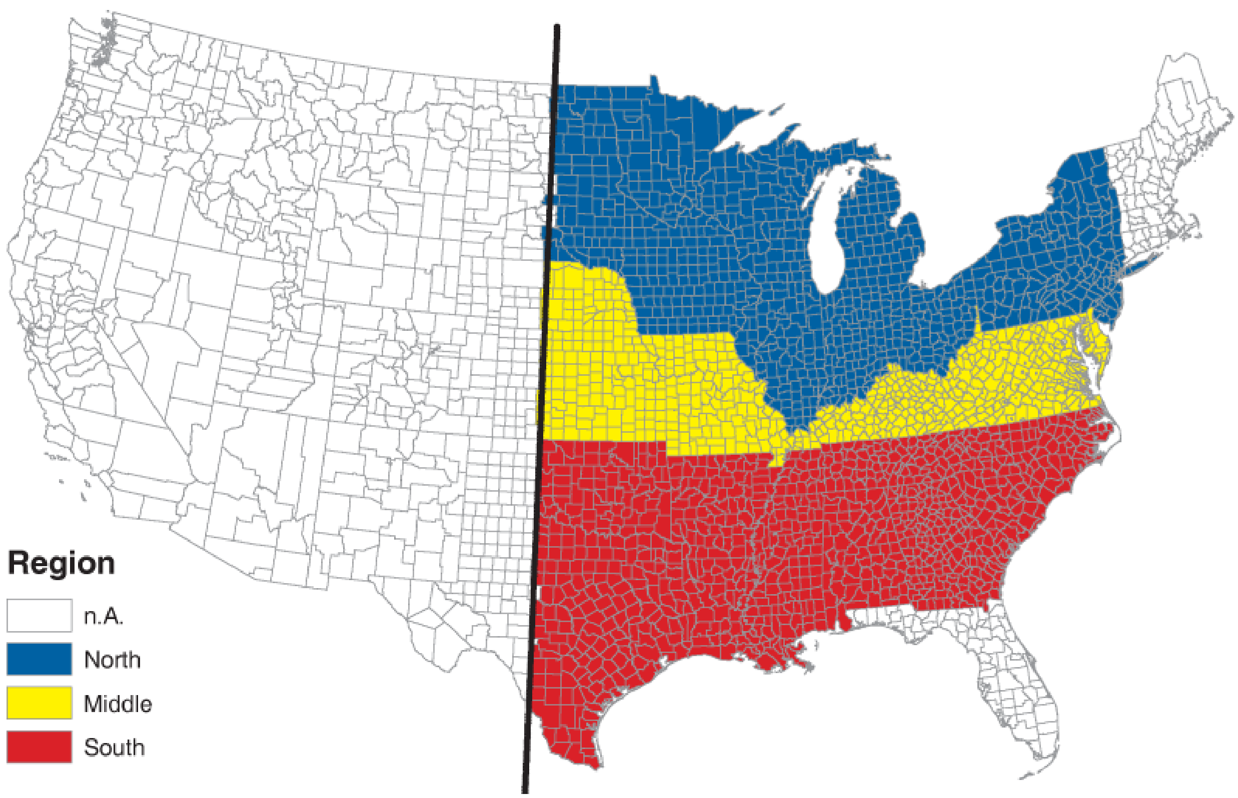

Figure 1 from the main paper for each subregion. Each plot also includes the empirical distribution of temperatures within each subregion as green histogram. These show how much warmer southern counties are in comparison to the northern counties.

When moving to hotter counties (columns on the right), there are two countervailing factors. First, the threshold where temperatures become harmful becomes slightly lower in warmer areas (the threshold in the piecewise linear regression is 30°C in the northern subsample, 29°C in the middle, and 28°C in the southern subsample. Second, the slope of the decline seems to be less steep in the warmer south. It is an empirical question which of the two effects will dominate. We hence conduct a policy experiment where we assume that all farmers are growing the corn varieties from the South and face their response function. We replicate the climate change impacts of Figure 2 in the main paper by evaluating the impacts for the *same counties*, but alternatively using only counties in the South to identify the regression coefficients. Predicted damages are approximately the same as shown in Figure A9.

Figure A8: Nonlinear Relation Between Temperature and Yields For Various Geographic Subregions

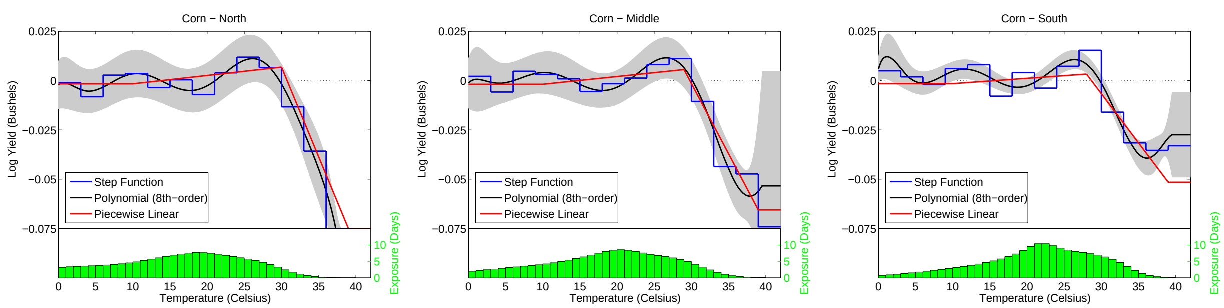

Notes: Graphs at the top of each panel display changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. The three panels from left to right, respectively, only use northern, middle, and southern counties in the estimation of the coefficients. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A9: Predicted Climate Change Impacts under the Hadley III Climate Model

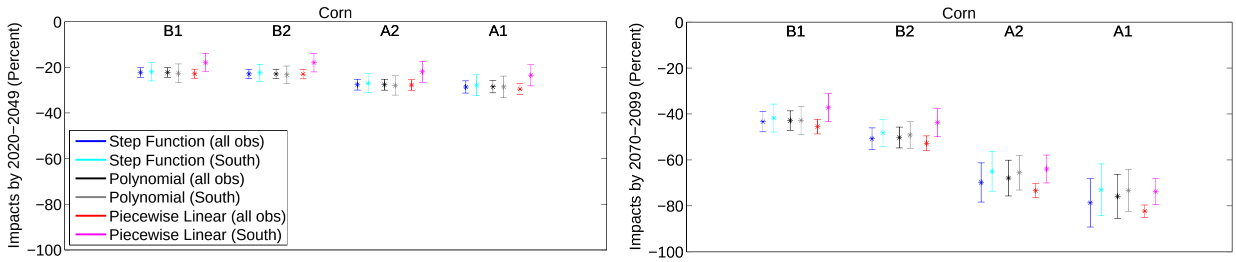

Notes: Graphs show predicted percent changes in crop yields under four emissions scenarios. The left panel displays predicted impacts in the medium-term (2020-2049) and the right panel in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The darker colors replicate Figure 2 in the main paper. Corresponding lighter colors only use the southern counties in the estimation of the regression coefficients but then evaluate the climate change impacts for the same counties as in Figure 2 of the main paper.

7.2 East versus West

Our baseline model uses corn and soybean yields in counties east of the 100 degree meridian. While some eastern states rely on irrigation (for example, 79 percent of the corn area in Arkansas was irrigated in 2007)[9], there is an essential difference between irrigation in the East and West. The West relies on large subsidized irrigation infrastructures to capture and carry large amounts of surface water, while the East relies predominantly on groundwater irrigation systems. Because groundwater irrigation is a possible mitigation strategy, we want to include such counties in the analysis, but not counties with irrigation systems that would seem unlikely to be built in the future (due to both physical and political constraints).

Figure A10 shows the estimated relationship for counties east of the 100 degree meridian (right column) and west of the 100 degree meridian (left column). While the effect of extreme temperature (difference in height of the downward-sloping portion) is different in the East and West for corn and soybeans, it is less so for cotton. We therefore pool all counties for cotton, which is highly irrigated even in the East. The percent of soybeans, corn, and cotton irrigated in the United States is 9%, 18%, and 58%, respectively.

[9]Census of Agriculture, Volume 1, Chapter 2, Table 26.

Figure A10: Nonlinear Relation Between Temperature and Yields For Eastern and Western United States

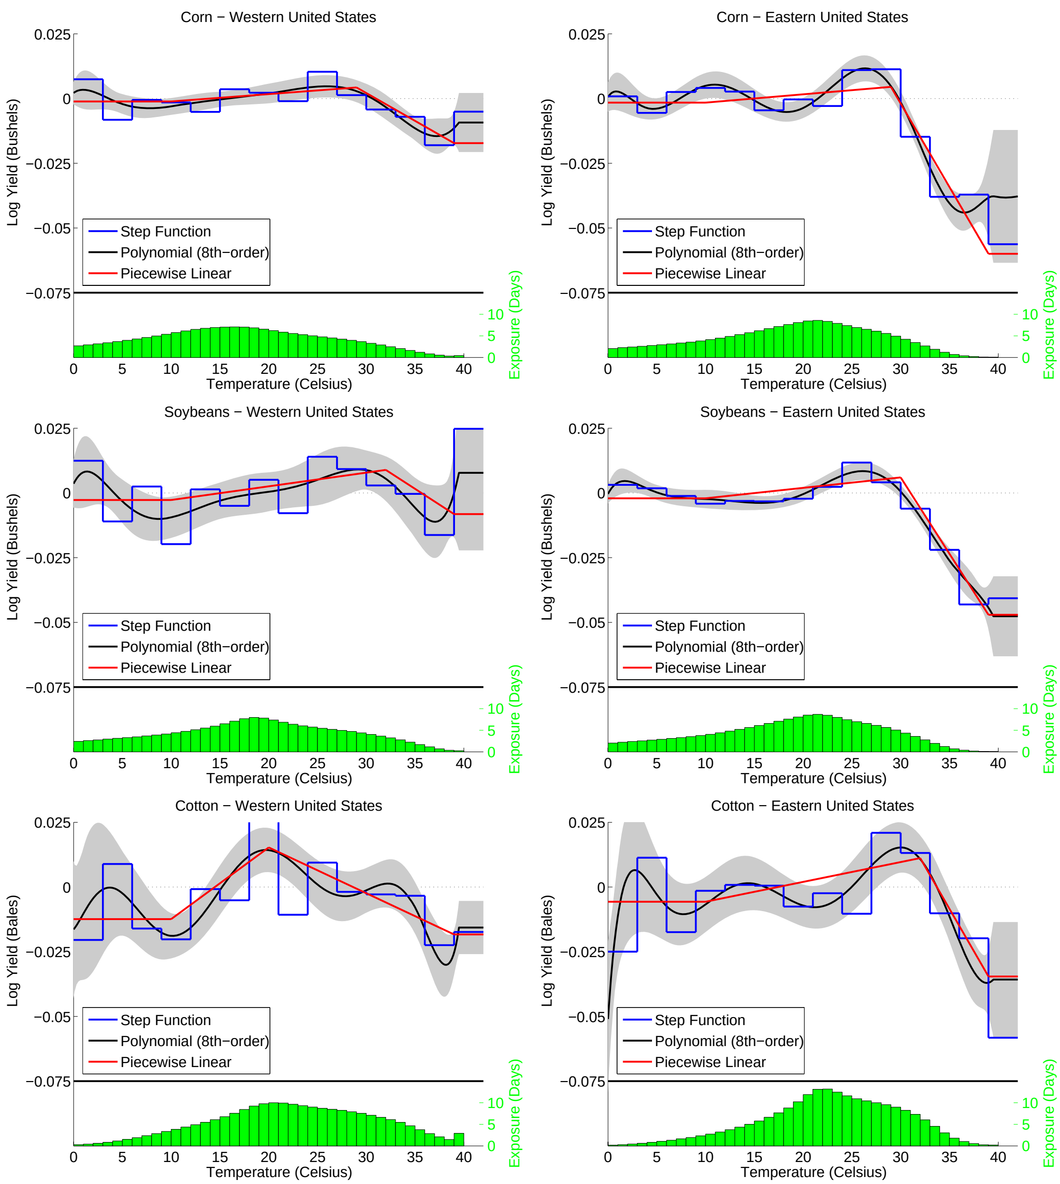

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. The three rows use corn, soybeans, and cotton yields. The left column uses counties west of the 100 degree meridian, while the right column uses counties east of the 100 degree meridian. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

8 Temporal Subsets

In this section we report regression results when the sample is split into two time periods of equal length, 1950-1977 and 1978-2005. The purpose here is to explore whether relative heat tolerance has changed over time as technological change boosted average yields by a factor of two to three over the 56 year time period, depending on the crop. Although average yields in the more recent panel are substantially greater than those in the earlier period, the *relative* relationships between temperature and yield growth are similar.

Since plants have not become more heat tolerant over time in our sample, we find very little difference in our predicted climate change impacts if the regression coefficients are identified using only the years 1978-2005 rather than the full sample. This comparison is shown in Figure A12. While we do not observe an increase in heat tolerance towards the

Figure A11: Nonlinear Relation Between Temperature and Yields for Various Temporal Subsets

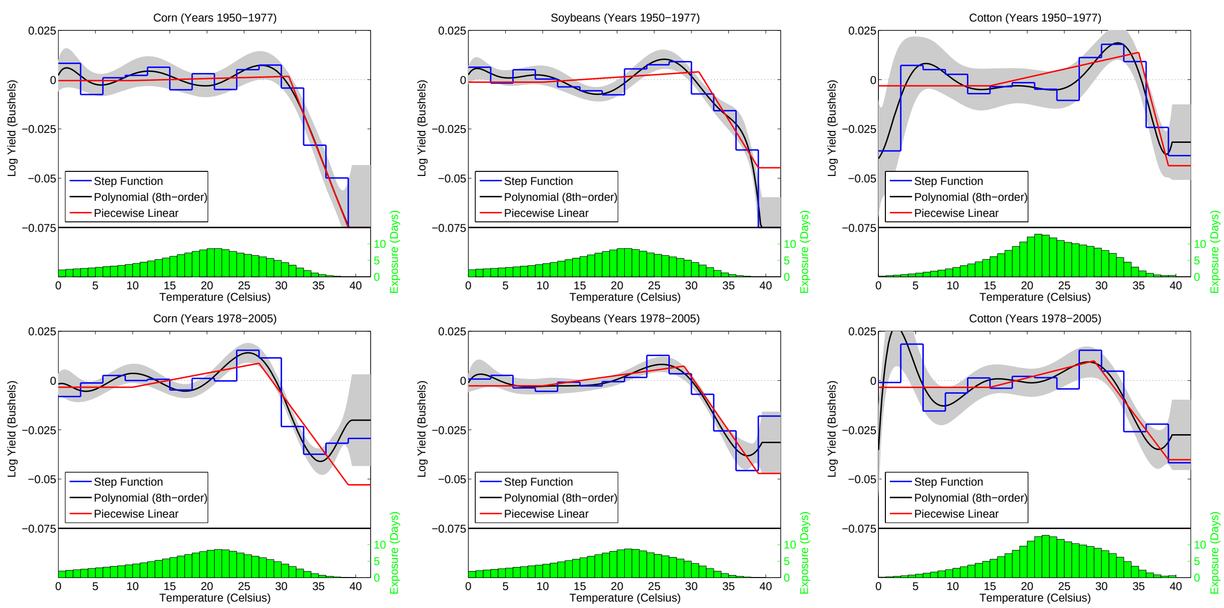

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. The top row uses only the first half of our panel in the estimation of the coefficients and the bottom row uses only the second half. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

end of our sample, a NSF funded study just completed a draft sequence of the corn genome, which might make it easier to develop new corn varieties with greater heat tolerance [10].

Figure A12: Predicted Climate Change Impacts under the Hadley III Climate Model for Various Temporal Subsets

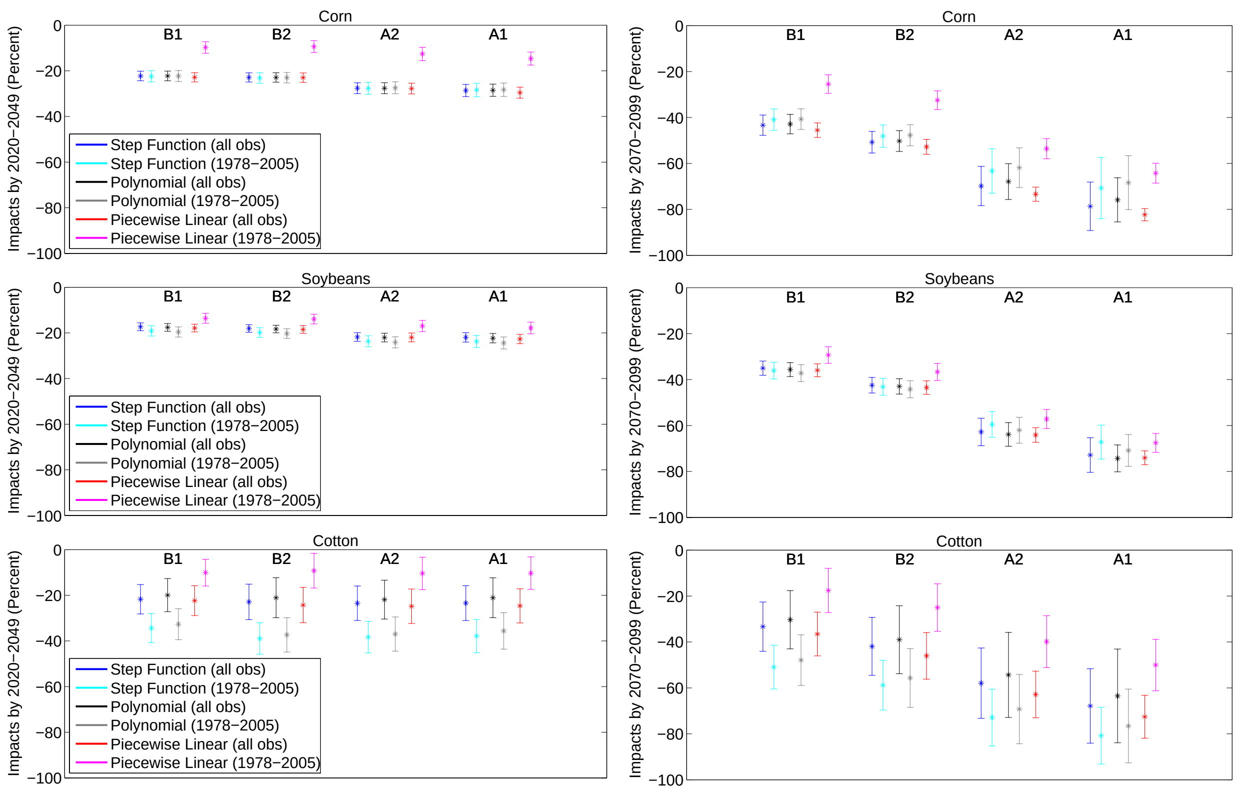

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The darker colors replicate Figure 2 in the main paper. Corresponding lighter colors only use the years 1978-2005 counties in the estimation of the regression coefficients but then evaluate the climate change impacts for the same counties as in Figure 2 of the main paper.

9 Temperature - Precipitation Interaction

All models we have considered so far include a quadratic specification of total growing-season rainfall, but do not consider interactions between temperature and precipitation. Here we report estimates from four regression models that divide the sample by quartiles of total precipitation in June and July. The estimates for corn have a similar shape to that of the

Figure A13: Nonlinear Relation Between Temperature and Yields and Precipitation Interactions

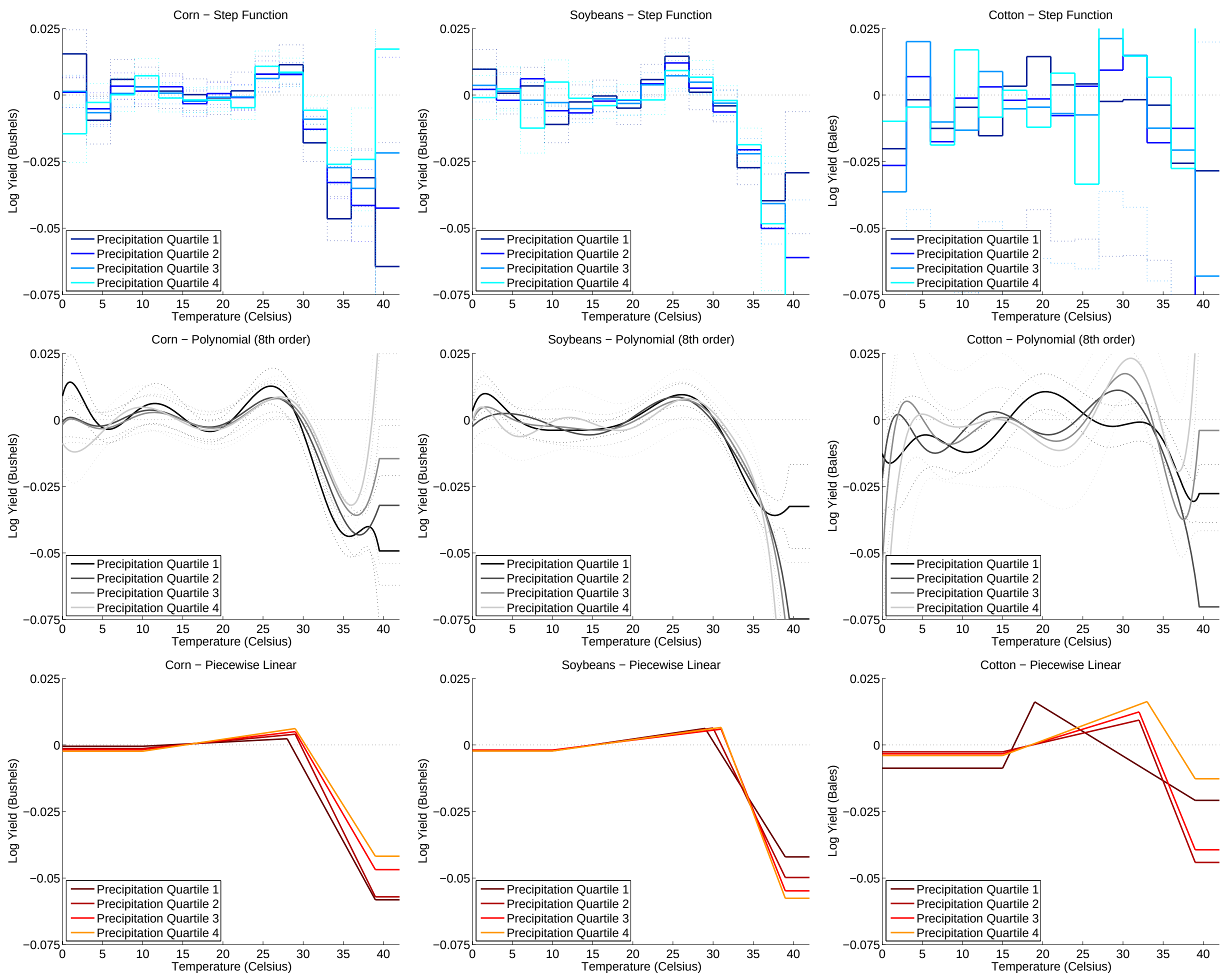

Notes: Graphs display changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The sample is split into four quartiles based on total precipitation in June and July. The 95% confidence band, after adjusting for spatial correlation, are added as dotted lines for the step function as well as the polynomial. Curves are centered so the exposure-weighted impact is zero.

pooled sample up to the critical temperature of 29°C. The decline above the threshold, however, is less steep for subsamples with greater precipitation. Note, however, how the confidence band widens for higher temperatures.

The same interaction effect does not hold for soybeans.

There is some evidence that precipitation partly mitigates damages from extreme temperatures, especially for corn. We did estimate models with richer interactions between temperature and rainfall, but these models do not predict yields out-of-sample significantly better than additively separable model reported above. It is possible that the relatively poor predictive power of precipitation in comparison to temperature stems from greater measurement error in the precipitation variable as spatial interpolation is more difficult. Since we do not find a significant correlation between temperatures and rainfall in the raw daily data, omitting temperature-rainfall interactions should not bias our predictions. That is, they should give the unbiased estimates of the average effects of temperature and rainfall.

10 Growing Season and Time Separability

An important assumption of the empirical model is the additive separability of temperature effects over the growing season. We fix the growing season to the months March through August for corn and soybeans in the main paper, even though northern regions tend to plant later than southern regions, and planting dates may vary from year to year depending on weather conditions. We explore the sensitivity of the results to various definitions of the growing season in Figure A14-Figure A16.

Each figure shows nine alternative specifications of the growing season together with the baseline (top left for corn and soybeans, and center middle for cotton). The estimated temperature effects appear similar regardless of how we shift the growing season. Columns in the first two rows vary the start date (March in the first column, April in the second, and May in the third) while rows vary the end date (for corn and soybeans: August in the top row and September in the second row; for cotton: September in the top row and October in the second row). The third row breaks the growing season into three two-month periods and still obtains similar results by only including temperature readings from a two-month period in the regression. This lends support to the assumption of additive separability. Note, however, how the confidence band becomes much larger if we limit the sample to only the first two months or the last two months. In case of the former, the temperature distribution included very few hot temperatures as shown in the green exposure histogram at the bottom.

In case of the latter, the exposure to temperatures below 10°C is limited.

There is some discussion in the agronomic literature that the effect of temperatures can vary over the season, especially during the corn flowering period [11, 12]. While the flowering period varies by region, it mostly falls within the month of July. Figure A17 therefore replicates our baseline models by including separate temperature variables for the months of July, while all other months are pooled.[10] The estimated relationship changes somewhat, but in different directions for corn and soybeans. The F-statistic against the null hypothesis that the 10 additional temperature variables for July are the same as the other months is 5.16 for corn, 6.54 for soybeans, and 4.10 for cotton. The statistics reject the null hypothesis that temperature effects are the same. Despite statistical significance, however, the model fit improves very little for both crops and the overall shape is still comparable. The R-square of the corn model without any weather variables (just fixed effects and time trends) is 0.6624. Adding the 14 temperature variables and two precipitation variables in our baseline model (step function) increases the R-square to 0.7654. When we add 10 temperature variables for July (omitting the bottom four categories due to lacking data), the R-square increases to just 0.7698.[11] If we include a separate set of temperature variables for each of the six months, adding an additional 70 temperature variables over our baseline model, the R-square again only increases slightly to 0.7774. More importantly though, the predicted climate change impacts are *not* statistically significantly different as shown in Figure A18.

[10] We would like to thank a referee for suggesting this sensitivity check. Since temperatures hardly ever drop below 10°C in July, we exclude these categories for July.

[11] The analogous numbers for soybeans are 0.4792 (no weather), 0.6253 (baseline model), 0.6322 (separate July variables added); and for cotton they are 0.2954 (no weather), 0.3728 (baseline model), 0.3866 (separate July variables added).

Figure A14: Nonlinear Relation Between Temperature and Yields for Various Growing Seasons for Corn

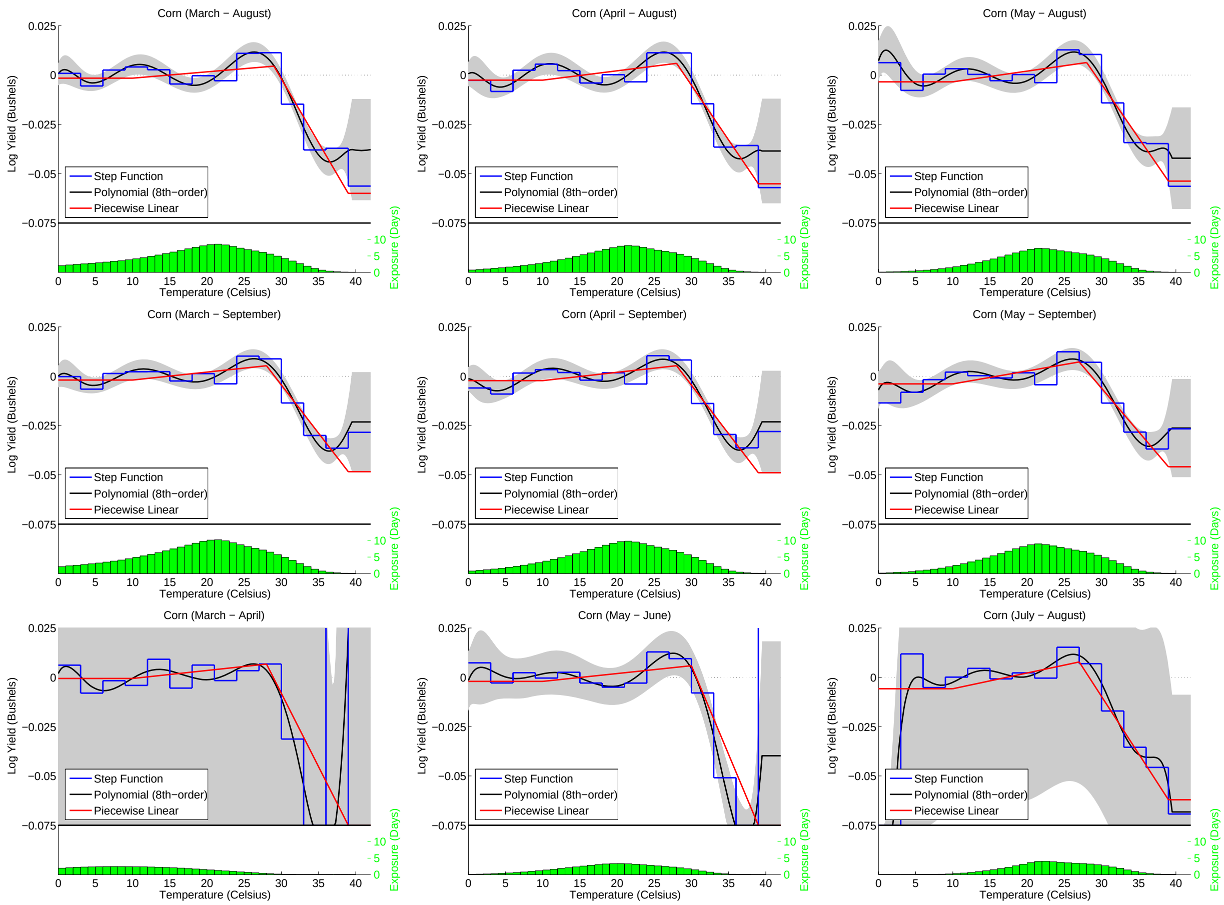

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. Each plot uses a different definition of the growing season. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A15: Nonlinear Relation Between Temperature and Yields for Various Growing Seasons for Soybeans

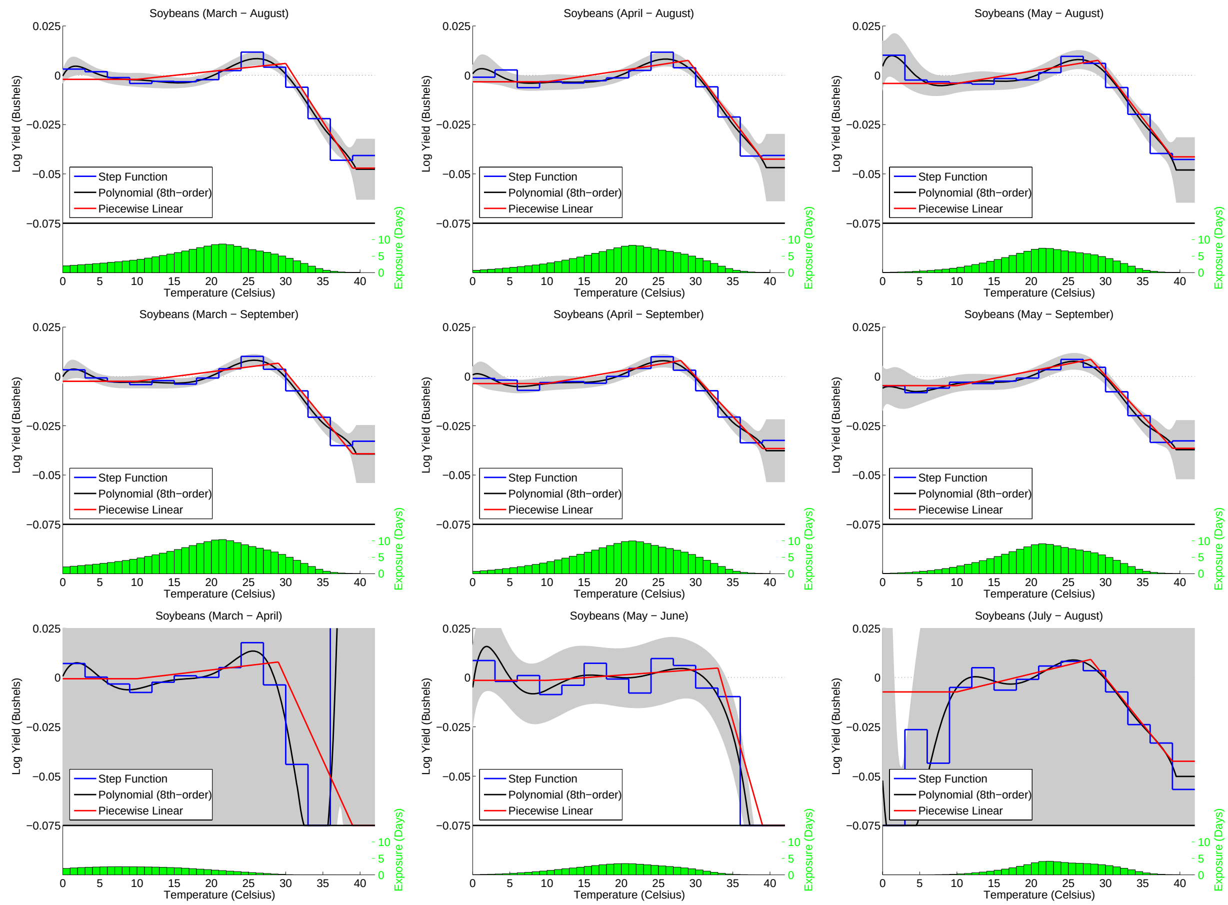

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. Each plot uses a different definition of the growing season. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A16: Nonlinear Relation Between Temperature and Yields for Various Growing Seasons for Cotton

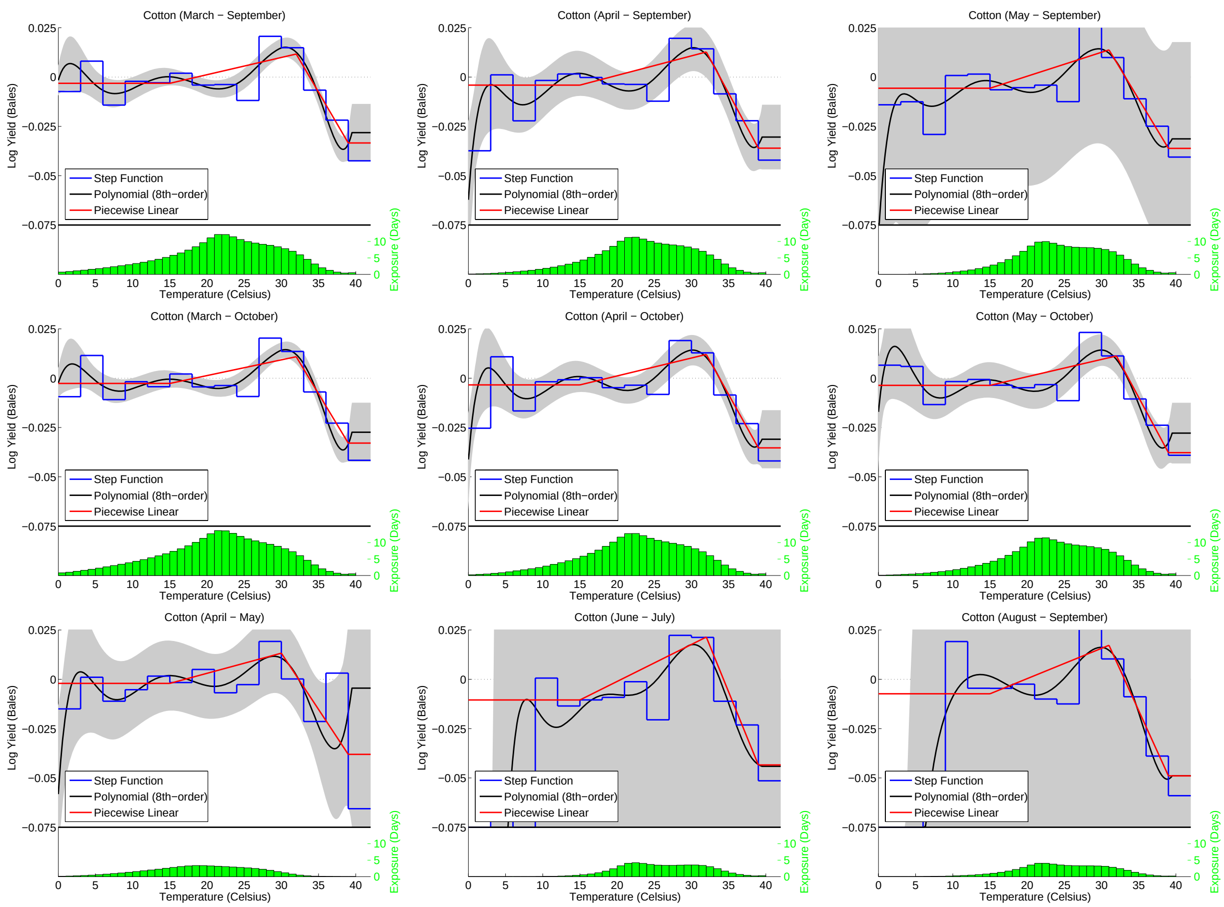

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. Each plot uses a different definition of the growing season. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A17: Nonlinear Relation Between Temperature and Yields in July Compared to Remaining Months

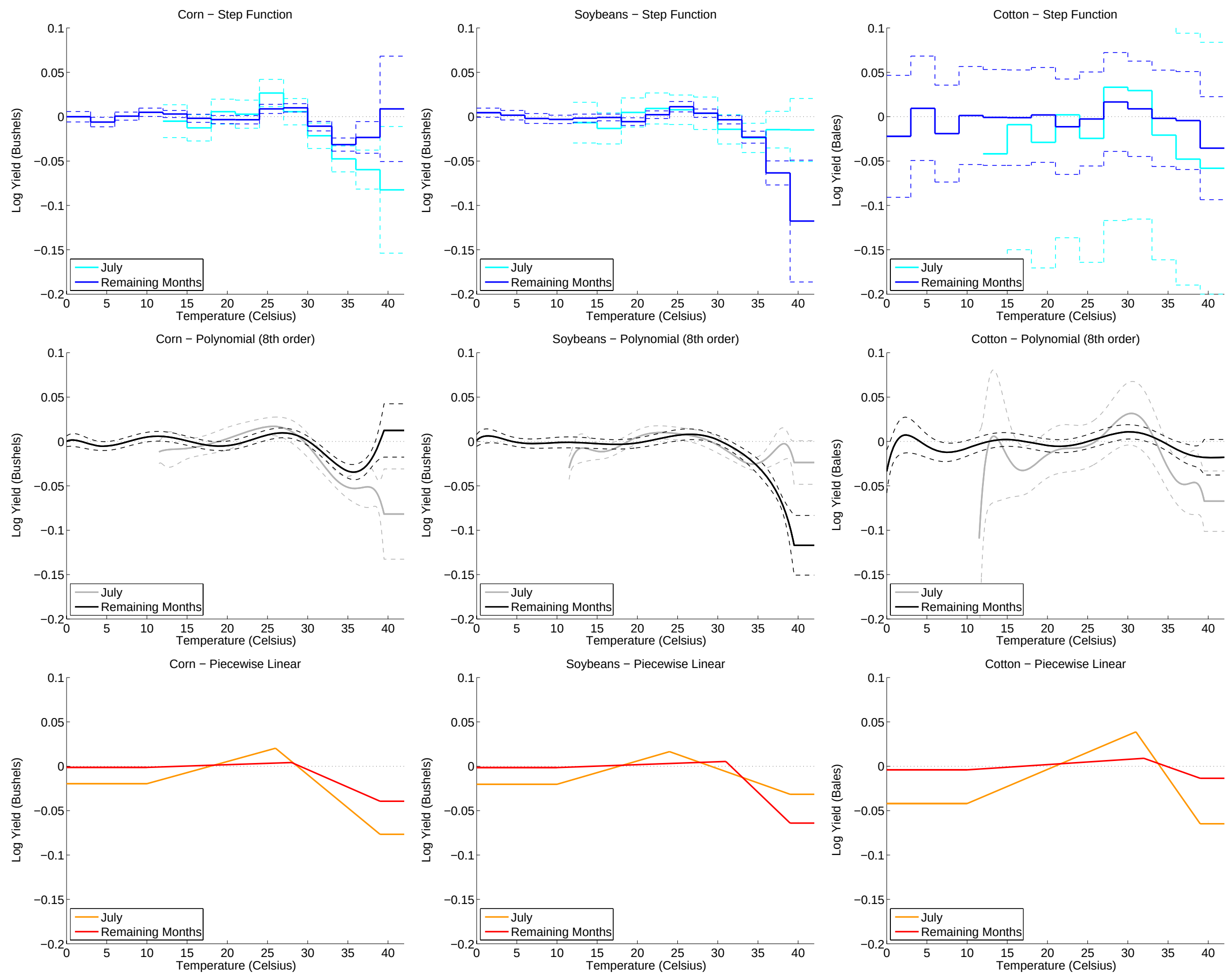

Notes: Graphs show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The growing season is split into July and the sum of the remaining months. The 95% confidence band, after adjusting for spatial correlation are added as dotted lines for the step function as well as the polynomial. Curves are centered so the exposure-weighted impact is zero.

Figure A18: Predicted Climate Change Impacts under the Hadley III Climate Model using Separate Temperature Variables for July

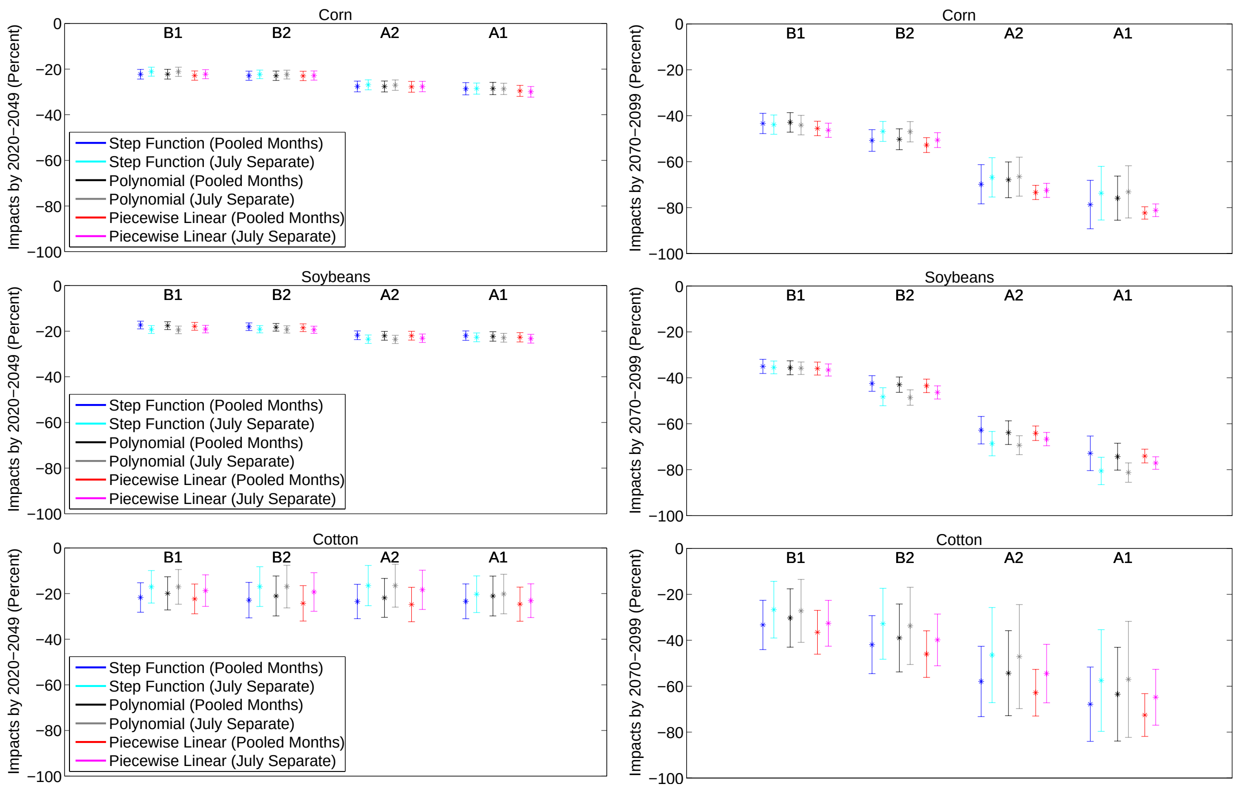

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The dark line replicate Figure 2 in the main paper. Lighter colors include a separate set of temperature variables for the month of July, while all remaining months are pooled.

11 Planted versus Harvested Area

The USDA yield data reports harvested bushels divided by harvested acres. This presents a potential selection problem during extremely bad years when yields are close to zero and farmers choose not to harvest. The yield measure may therefore underestimate the harmful effects of heat waves. In Figure A19 we consider an alternative yield measure that divides the total production quantity for corn by acres *planted* instead of acres harvested. If we use the alternative yield measure we obtain similar results. The corresponding climate change impacts are insensitive to the chosen yield definition as shown in Figure A20. Because the data on planted acres is missing for 20% of our data for corn, we use production per harvested acre in the main paper.

Figure A19: Nonlinear Relation Between Temperature and Yields using Production per Area Planted

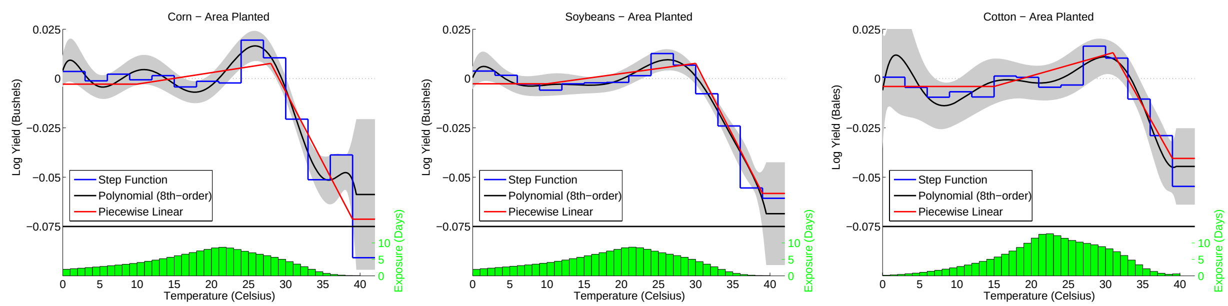

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular 1°C temperature interval where we sum the fraction of a day temperatures fall within each interval. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial regression. Yield is total production divided by the area *planted*. Curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A20: Predicted Climate Change Impacts under the Hadley III Climate Model using Production per Area Planted

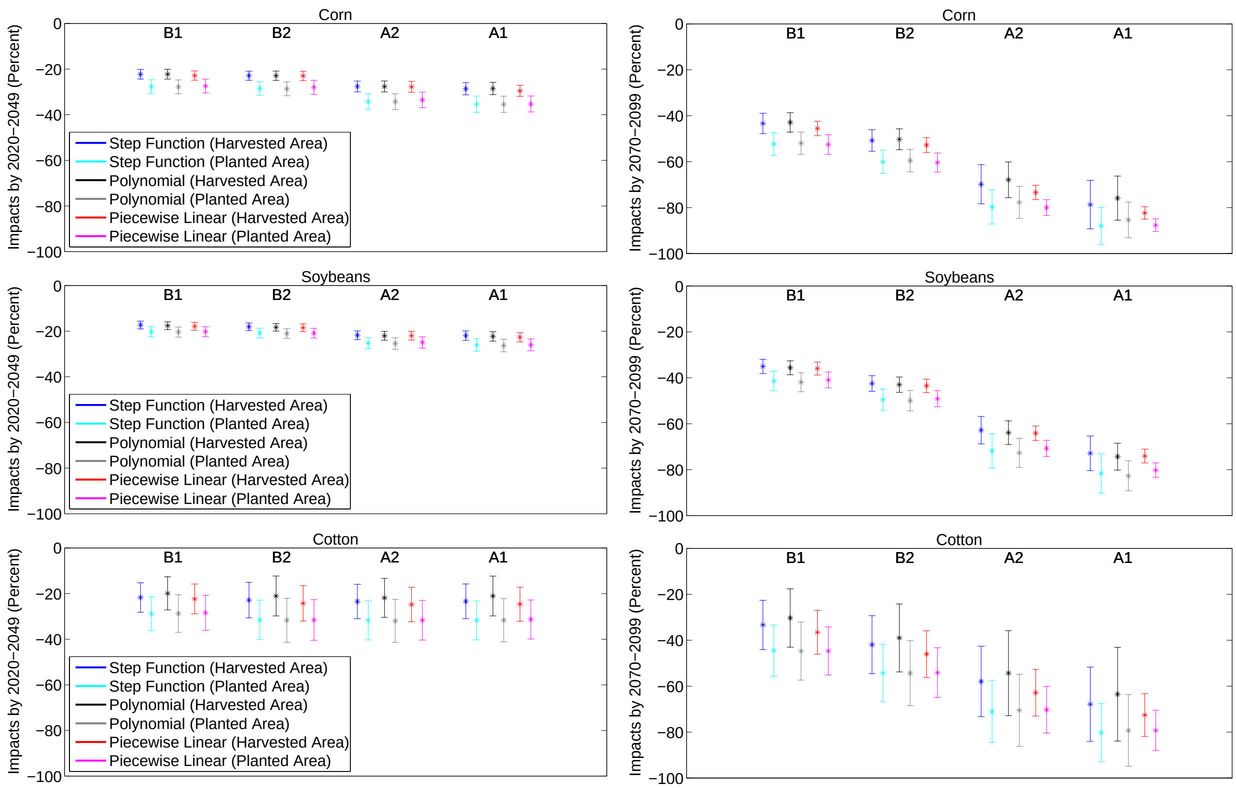

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The dark line replicates Figure 2 in the main paper where yield is total production divided by the area harvested. Lighter colors instead use total production divided by area *planted*.

12 Year Fixed Effects

All regression so far included quadratic time trends by state. A smooth trend will not be able to pick up sudden discrete jumps, i.e., after the introduction of a new crop variety with a significant yield boost or other temporal shocks. We therefore replicate the analysis with year fixed effects instead. Results are shown in Figure A21, which are similar to our initial estimates. The corresponding climate change impacts are generally insensitive to the chosen interpolation method as shown in Figure A22, but they do diverge somewhat for smaller temperature increases for soybeans.

Figure A21: Nonlinear Relation Between Temperature and Yields using Year Fixed Effects

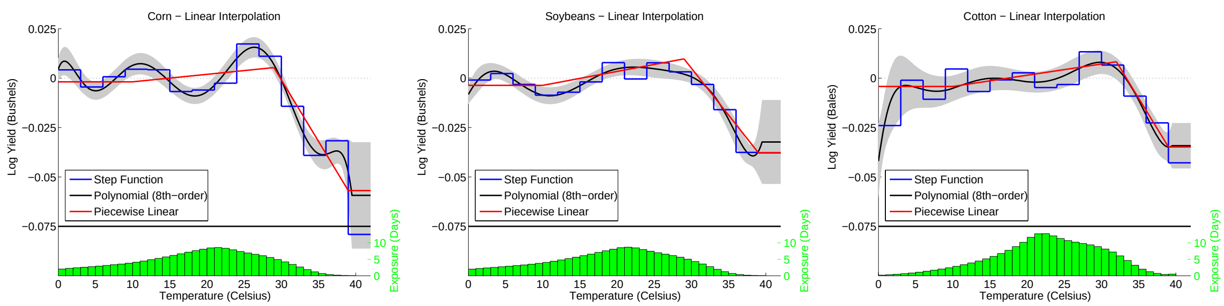

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular temperature. Regressions include year fixed effects instead of quadratic time trends by state. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial Regression. Estimated growth curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A22: Predicted Climate Change Impacts under the Hadley III Climate Model using Year Fixed Effects

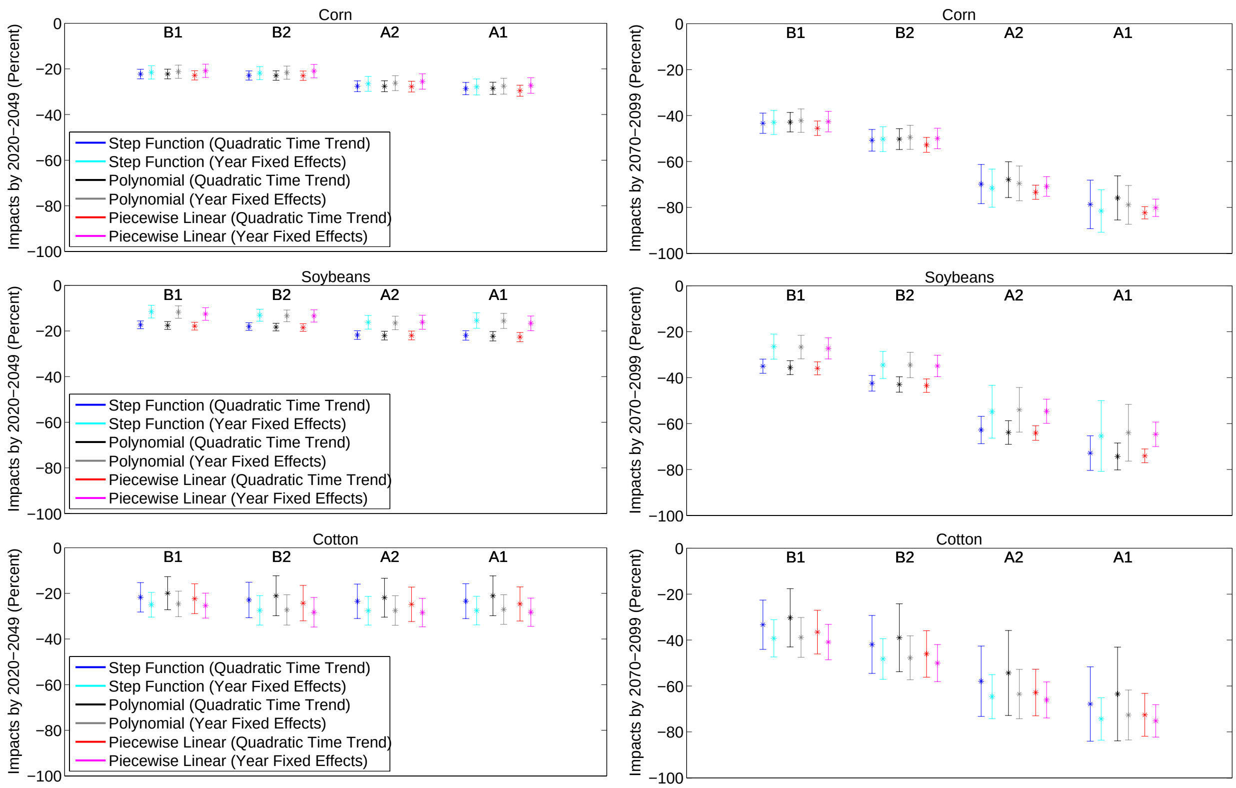

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The dark line replicate Figure 2 in the main paper where we use quadratic time trends by state. Lighter colors instead use year fixed effects in the regression equation.

13 Linear Within-Day Temperature Interpolation

In this section we test the sensitivity of results to the chosen sinusoidal interpolation between minimum and maximum temperature within each day. We replicate the analysis using a linear interpolation instead. Results are shown in Figure A23, which are similar to our initial estimates. The corresponding climate change impacts are insensitive to the chosen interpolation method as shown in Figure A24.

Figure A23: Nonlinear Relation Between Temperature and Yields using a Linear Interpolation between Minimum and Maximum Temperature within Each Day

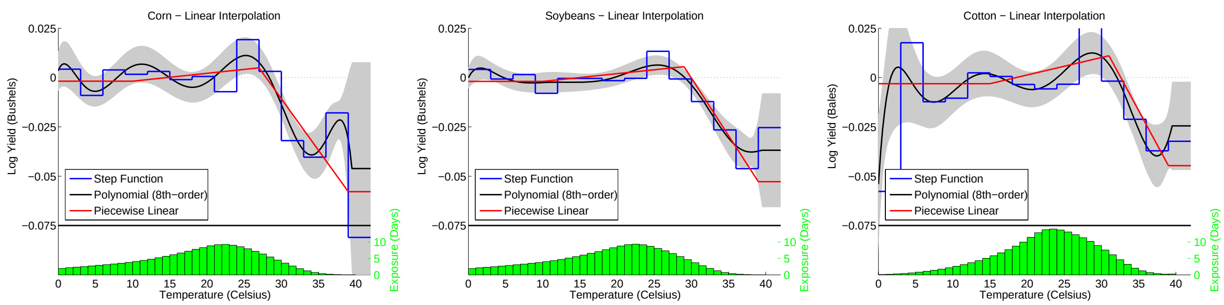

Notes: Graphs at the top of each panel show changes in log yield if the crop is exposed for one day to a particular temperature. The distribution of temperatures within a day is derived by using a linear interpolation between minimum and maximum temperature rather than a sinusoidal interpolation. The 95% confidence band, after adjusting for spatial correlation, is added as grey area for the Polynomial Regression. Estimated growth curves are centered so the exposure-weighted impact is zero. Histograms at the bottom of each panel display the average temperature exposure among all counties in the data.

Figure A24: Predicted Climate Change Impacts under the Hadley III Climate Model using a Linear Interpolation between Minimum and Maximum Temperature within Each Day

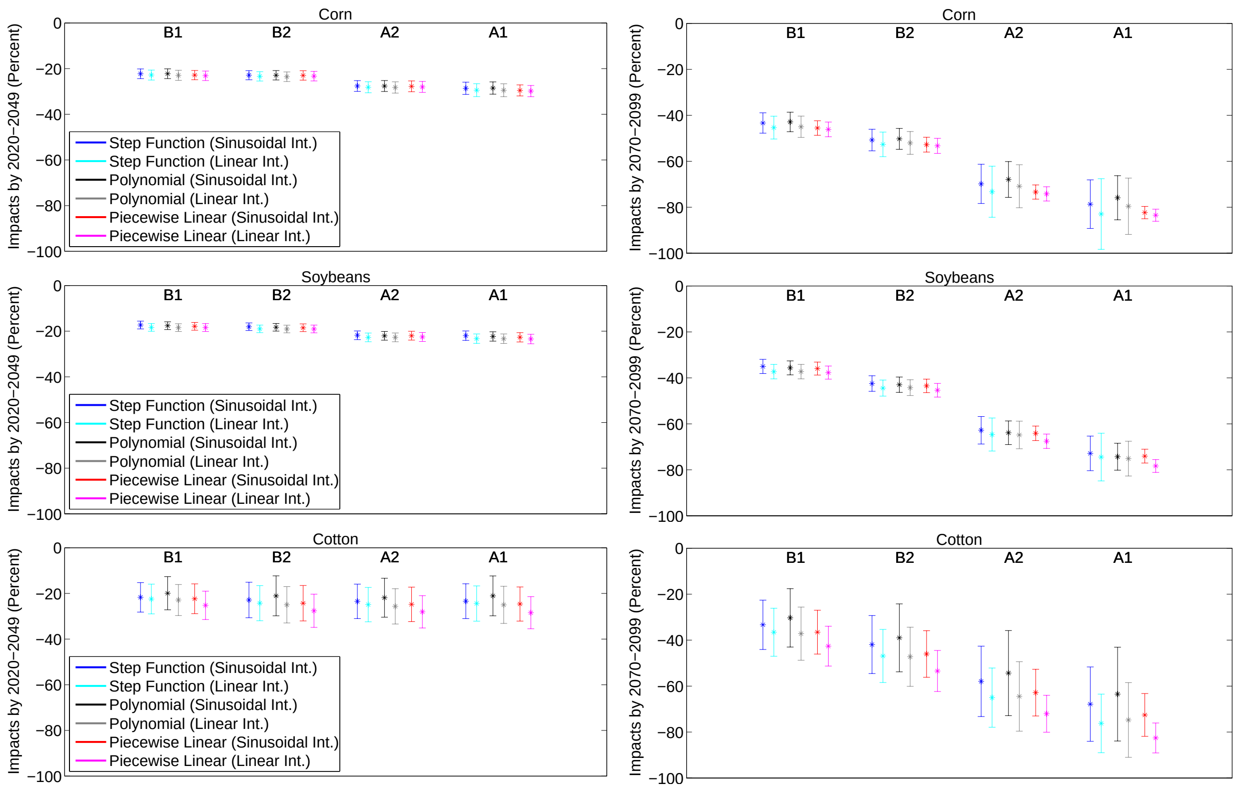

Notes: Graphs show predicted percentage changes in crop yields under four emissions scenarios. Left panels display predicted impacts in the medium-term (2020-2049), right panels in the long-term (2070-2099). A star indicates the point estimates and whiskers show the 95% confidence interval after adjusting for spatial correlation. The dark line replicate Figure 2 in the main paper where we used a sinusoidal approximation between minimum and maximum temperature. Lighter colors instead use a linear interpolation between minimum and maximum temperature.

14 Uniform Climate Change Scenarios

Table A5 reports predicted yield impacts under a range of uniform temperature changes rather than climate change scenarios from the Hadley III climate model. Impacts on total production are based on the most flexible functional form (step-function) and correspond to the blue line of the non-uniform climate change scenarios in Figure 2 of the main paper.

Table A5: Predicted Climate Change Impacts under Uniform Climate Change

	Corn		**Soybeans**		**Cotton**	
	Impact	**s.e.**	**Impact**	**s.e.**	**Impact**	**s.e.**
Temperature +1°C	-6.38	(0.53)	-3.17	(0.46)	-2.59	(1.57)
Temperature +2°C	-14.87	(1.02)	-8.89	(0.90)	-8.34	(3.08)
Temperature +3°C	-24.76	(1.47)	-16.98	(1.25)	-16.30	(4.54)
Temperature +4°C	-35.30	(1.84)	-26.88	(1.51)	-25.48	(6.01)
Temperature +5°C	-45.75	(2.21)	-37.66	(1.73)	-35.61	(7.36)
Temperature +6°C	-55.88	(2.66)	-48.82	(1.97)	-46.72	(8.39)
Temperature +7°C	-65.19	(3.12)	-59.16	(2.24)	-58.28	(8.72)
Temperature +8°C	-73.42	(3.56)	-68.18	(2.53)	-69.08	(8.33)
Temperature +9°C	-80.26	(3.87)	-75.66	(2.80)	-78.08	(7.52)
Temperature +10°C	-85.86	(4.07)	-81.69	(2.99)	-85.06	(6.43)
Precipitation -50%	-10.44	(1.47)	-11.80	(1.29)	+6.38	(2.23)
Precipitation -40%	-7.43	(1.12)	-8.52	(0.98)	+5.13	(1.64)
Precipitation -30%	-4.85	(0.79)	-5.66	(0.69)	+3.86	(1.13)
Precipitation -20%	-2.73	(0.49)	-3.26	(0.43)	+2.58	(0.69)
Precipitation -10%	-1.10	(0.23)	-1.36	(0.20)	+1.30	(0.32)
Precipitation +10%	+0.57	(0.20)	+0.81	(0.18)	-1.30	(0.28)
Precipitation +20%	+0.60	(0.38)	+1.07	(0.35)	-2.61	(0.53)
Precipitation +30%	+0.09	(0.55)	+0.76	(0.53)	-3.93	(0.78)
Precipitation +40%	-0.95	(0.73)	-0.10	(0.71)	-5.24	(1.02)
Precipitation +50%	-2.50	(0.92)	-1.49	(0.92)	-6.57	(1.28)

Notes: Table reports predicted yield changes and standard errors (in parentheses) under unform temperature changes and precipitation change. Predictions are based on our baseline (step function) model.

References

[1] Schlenker W, Roberts MJ (2006) Nonlinear Effects of Weather on Corn Yields. *Rev Agr Econ* 28:391-398.

[2] Thom HCS (1966) Normal Degree Days Above Any Base by the Universal Truncation Coefficient. *Mon Weather Rev* 94:461-465.

[3] Snyder RL (1985) Hand Calculating Degree Days. *Agr Forest Meterol* 35:353-358.

[4] Nakicenovic N, Swart R (2000) *Special Report on Emissions Scenarios* (Cambridge University Press, Cambridge).

[5] Conley TG (1999) GMM Estimation with Cross Sectional Dependence. *J Econometrics* 92:1-45.

[6] Ritchie JT, NeSmith DS (1991) in *Modeling Plant and Soil Systems. Agronomy 31*, eds Hanks J, Richie JT (American Society of Agronomy, Madison), pp 5-29.

[7] Schlenker W, Hanemann WM, Fisher AC (2006) The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions. *Rev Econ Stat* 88:113-125.

[8] Mendelsohn R, Nordhaus WD, Shaw D (1994) The Impact of Global Warming on Agriculture: A Ricardian Analysis. *Am Econ Rev* 84:753-771.

[9] Deschênes O, Greenstone M (2007) The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *Am Econ Rev*, 97:354-385.

[10] National Corn Growers Association, February 27 2008, http://monsanto.mediaroom.com/index.php?s=43&item=576.

[11] Hallauer AR (2000) *Speciality Corns* (CRC Press, Boca Raton).

[12] Thomison P (2002) *Drought and Heat Stress Effects On Corn Yield Potential* (Crop Observation and Recommendation Network, Ohio State University).