Adaptation to Climate Change: Evidence from US Agriculture
Online Appendix

Marshall Burke and Kyle Emerick

A Online Appendix Materials

A.1 Extended model of Bayesian learning

We simulate the extended learning environment described in the main text where a farmer learns from both her county and a single neighboring county. We generate a simulated temperature series of 100 years for both counties, where the true mean temperature in both counties is 24, the variance of annual temperature is 6, the farmer's prior on average temperature in both counties is 23.5, and both priors have a variance of 3. We vary both the correlations between annual temperatures and beliefs about mean temperatures and draw each weather series from a multivariate normal distribution. We repeat this exercise 25 times and calculate the average belief about the farmer's own county for each time period.

Figure A.1 shows the farmer's belief about average temperature in their own county over time. Observing temperature realizations from the neighboring county hastens the convergence of the belief to the true mean the most when average temperatures are highly correlated and annual realizations of temperature are uncorrelated. Nevertheless, at least for the range of parameter values explored in this simulation (which we believe spans the range of parameter values in the data), the added benefit of observing nearby counties appears modest.

Figure A.1: *Simulated belief's about mean temperature. Each plot shows two sets of beliefs. The blue line corresponds to beliefs when the farmer only observes her own county. The black line corresponds to beliefs when the farmer observes both her own county and the neighboring county. The true average temperature in the farmer's county is 24 (the horizontal dashed line). Rho refers to the correlation between annual temperature in the two counties and v refers to the correlation between average temperatures.*

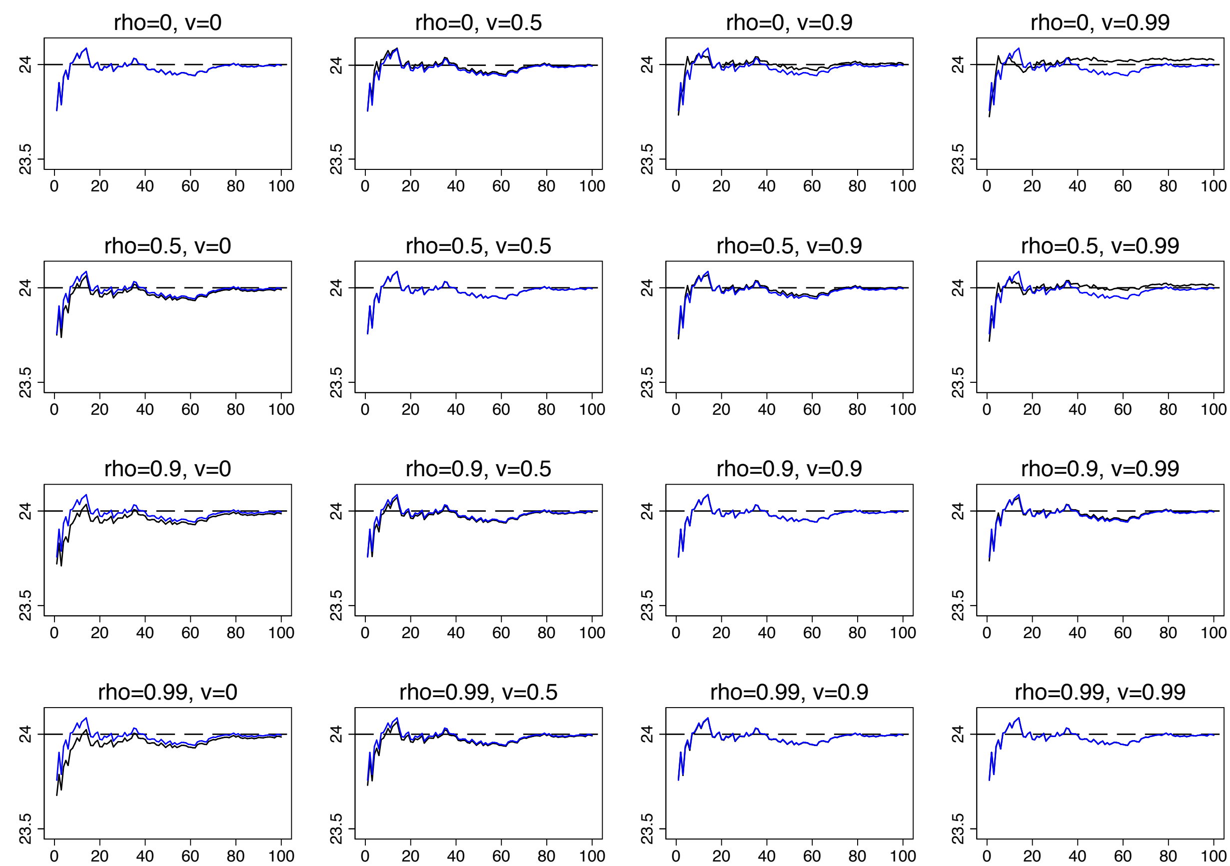

A.2 Understanding changes in climate and agriculture over time

Figure A.2 plots changes in GDD 0-29C and GDD >29C between 1980-2000 for our sample counties. The left plot shows that while increases in "beneficial" and "harmful" GDD are positively correlated, many counties experienced increases in one and decreases in the other. The right panel plots the relationship between change in log corn yields and change in harmful GDD >29C over the same period. Because both figures show large outliers in terms of either temperature or log yields (the $\sim$ 10 points plotted as white circles in the figure), and we run regressions with and without these outliers to make sure they are not driving our results. Figure A.3 maps these changes in GDD, showing that extreme-heat outliers are clustered among a few counties in southern Texas.

Figure A.2: *Changes in GDD and corn yield for corn-growing counties east of the 100th meridian. Left panel: changes in GDD 0-29C and GDD>29C over the 1980-2000 period. Middle panel: change in log corn yields and GDD>29C over the same period. Right panel: changes in GDD>29C, 1955-75 versus 1980-2000. To check robustness we run the long differences regressions with and without the points shown as white circles.*

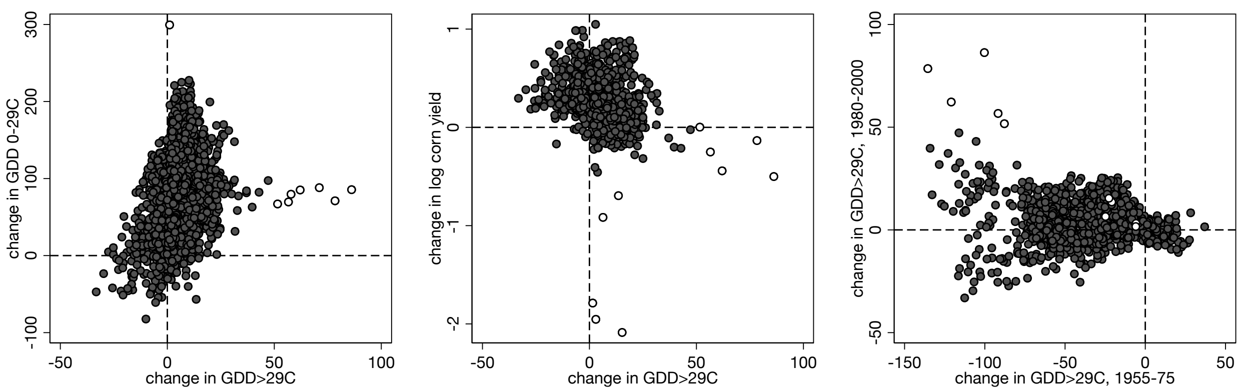

There are two potential concerns with the variation in temperature we are using in the long differences. The first is that state fixed effects could absorb most of the meaningful variation in temperature changes over time, and the second is that the apparent long-run changes in temperature might just reflect short-run variation around endpoint years - e.g. single hot or cold years that create large differences between endpoints but do not reflect underlying long-term changes in temperature. If this latter concern were true, then the panel and long difference approaches will mechanically deliver estimates of yield responses that are similar to each other (albeit with the LD being much noisier), which in turn would lead us to erroneously conclude that there had been "no adaption" when in fact there was no underlying trend to adapt to.

To address these concerns, we begin by more carefully characterizing the variation in extreme heat in the long differences and the panel, and comparing this variation to the

Figure A.3: *Map of changes in GDD 0-29C and GDD above 29C between 1980-2000, for corn-growing counties east of the 100th meridian. Rightmost panel re-plots the change in GDD >29 dropping the outliers indicated in Figure A.2.*

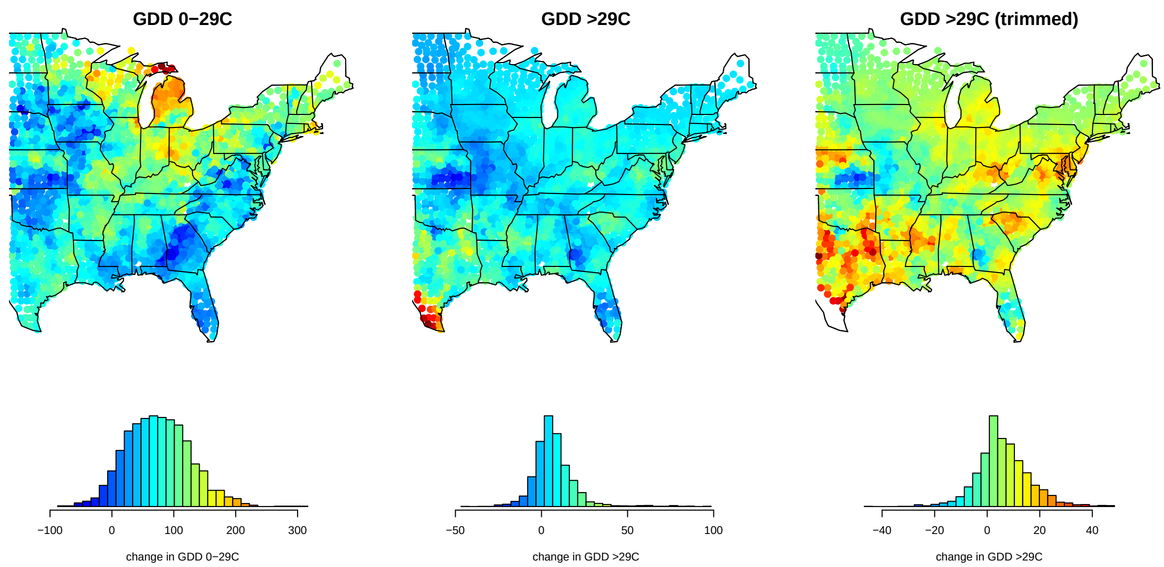

projected future changes in extreme heat. Table A.1 shows the amount of variation left in our extreme heat variable after accounting for the other climate variables and various sets of fixed effects, in both the panel and the long differences. The variation in long-run changes in extreme heat is smaller than the inter-annual variation in extreme heat, but the other climate controls and fixed effects absorb a smaller percentage of the variation in the long differences (as shown by the 3rd and 4th columns in the table) than in the panel.

Figure A.4 relates the distribution of observed changes in GDD>29 to the extent of variation in projected changes in GDD>29, plotting the raw distributions (left panel) and the residualized distributions (right panel). The conditional distribution of future changes is calculated for each of the 18 climate models as the residuals from a regression of GDD>29C on GDD0-29C and a piecewise function of precipitation (i.e. the other climate variables in all of our regressions). Given our empirical approach, we are most interested in the overlap in the conditional distributions, and the Figure demonstrates the substantial overlap between the variation we are exploiting in the long differences, the variation we exploit in the panel, and the variation we use in the projections (after accounting for projected changes in the other climate variables). This gives us additional confidence that our projections are not wild extrapolations from historical experience.

To address concerns that the "changes" in temperature we observe over time are indeed meaningful and not a function of short-run variation around endpoint years, we first estimate the trend in temperature and precipitation from 1978-2002 for each county in our main sample by running the regression

$$ln(clim_t) = \alpha + \beta t + \varepsilon_t, \tag{1}$$

Figure A.4: *Distributions of GDD > 29 for the 1980-2000 period (red lines) and as projected for 2050 across 18 climate models for the A1B scenario (blue lines). Left panel: raw changes. Right panel: changes conditional on other climate variables (GDD0-29C, precipitation) and on various sets of fixed effects as indicated. Distributions are area weighted, as in our main regressions.*

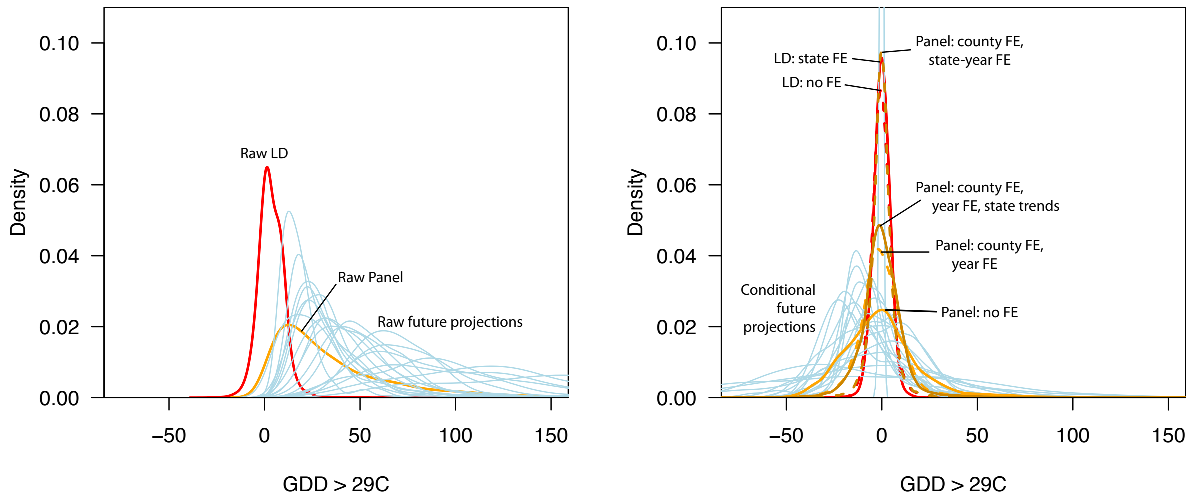

where t is the sample year. Results for our main GDD>29C variable are shown in Figure A.5 for the main corn belt states. Plots represent the distribution of annual percentage changes in GDD>29C across counties within a given state (i.e. the kernel density of βs estimated in Equation (1)), and show that annual changes in extreme heat vary by 2-4 percentage points within states. This represents substantial variation over our 20 year estimation period. For instance, estimates for Iowa suggest that changes over 20 years ranged from 80% declines in exposure to extreme heat to slight increases in exposure; estimates for Illinois range from 40% decreases to 70% increases.

Second, we show in a simulation that the observed distribution of temperature changes over our study period is highly unlikely to be generated by a time series with a fixed mean. For each county in our data, we calculate the observed mean μ_i and standard deviation σ_i of GDD>29C between 1978-2002, and then use these parameters to generate 1000 simulated panel datasets, where the observation for $GDD > 29_{it}$ is a draw from a normal distribution $\sim N(\mu_i, \sigma_i)$. For each of these simulated panels we then compute long differences for each county (differencing the 5-year averages at the endpoints each county's time series, as in our main exercise). These long differences are therefore generated from data with no "permanent" change in temperature, with variation in the LD by construction coming only from random variation in temperature around the endpoints. We can then compare the distribution of these simulated changes to our actual observed distribution of ΔGDD>29 to understand whether the changes we observed were likely generated from data with no "permanent" change in temperature.

The results are shown in Figure A.6, with the observed distributions of ΔGDD>29 shown in red and the 1000 simulated distributions shown in grey (the right panel is for 1980-2000, the left panel repeats the exercise for 1955-1975 with corresponding data). This exercise suggests the observed changes over time are extremely unlikely to be generated from data with a fixed mean. The distribution of observed changes over 1955-1975, a period of substantial cooling in the central US, is far to the left of all of the simulated distributions for that period; the observed distribution in 1980-2000, a period of substantial average warming across the US, is shifted substantially to the right of the simulated distributions for that period.

Figure A.5: *Distribution of estimated annual growth in GDD > 29 for counties in 13 corn belt states. Horizontal axis for each plot is the estimated annual growth (% per year) in GDD > 29 for 1978-2002. Vertical axis is kernel density.*

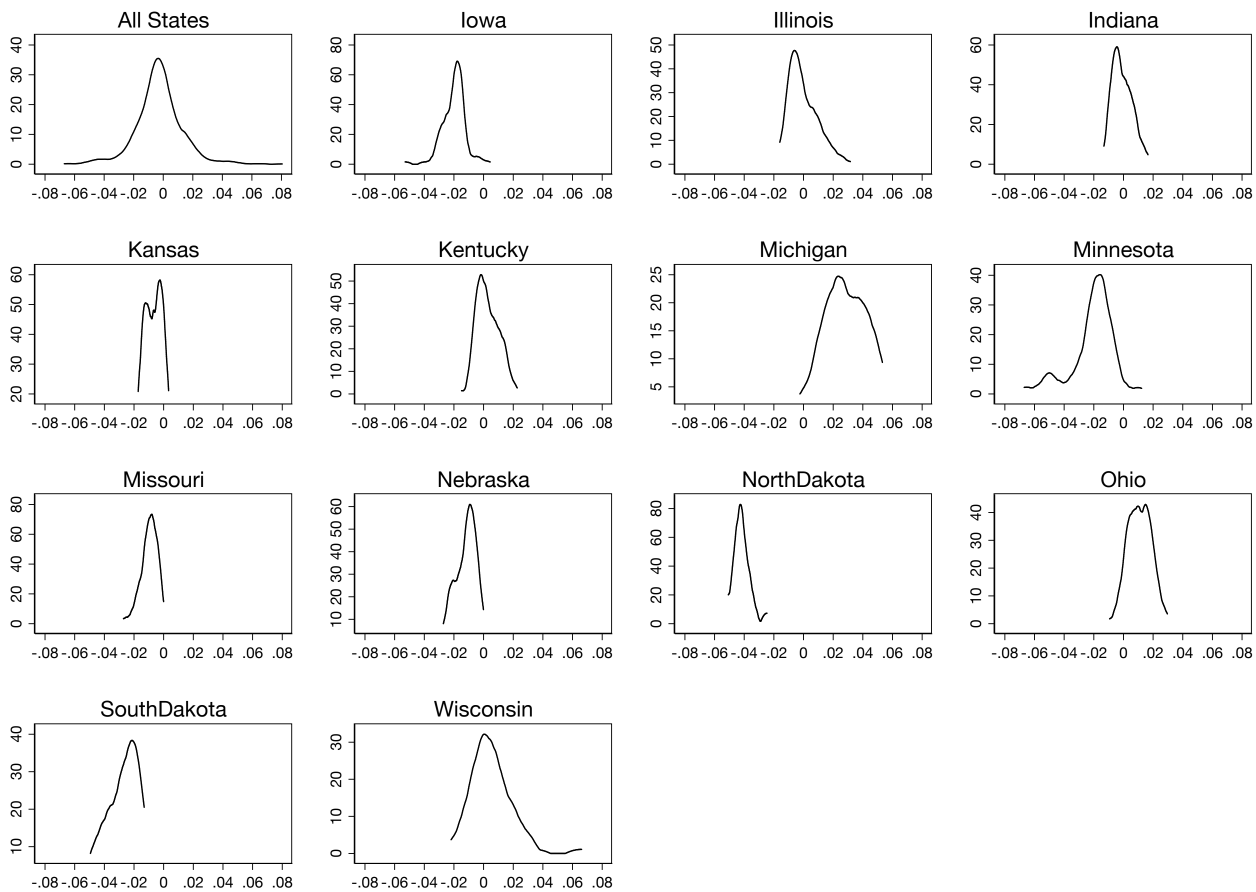

Table A.1: *Yield response to GDD>29 across different panel and long difference models, and variation in GDD>29 after accounting for fixed effects and other climate controls in these models.*

	DV: log yield		DV: GDD>29			
	$\hat{\beta}$	$\hat{se}$	R^2	σ	share>5	share>10
Panel						
None	-0.0036	0.0005	0.00	62.83	0.53	0.50
Climate	-0.0057	0.0010	0.70	33.11	0.42	0.35
Climate, state FE	-0.0067	0.0009	0.80	27.10	0.34	0.26
Climate, county FE, year FE	-0.0056	0.0007	0.91	20.79	0.32	0.21
Climate, county FE, state trends	-0.0055	0.0007	0.91	20.66	0.32	0.21
Climate, county FE, state-year FE	-0.0062	0.0007	0.97	9.78	0.22	0.10
Long Differences						
None	-0.0077	0.0017	0.00	9.44	0.31	0.14
Climate	-0.0053	0.0010	0.34	8.23	0.27	0.12
Climate, state FE	-0.0044	0.0008	0.60	6.63	0.17	0.05

The first two columns display the estimated effect of GDD $>$ 29 on log of corn yield and its standard error, estimated using alternate versions of Equation (10) or its panel analog; all estimates are statistically significant at the 1% level. Remaining columns pertain to regressions of GDD $>$ 29 on GDD$<$29, a piecewise linear function of precipitation, and the listed fixed effects, and show the R^2 of that regression, the standard deviation of the GDD>29 residuals, and the percentage of those residuals larger than 5C or 10C. Panel regressions are based on 48,465 county-year observations, and long difference regressions based on 1,531 county observations.

Figure A.6: *The observed distribution of changes in GDD>29 (in red), compared to 1000 distributions generated from data with the same variance but no change in mean (shown in grey). The left panel is real and simulated changes between 1955-1975, the right panel for 1980-2000.*

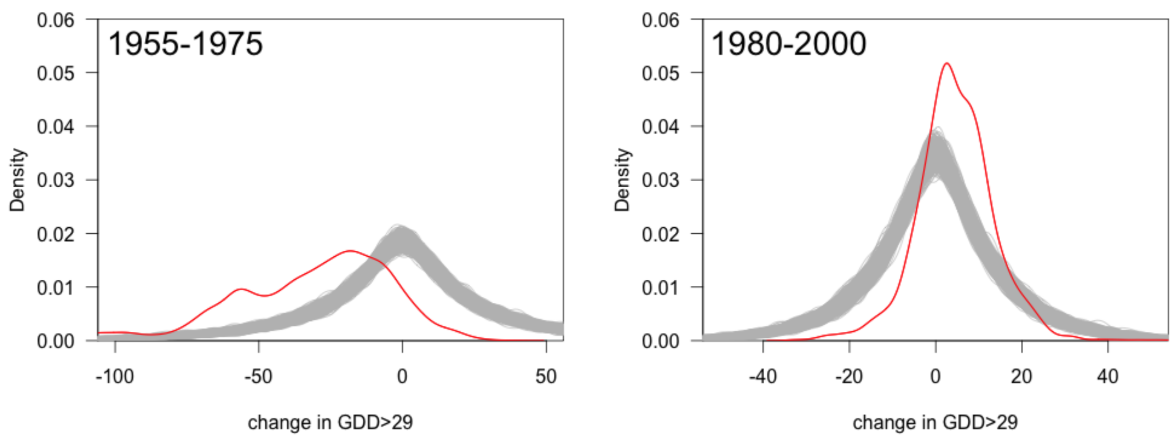

A.2.1 Can the long-differences uncover a long-run response, if it actually exists?

We now explore the conditions under which the panel and long differences estimates actually identify different responses if the "true" short-run and long-run response are actually different. Consider the following data generating process:

$$y_{it} = \alpha_i + \delta_1 \bar{T}_i + \delta_2 \nu_{it} + \epsilon_{it} \tag{2}$$

Change in the outcome in county i in year t responds to both average climate $\bar{T}_i$ as well as to short run variation about that average ν_{it}. The latter can be thought of as being the weather draw the farmer receives after making her planting and input decisions. If $\delta_2 < 0$, the typical "adaptation" assumption is $|\delta_1| < |\delta_2|$, i.e. outcomes respond differently to changes in average temperature and to short-run variation around that average. Our goal is to identify these differential effects.

Let T_{it} represent the temperature in a county-year that is actually observed, and imagine that it's made up of three pieces:

$$T_{it} = T_{i0} + f_i(t) + \nu_{it} \tag{3}$$

where T_{i0} is some baseline average temperature before any warming starts, $f_i(t)$ is a county-specific warming trend that might not be linear, and ν_{it} is mean-zero random year to year weather.

In the "typical" panel model with county FE and county trends, i.e.:

$$y_{it} = \alpha + \beta T_{it} + c_i + \theta_i * t + \varepsilon_{it} \tag{4}$$

if the true trend in temperature is close to linear then the only variation left over in T_{it} is going to be ν_{it}, and so β will be estimated from weather variation alone. Now construct our "long difference" estimate between two periods centered around years a and b. We have that the change in average temperature between these endpoints is:

$$\begin{aligned}\Delta\bar{T}_i &= \bar{T}_{ib} - \bar{T}_{ib} && (5)\\ &= (T_{i0} - T_{i0}) + f_i(b) - f_i(a) + (\bar{\nu_{ib}} - \bar{\nu_{ia}}) && (6)\\ &= f_i(b) - f_i(a) + \Delta\bar{\nu}_i && (7)\end{aligned}$$

So some part of the $\Delta\bar{T}_i$ we observe is coming from the true underlying trend in average temperature and some is coming from random noise from the weather. (If we assume a linear annual temperature trend in each county of θ_i, then this is just $\Delta\bar{T}_i = \theta_i * (b - a) + \Delta\bar{\nu}_i$). So in the long differences regression:

$$\Delta\bar{y}_i = \alpha + \beta_{LD}\Delta\bar{T}_i + \epsilon_i \tag{8}$$

β_{LD} will be estimated from a combination of short-run and longer-run variation.

To explore the consequences of this, we set $\delta_1 = -10$ and $\delta_2 = -20$, fix the variation in trends in average temperature across counties, and slowly increase the variation in year-to-year weather ν_{it}. Results are shown in Figure A.7. As $\Delta\bar{T}_i$ is made up more and more of changes in weather rather than changes in underlying average temperature, then it becomes increasingly difficult to recover the true long-run response and $\hat{\beta_{LD}} \rightarrow \delta_2$ (See Figure A.7).

This gives us a useful prediction we can take to the data. Since $\bar{\nu_{it}} \rightarrow 0$ as more and more years are included in the average, then it is mechanically the case that given an underlying trend in average temperature (e.g. +0.1C per year), the proportion of $\Delta\bar{T}_i$ that is made up of $\Delta\bar{\nu_{it}}$ also goes to zero, because $\Delta\bar{\nu_{it}} \rightarrow 0$. This implies that as we average our endpoints over more and more years, our estimated β_{LD} should converge to the "true" value (δ_1 in this case). This can be easily seen in the same simulation: as shown in Figure A.8, fixing both the magnitude of the underlying temperature trend and the variance in weather at some value, averaging the endpoints over more years causes LD estimates to converge to the "true" value of $\delta_1 = -10$.

We explore the corresponding result as we average endpoints over increasing numbers of years in our data. To maximize the amount of years we can average over without having overlapping periods, we set the center of the two endpoints at 1963 and 1992 (our data are from 1950-2005), and vary the number of years each endpoint is averaged over from 3 to 27. Thus the estimate using the largest amount of data differences the average in years 1979-2005 and the average in years 1950-1976, and the estimate using the least data differences the averages in years 1991-1993 and 1962-1964.

The results are shown in Figure A.9. If the true response to longer-run temperature changes was smaller in absolute value than the response to weather, then estimated coefficients should get smaller as we use more years in the endpoints. If anything we see the opposite – coefficient point estimates get slightly more negative – suggesting that our β_{LD} estimates using shorter endpoints are not biased away from zero and if anything are conservative estimates of the true effect.

Figure A.7: *Simulation results for the behavior of* $\hat{\beta_{LD}}$ *as the percentage of the variation in observed temperature that is due to the weather is decreased. The true response to short-run variation (weather) is -20 and the true response to long-run change is -10.*

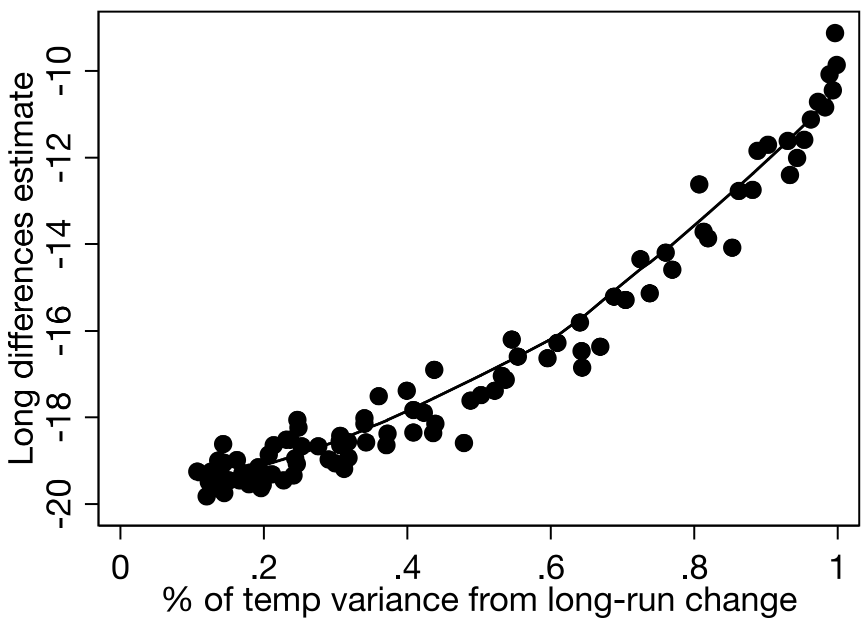

Figure A.8: *Estimates of β_{LD} while fixing both the underlying change in temperature and the variance in the weather, but varying the number of years the long-difference endpoints are calculated from.*

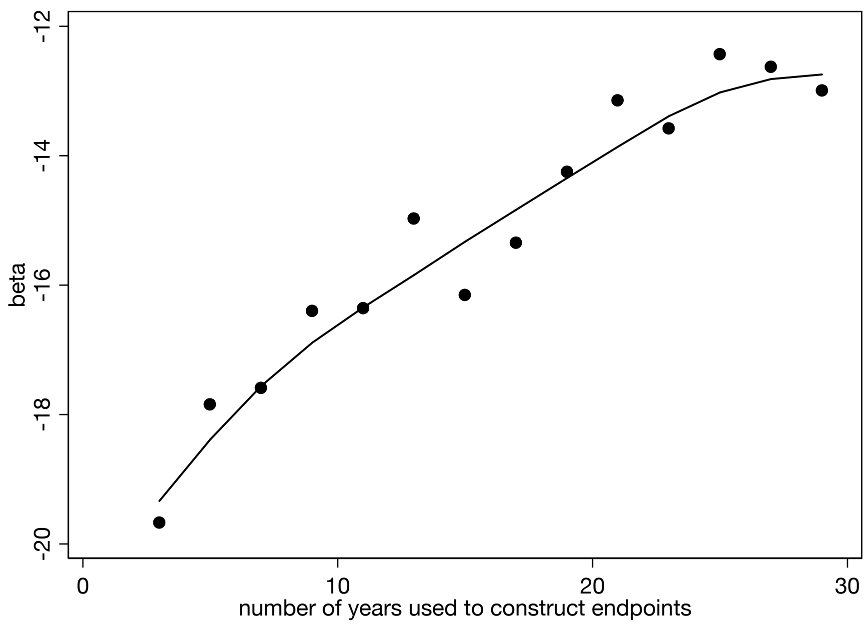

Figure A.9: *Estimates of* $\beta_{GDD>29}$ *from our actual data, as a function of the number of years used to construct the long differenced endpoints. Endpoints are centered at 1963 and 1992 to maximize the amount of data we can use. Vertical bars give 95% CI based on state-level clustering.*

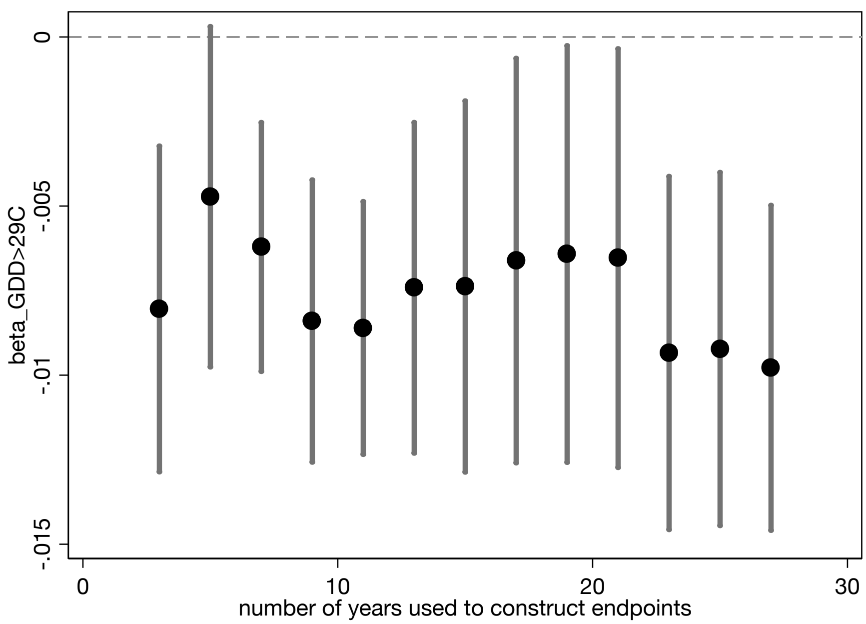

A.3 Correlates of trends in extreme heat

Tables A.2 and A.3 investigate the sensitivity of our main long difference results to controlling for county characteristics. Table A.2 shows that changes from 1980-2000 in GDD>29 are not strongly correlated with baseline measures of population density, various measures of farm land use, and characteristics of the soil. Table A.3 shows regression coefficients when controlling for these variables. Adding controls does not substantially change our results.

Table A.2: Coefficients and p-values of univariate regressions of county characteristics on change in extreme heat exposure

	Coefficient	p-Value
Pop density 1980	1.96	0.16
Farm area 1978	889.29	0.14
Corn area 1980	273.31	0.31
County area	0.67	0.56
Irrigated area 1982	338.62	0.15
Farm value 1978	16.49	0.15
Percent of soil that is clay	-0.04	0.49
Water capacity of soil	-0.01	0.60
Percent of soil that is high quality	0.13	0.20
Income per capita 1978	7.59	0.70

Table displays coefficients and p-values from regressions of each county characteristic on change in GDD> 29 from 1980-2000 and state fixed effects. Standard errors for each regression are clustered at the state level.

A.4 Robustness to outliers

Robustness of our corn yield results to dropping outliers is explored in Table A.4. Point estimates decline slightly when outliers are dropped – not surprising given that nearly all of the outliers experienced both yield declines and large increases in exposure to extreme heat – but coefficients are statistically indistinguishable from estimates on the full sample.

Robustness of the results on alternate adaptation margins to dropping outliers is shown in Table A.5. Here dropping the 5 extreme heat outliers (0.003% of the sample) does have a substantial effect on farm area, on the number of farms, and on farm land values. When these outliers are dropped, extreme heat coefficients on these variables drops by at least 60-70% and becomes statistically insignificant. For the reason we focus on the results for the trimmed sample in the main text.

Table A.3: Robustness of long difference results to addition of county control variables

	(1)	(2)
GDD below threshold	0.0000	0.0003*
	(0.0003)	(0.0001)
GDD above threshold	-0.0046***	-0.0043***
	(0.0010)	(0.0008)
Precip below threshold	0.0429***	0.0300***
	(0.0154)	(0.0104)
Precip above threshold	0.0022	0.0028***
	(0.0015)	(0.0009)
Observations	1525	1525
R squared	0.387	0.637
Fixed Effects	None	State
Controls	Yes	Yes

Table displays long difference regression results from 1980-2000. Standard errors for each regression are clustered at the state level. Temperature threshold is 29°C and precipitation threshold is 42 cm. Control variables are population density in 1980, total farm area in 1978, total corn area in 1980, total county area, irrigated area in 1982, average farm value in 1978, percent of soil that is clay, water capacity of soil, percent of soil that is high quality, and income per capita in 1978. Column 2 also includes state fixed effects. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

Table A.4: Robustness of corn yield results to dropping outliers

	(1)	(2)	(3)	(4)
	full	trimmed	full	trimmed
GDD below threshold	0.0002	0.0002	0.0003*	0.0002
	(0.0002)	(0.0002)	(0.0002)	(0.0001)
GDD above threshold	-0.0044***	-0.0043***	-0.0037***	-0.0032***
	(0.0008)	(0.0009)	(0.0009)	(0.0009)
Precip below threshold	0.0297**	0.0309**	0.0115**	0.0117**
	(0.0125)	(0.0130)	(0.0046)	(0.0045)
Precip above threshold	0.0034***	0.0034***	0.0029***	0.0030***
	(0.0008)	(0.0008)	(0.0007)	(0.0007)
Constant	0.2397***	0.2403***	0.2400***	0.2409***
	(0.0124)	(0.0125)	(0.0115)	(0.0118)
Observations	1531	1521	1531	1521
R squared	0.610	0.624	0.602	0.617
Fixed Effects	State	State	State	State
T threshold	29	29	28	28
P threshold	42	42	50	50

All regressions use log of corn yields as the dependent variable, and use temperature and precipitation thresholds as indicated at the bottom of the table. Columns 1 and 3 are on the full sample, columns 2 and 4 drop the outliers indicated in Figure A.2. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

Table A.5: Robustness of results on alternate adaptations to removal of outliers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Corn area	Corn area	Corn %	Corn %	Farm area	Farm area	# farm	# farm
GDD below threshold	0.0013	0.0010	0.0003***	0.0003***	-0.0000	-0.0001	-0.0001	-0.0002
	(0.0012)	(0.0012)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0002)
GDD above threshold	-0.0038	-0.0005	-0.0007**	-0.0009**	-0.0010**	0.0000	-0.0023*	-0.0007
	(0.0048)	(0.0038)	(0.0003)	(0.0004)	(0.0004)	(0.0004)	(0.0011)	(0.0010)
Precip below threshold	0.0404	0.0264	-0.0006	-0.0004	0.0067	0.0037	0.0077**	0.0021
	(0.0485)	(0.0637)	(0.0030)	(0.0034)	(0.0043)	(0.0035)	(0.0036)	(0.0029)
Precip above threshold	-0.0079	-0.0051	-0.0015	-0.0016	-0.0001	0.0007	-0.0023	-0.0013
	(0.0067)	(0.0063)	(0.0010)	(0.0010)	(0.0008)	(0.0007)	(0.0032)	(0.0033)
Constant	-0.0273	-0.0130	-0.0164***	-0.0174***	-0.0649***	-0.0614***	-0.1881***	-0.1836***
	(0.0661)	(0.0687)	(0.0045)	(0.0045)	(0.0079)	(0.0075)	(0.0156)	(0.0157)
Observations	1516	1511	1521	1516	1528	1523	1531	1526
Mean of Dep Variable	0.073	0.075	0.002	0.002	-0.069	-0.068	-0.202	-0.202
R squared	0.642	0.645	0.418	0.418	0.387	0.399	0.478	0.488
Fixed Effects	State	State	State	State	State	State	State	State
T threshold	29	29	29	29	29	29	29	29
P threshold	42	42	42	42	42	42	42	42

Dependent variable is difference in log of corn acres (Columns 1-2), difference in share of agricultural area planted to corn (Columns 3-4), difference in total log farm area (Columns 5-6), and difference in log number of farmers (Columns 7-8). All regressions are long differences from 1980-2000, and even numbered regressions have the sample trimmed of outliers indicated in Figure A.2. All regressions are weighted by average agricultural area from 1978-2002. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels..

A.5 Choice of time period

Our main specification focuses on changes in climate and yields for the 1980-2000 period. We focus on this period for a few reasons. First, relative to earlier periods, and as shown in Figure A.11, global warming had begun in earnest over this period, and counties had experienced on average much more warming. Importantly, many more counties had experienced at least 1C warming over the period, making this period more representative of the warming that climate models predict will occur over the next few decades and thus a better baseline with which to project future impacts. Second, prior to 1980, scientific opinion was relatively split as to whether the future climate would be cooler or warmer than the current climate, and in fact there was significant concern about "global cooling" (e.g. Gwynne (1975)). Growing scientific and public recognition of "global warming" during the 1980's and 1990's – i.e. a recognition that increasing greenhouse gas emissions would lead to future warming – again makes this period more relevant for projecting future impacts because there was recognition that the climate was warming and would continue to warm.

Nevertheless, Figure 4 directly compares our benchmark 1980-2000 estimate to estimates using alternate time periods and differencing lengths, and shows that these alternate estimates are largely indistinguishable from our main estimate. Figure A.10 displays point estimates and their confidence intervals from each of these regressions (rather than comparisons with the 1980-2000 estimate); all of these estimates are negative, and in only 8 out of 39 cases do we fail to reject no effect of extreme heat on corn yields.

As a final robustness test on our choice of time period, we vary the length of the endpoints over which our two periods are averaged. We begin with our 1980-2000 period, and average our endpoints over ten years instead of five – i.e. the long difference is now 1995-2005 average minus 1975-1985 average. We include in the sample any counties that reported growing corn in at least one year in both periods, or restrict the sample to counties that grew corn in all years in both averaging periods. Results are given in Columns 1 and 2 of Table A.6. Coefficients on extreme heat in both specifications are slightly more negative than our baseline estimates and highly significant.

Finally, we utilize our full 1950-2005 sample, split it into 28-year periods (1950-1977 and 1978-2005), average both climate and crop yields within each period, difference these averages, and then run our basic long differences specification on these two time periods. This is equivalent to smoothing our data with a 28-year running mean, and then differencing between the years 1991 and 1964. We similarly restrict the sample to include either all counties reporting growing corn in at least one year in both periods (column 3), or successively limit the sample to counties with at least 40, 50, or 56 observations (columns 4-6).

The coefficient on GDD above 29C is again large, negative, and highly significant across all specifications. Point estimates are in fact substantially more negative than for our baseline 1980-2000 period. One explanation for this is that farmers have become less sensitive to temperature over time, with our main 1980-2000 specification focusing on a later (and thus less sensitive) period. But both Figure 4 and Figure A.12 (see discussion below) show that there is little evidence that temperature sensitivities have declined over time. We can also run the panel model for the full 1950-2005 period (shown in column 7 of Table A.6), and we find that the panel coefficient on extreme heat is somewhat more negative that for the 1980-2000 period but not substantially so. An alternate explanation is that if measurement

Figure A.10: *Long difference estimates under various starting years and differencing lengths. Dots are point estimates and whiskers are 95% confidence intervals.*

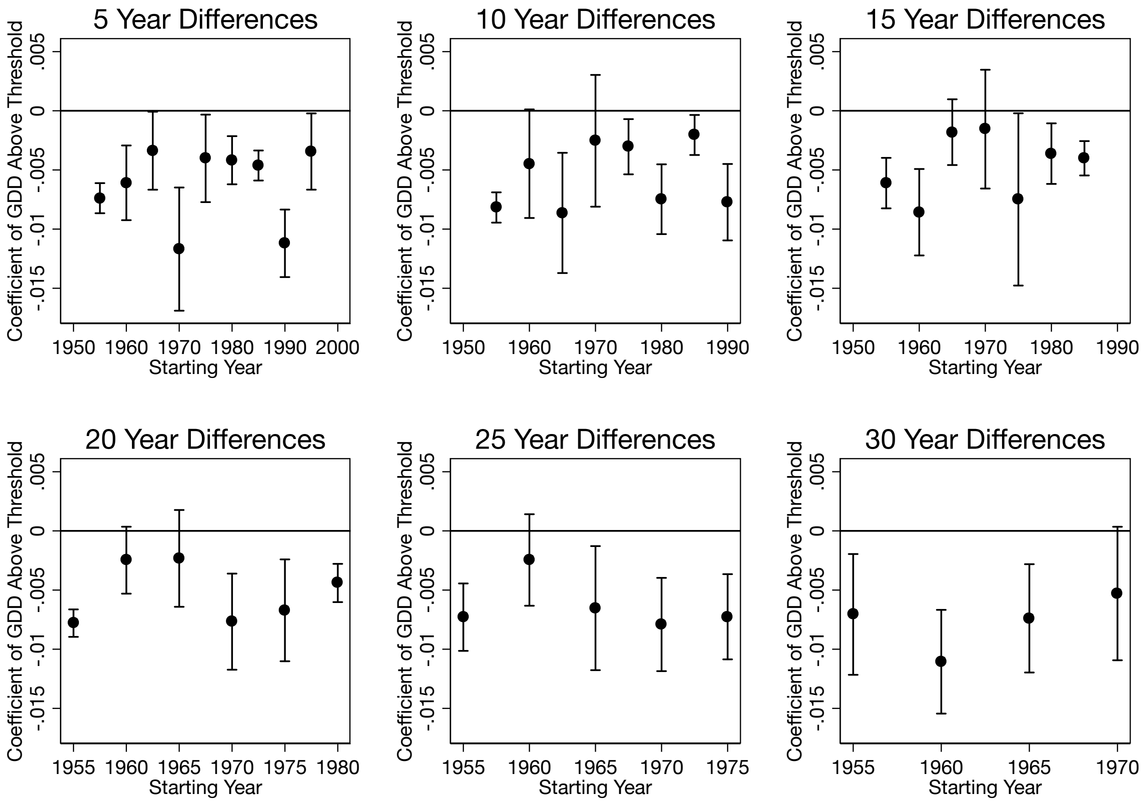

error in temperature is uncorrelated across years, then averaging over more years will reduce attenuation bias, resulting in larger (in absolute value) coefficients. While this explanation is hard to either support or rule out with the data, it appears more plausible than declining sensitivities.

Nevertheless, we cannot reject that the long differences estimates for the full period are the same than the panel estimates over the same period, and so these results do not suggest a qualitative or quantitatively different conclusion from that which we draw from our baseline specification. We view these results as yet more evidence that farmers have been unable to adapt very effectively in the long run, and these results suggest that our baseline estimates are somewhat conservative in terms of levels effects of extreme heat on yields..

We conduct analogous exercise for our panel results, to ensure that our panel estimates are also not being driven by our choice of time period. Since our data span five decades, we estimate our main panel regressions for each decade from the 1950's to the 1990's (results from running the panel on the full dataset are given in the last column in Table A.6). In Figure A.12 we show the coefficient on GDD above 28C and its 95% confidence interval for each of these five regressions. The estimates vary only slightly between decades and

Table A.6: Long differences regressions with endpoints averaged over longer periods.

	(1) 1980-2000	(2) 1980-2000	(3) 1950-2005	(4) 1950-2005	(5) 1950-2005	(6) 1950-2005	(7) panel 50-05
GDD below	0.0000	0.0000	0.0008**	0.0008**	0.0009**	0.0010**	0.0004***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0001)
GDD above	-0.0050***	-0.0047***	-0.0090***	-0.0091***	-0.0102***	-0.0096**	-0.0065***
	(0.0010)	(0.0011)	(0.0025)	(0.0027)	(0.0030)	(0.0034)	(0.0006)
Precip below	0.0234***	0.0245***	0.0480***	0.0460***	0.0457***	0.0450***	0.0172***
	(0.0045)	(0.0045)	(0.0047)	(0.0049)	(0.0049)	(0.0046)	(0.0022)
Precip above	0.0026	0.0027	0.0021	0.0023	0.0022	0.0020	-0.0020***
	(0.0019)	(0.0019)	(0.0033)	(0.0035)	(0.0036)	(0.0039)	(0.0003)
Constant	0.2805***	0.2809***	0.5206***	0.5232***	0.5237***	0.5117***	2.7418***
	(0.0074)	(0.0080)	(0.0101)	(0.0110)	(0.0123)	(0.0132)	(0.2234)
Observations	1950	1451	2241	1711	1262	956	107290
Mean of Dep Var.e	0.31	0.31	0.57	0.57	0.58	0.57	4.42
R squared	0.627	0.670	0.687	0.719	0.722	0.734	0.821
Period 1 years	1975-1985	1975-1985	1950-1977	1950-1977	1950-1977	1950-1977	
Period 2 years	1995-2005	1995-2005	1978-2005	1978-2005	1978-2005	1978-2005	
Min. yrs in sample	2	All	2	40	50	All	Any

All regressions use log of corn yields as the dependent variable, and use the 29C temperature and 42cm precipitation thresholds. The sample period is either 1975-2005 (columns 1-2) or 1950-2005 (columns 3-7), with endpoints averaged over the two different periods given in the bottom of the table. As indicated in the last line of the table, samples either include counties that report at least one year of growing corn in both periods (columns and 3), counties that grew corn in all years in the sample (columns 2 and 6), or counties that grew corn in at least 40 or 50 years of the 56-year sample (columns 4 and 5). The final model (column 7) is a panel model over the full 1950-2005 period. Standard errors are clustered at the state level, and asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

Figure A.11: *Distribution of the change in average growing season temperature across our sample counties, for the period 1960-1980 (dotted line) or the period 1980-2000 (solid line).*

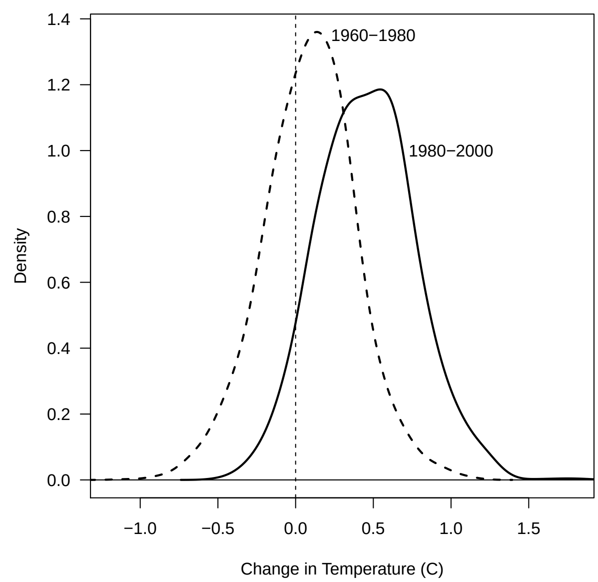

there is no clear pattern suggesting that corn yields have become less sensitive to short-term deviations in weather over time.

While this unchanging sensitivity of yield to extreme heat over time could be interpreted as additional evidence of a lack of adaptation (as in Schlenker and Roberts (2009)), we note that whether responses to short-run variation have changed over time is conceptually distinct from whether farmers have responded to long-run changes in average temperature. In particular, there is no reason to expect farmers to respond similarly to these two different types of variation. Indeed, farmers could adapt completely to long-run changes in temperature such that average yields do not change – e.g. by adopting a new variety that on average performs just as well in the new expected temperature as the old variety did under the old average temperature – but still face year-to-year variation in yield due to random deviations in temperature about its new long-run average. As such, we view this exercise more as a test of the robustness of the panel model than as evidence of (a lack of) adaptation per se.

A.6 Measurement error

As discussed in the main text, one concern is that fixed effects estimators are more likely than long differences estimates to suffer attenuation bias if climate variables are measured with error. Following Griliches and Hausman (1986), we compare fixed effectsand first difference estimates with random effects estimates, with the expectation that if measurement error in our climate variables is a problem, then estimates from a random effects estimation should be larger in absolute value than the fixed effects estimates which in turn should be larger than estimates using first differences.

Figure A.12: *Panel estimates of the effect of extreme heat on log corn yields by decade. Figure shows point estimate and 95% confidence interval for regressions run separately for each decade. The black line is the coefficient on extreme heat from our baseline panel regression (Column 3 in Table 1) . All regressions include county and time fixed effects and are weighted by average corn area in the county during the relevant decade, with errors clustered at the state level.*

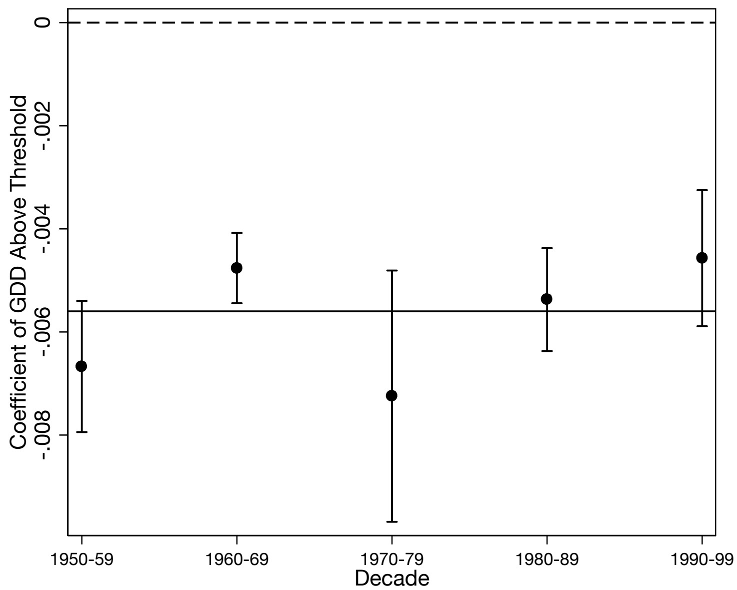

Table A.7 presents the results of a horse race between these three estimators. The first column presents unweighted fixed effects estimates. The random effects estimates in Column 2 are remarkably similar to the fixed effects estimates. The main coefficient of interest for GDD above 28° is *smaller* in absolute value by a modest 7%. Column 3 shows that the first difference estimator also produces a very similar effect of increases in temperatures above 28° on yields. Results suggest that measurement error is not responsible for the lack of difference between fixed effects estimators and long differences that we observe in the data.

A.7 Functional Form

Our use of growing degree days to capture nonlinearities is primarily motivated by results from the agronomy literature suggesting that plant growth increases linearly with temperature up to a certain threshold level, and then declines with further temperature increases. Figure 3 in the main text shows that our results produce this relationship. Our piecewise linear approximation will be misspecified in the presence of strong nonlinearities within the ranges from 0-29 and 29 and above. Schlenker and Roberts (2009) show that the piecewise linear relationship achieved with growing degree days estimates that use either a higher order polynomial or a set of temperature bins measuring the days of exposure to various temperature ranges. These results strongly suggest that use of growing degree days is not affected

Table A.7: Understanding measurement error through the comparison of panel estimators

	(1) Fixed Effects	(2) Random Effects	(3) First Difference
GDD below threshold	0.0004***	0.0003***	0.0003***
	(0.0001)	(0.0001)	(0.0001)
GDD above threshold	-0.0045***	-0.0042***	-0.0045***
	(0.0005)	(0.0005)	(0.0004)
Precip below threshold	0.0045**	0.0040**	0.0054***
	(0.0018)	(0.0017)	(0.0018)
Precip above threshold	-0.0011*	-0.0011	-0.0009
	(0.0006)	(0.0007)	(0.0006)
Constant	3.2154***	3.6483***	0.0703**
	(0.2877)	(0.1648)	(0.0343)
Observations	48465	48465	45405
R squared	0.463		0.494
Fixed Effects	Cty, Yr	Yr	Yr
T threshold	28	28	28
P threshold	50	50	50

All regressions use log of corn yields as the dependent variable. All regressions are unweighted. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

by misspecification due to nonlinearity.

Nevertheless, we re-estimate both our main panel and long difference specifications using three degree bins. In both models we include the same functions of precipitation as were included in the main specifications. Figure A.13 shows the results. Both the panel and long difference specifications using temperature bins produce similar results to those using growing degree days, consistent with Schlenker and Roberts (2009). Our use of a piecewise linear function of growing degree days does not seem to misrepresent the relationship between temperature and yields.

A.8 Effects on soy productivity

Estimates of the impact of extreme heat on (log) soy yields are shown in Figure A.14. The horizontal line in each panel is the 1978-2002 panel estimate of β_2 for soy which is -0.0047. The thresholds for temperature and precipitation are 29° and 50 cm, which are those that produce the best fit for the panel model. The average response to extreme heat across the 39 estimates is -0.0032, giving a point estimate of longer run adaptation to extreme heat of about 30%. This estimate is slightly larger but of similar magnitude to the corn estimate, and we are again unable to reject that the long differences estimates are different than the panel estimates. As for corn, we conclude that there is limited evidence for substantial adaptation of soy productivity to extreme heat.

Figure A.13: *Relationship between corn yields and temperature. Estimates represent the change in log corn yield under an additional day of exposure to a given °C temperature, relative to a day spent at 0-3°C. Estimates of 3°C temperature bins are used for long difference and panel versions of binned regressions. Dots represent midpoints of bins. GDD regressions are identical to those in Figure 3 of the main text. The shaded area is the confidence interval of the long difference estimates where temperature is measured with GDD.*

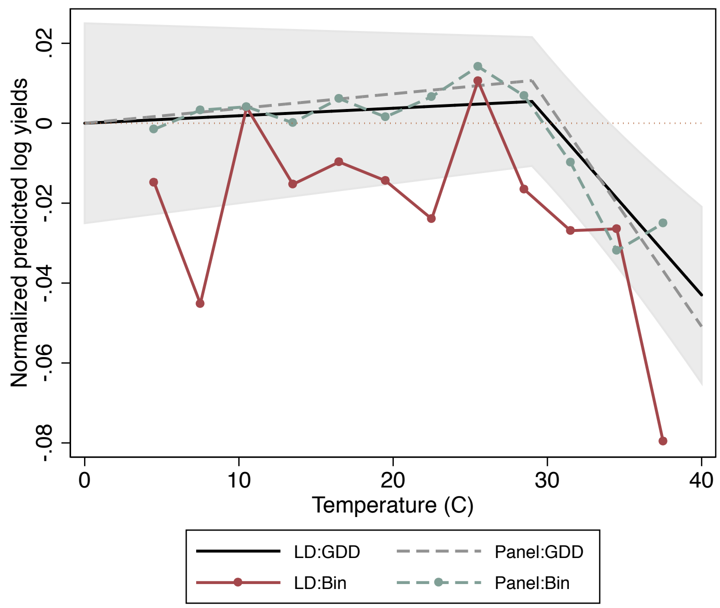

A.9 Revenues and profits

A basic concern with our crop yield results is that they could hide alternate adjustments that help farmers maintain profitability in the face of a changing climate. The US Agricultural Census, conducted roughly every 5 years, contains data on overall farm revenues and expenses for the year in which the census is conducted. A basic measure of profits for a given year can be constructed by differencing these two variables (i.e. $\text{profits}_{2000} = \text{revenues}_{2000}$ - costs_{2000} for years in which data are available, and and this approach as recently been used in similar settings (Deschênes and Greenstone, 2007).

We choose not to focus on such a profit measure for two reasons. The first is a concern that costs are not fully measured, and that unmeasured costs might respond to climate shocks in a way that would bias the above profit measure. In particular, expense data do not appear to include the value of own or family labor, which could respond on the intensive or extensive margin in the face of a drought or heat event (e.g. if a crop fails and is replanted).[1] The second concern is that both costs and revenues will likely respond to annual variation

[1] In recent years, the value of own labor appears to represent about 10% of operating costs for corn, based on cost estimates available at http://www.ers.usda.gov/data-products/commodity-costs-and-returns.aspx. Hired labor expenditures are minimal for corn.

Figure A.14: Effects of extreme heat on soy yields under various starting years and differencing lengths, as compared to the point estimate from a 24-year panel estimated over 1978-2002 displayed by the horizontal line in each figure panel.

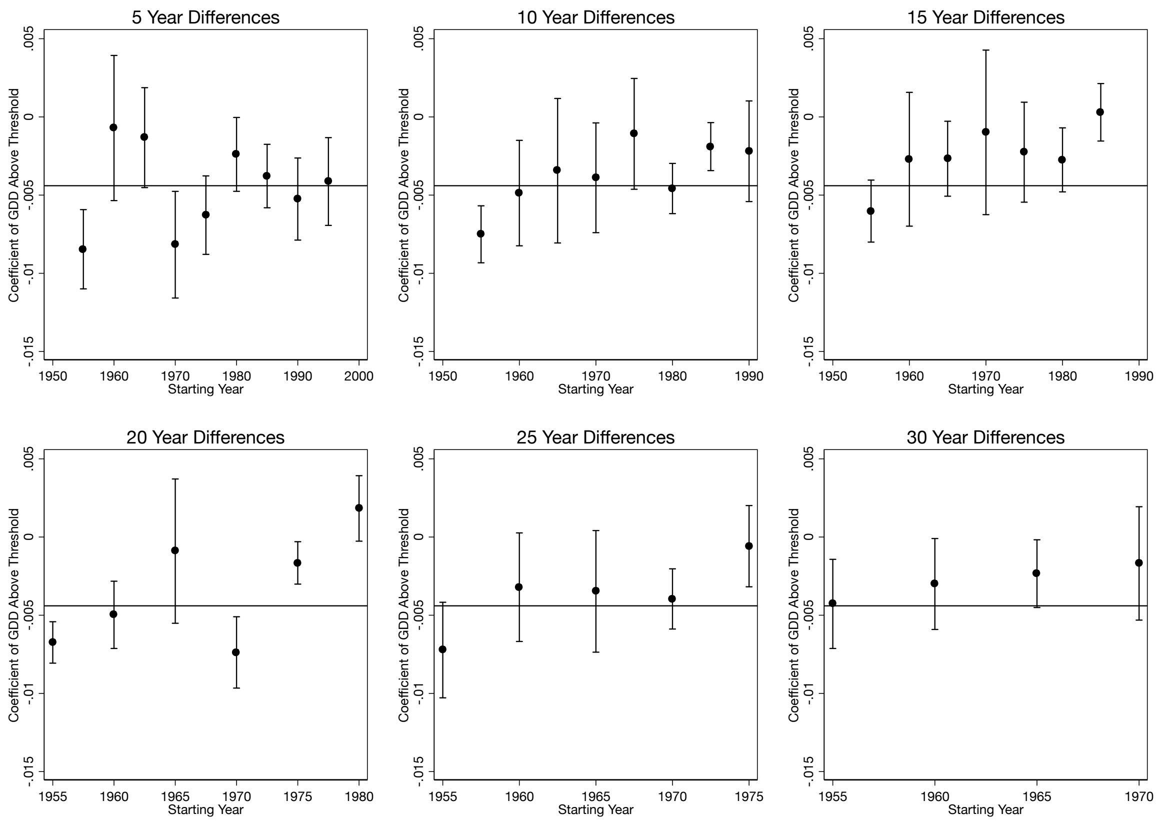

in climate, but data are only available for 5-year snapshots. Given that our differencing approach seeks to capture change in average farm outcomes over time, differencing two of these snapshots might provide a very noisy measure of the overall change in profits.

Regressions appear to confirm that profit measures are quite noisy. Agricultural census data on expenditures and revenues are available in 1978, 1982, 1987, 1992, 1997, and 2002. We construct a measure of the change in log profits as:

$$\Delta logprofits_{1980-2000} = ln(profit_{1997} + profit_{2002})/2 - ln(profit_{1978} + profit_{1982})/2 \quad (9)$$

When we re-estimate our main specification with $\Delta logprofits_{1980-2000}$ as the dependent variable, the coefficient on extreme heat using the untrimmed sample is $\beta_2 = -0.0013$, with 95% CI of [-0.010, 0.007], and using the trimmed sample we have $\beta_2 = -0.0054$, with 95% CI of [-0.014,0.003]. This means we can't reject that there is no effect on profits, and similarly can't reject that the effect of extreme heat on profits is a factor of 3 larger (and more negative) than the effect on corn yields – i.e. that each additional day of exposure to temperatures above 29C reduced *annual* profits by 1.4%. This does not provide much insight on the relationship between extreme heat exposure and profitability.

We take two alternate approaches to exploring profitability impacts that help to address these measurement issues. The first is to construct a measure of revenues using annual yield data, which we multiply by annual data on state-level prices to obtain revenue-per-acre for a given crop. Summing up these revenues across crops then provides a reasonable measure of annual county-level crop revenues, which will be underestimated to the extent that not all contributing crops are included. The effect of climate variation on this revenue measure is given in the main text, and we find minimal difference between panel and long difference estimates of impacts on expenditures.

Our second approach proceeds with the available expenses data from the ag census to examine the impact of longer-run changes in climate on different input expenditures, where we attempt to capture changes in *average* expenditures by averaging two census outcomes near each endpoint and then differencing these averaged values.[2] As shown in Table A.8, we find little effect of long-run trends in climate on expenditures on fertilizer, seed, chemical, and petroleum. While we do not wish to push these expenditure data too far given the noisy way in which the long differences are constructed, we interpret these as further evidence that yield declines are economically meaningful and not masking other adjustments on the expenditure side that somehow reduce profit losses.

Table A.8: Effects of Climate Variation on Input Expenditures

	(1) Fertilizer	(2) Seed	(3) Chemicals	(4) Petroleum
GDD below threshold	0.0005	0.0008**	0.0011*	0.0002
	(0.0004)	(0.0004)	(0.0006)	(0.0004)
GDD above threshold	-0.0007	-0.0009	-0.0001	-0.0009
	(0.0015)	(0.0013)	(0.0034)	(0.0011)
Precip below threshold	0.0141	-0.0105	0.0392***	-0.0016
	(0.0229)	(0.0125)	(0.0115)	(0.0087)
Precip above threshold	-0.0016	-0.0021	0.0004	0.0010
	(0.0019)	(0.0024)	(0.0036)	(0.0019)
Constant	0.3215***	0.7295***	0.6993***	0.0281
	(0.0276)	(0.0217)	(0.0338)	(0.0237)
Observations	1528	1519	1523	1518
R squared	0.532	0.313	0.460	0.258
Fixed Effects	State	State	State	State
T threshold	29	29	29	29
P threshold	42	42	42	42

Dependent variable is difference in log of input expenditure per acre. All regressions are long differences from 1980-2000. All regressions are weighted by average agricultural area between 1978-1982. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

[2] As with the profit measure described above, the change in fertilizer expenditures over the period are constructed as: Δfertilizer expenditure$_{1980-2000}$ = (fert$_{1997}$ + fert$_{2002}$)/2 - (fert$_{1978}$ + fert$_{1982}$)/2

A.10 Exit from agriculture

As an extension to our basic long difference results on how the number of farms change in response to climate variation, we adopt an empirical strategy similar to that of Hornbeck (2012). We use the six agricultural censuses from 1978-2002 to estimate whether the number of farms grew differently between areas that were differentially exposed to extreme heating from 1970-1980. We first take the difference between average annual GDD above 29° from 1976-1980 and average annual GDD above 29° from 1966-1970. We then define extreme heating as an indicator variable for this difference being above a certain value. The econometric specification is,

$$ln(farms)_{ist} - ln(farms)_{is1978} = \beta_t * Extremeheat_{is} + \alpha_{st} + \varepsilon_{ist}, \tag{10}$$

where $Extremeheat_{is}$ is an indicator variable for a large change in GDD above 29. An important note is that the census defines a farm to be any place where at least \$1000 in agricultural products was sold during that year. Table A.9 reports estimates with and without state-specific time fixed effects. The state specific time-effect eliminates all state-specific factors varying over time. For instance, if heating was more heavily concentrated in some states and those states had different policies over time, the state-specific time effects would control for this correlation. We also show two different definitions of extreme heat. In the first definition it is defined as an indicator for an increase in GDD above 29° of 10 or more. This results in approximately 48% of counties being classified as having been exposed to heating. The second definition uses a stricter cutoff of 20. This results in 28% of counties being classified as exposed to heating. Each coefficient β_t measures the predicted percentage difference in the number of farms in year t between the counties that warmed from 1970-1980 and those that did not. For instance in Column 2, the number of farms in 1982 is predicted to be 2.75% lower in counties that heated substantially from 1970-1980. This predicted difference increases to 3.5% in 1987. The predicted difference in the number of farms generally becomes smaller in the later years of 1997 and 2002 which is consistent with some longer term adjustments back towards pre-warming degree of farming activity. This interpretation must be made with caution given the large standard errors in these years. The pattern of coefficients suggests that simply not farming may be an important immediate adaptation to climate change.

A.11 Additional evidence on selection

The potential of exit from agriculture and migration as responses to climate change highlights an important potential issue with our estimates of the effects of long-term climate trends on yields. If exit/migration is selective, then the appearance of a lack of adaptation in the data could be due to a selection effect where the most productive farmers recognize the changing climate and leave agriculture. In this case the appearance of a lack of adaptation in the data could be due to the change in the ability of the farming population that results from climate change. This possibility would become especially problematic if farmers that were more productive and had access to better quality land also had a larger opportunity cost of being in farming. If selection of this type is driving our estimates then we should see characteristics that are correlated with productivity changing differentially between places

Table A.9: Estimated Differences in Log Number of Farms by Amount of Warming

	Extreme Heat=Change GDD > 10		Extreme Heat=Change GDD > 20	
	(1)	(2)	(3)	(4)
1982*Extreme Heating	-0.0585***	-0.0275**	-0.0741***	-0.0230***
	(0.0177)	(0.0107)	(0.0231)	(0.0057)
1987*Extreme Heating	-0.0579***	-0.0352**	-0.0727***	-0.0455**
	(0.0190)	(0.0166)	(0.0205)	(0.0216)
1992*Extreme Heating	-0.0351	-0.0396	-0.0460**	-0.0430**
	(0.0223)	(0.0240)	(0.0191)	(0.0179)
1997*Extreme Heating	0.0051	-0.0155	-0.0016	-0.0221
	(0.0296)	(0.0318)	(0.0216)	(0.0171)
2002*Extreme Heating	0.0174	-0.0169	0.0045	-0.0617*
	(0.0351)	(0.0299)	(0.0352)	(0.0318)
Observations	12120	12120	12120	12120
Mean of Dep Variable	-0.13	-0.13	-0.13	-0.13
R squared	0.617	0.681	0.618	0.681
State by Year Fixed Effects	No	Yes	No	Yes

Data are for US counties east of the 100th meridian. Dependent variable in all specifications is difference between log number of farms in year t and log number of farms in 1978. Coefficients represent estimated differences in log number of farms between counties that experienced extreme heating from 1970-1980 and those that did not. Extreme heating defined as indicator for increases in GDD above 29 greater than cutoff value of 10 (Columns 1-2) or 20 (Columns 3-4). All regressions are weighted by county farm area in 1978. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

that heated and those that did not. In Table A.10 we regress the percentage of farms owning more than $20,000 in equipment on our same climate variables. Since the percentage of farms owning valuable equipment is positively correlated with yields, if selection is driving our results we should expect to see a large decrease as a response to increases in extreme temperatures. The results are not consistent with this story. The long differences estimate is negative, but small and not statistically significant from zero. The panel estimate is positive, small in magnitude and marginally statistically significant. While we obviously can not fully rule out selective migration, these regressions are suggestive that it is not driving our yield results.

A.12 Why no adaptation

To provide additional evidence on whether an absence of learning was what constrained adaptation, we check for mean reversion in temperatures. Even if there are "real" temperature changes during a given decade, the longer-term mean might not change if a period of warming was then followed by a period of cooling. If farmers know about this cyclicality in temperature, it therefore might make sense to not adapt to a temperature increase. To test for mean reversion in temperature we compare our ΔGDD>29 over the main 1980-2000

Table A.10: Effects of climate variation on equipment ownership.

	(1) Diffs, 1978-1997	(2) Panel, 1978-2002
GDD below threshold	0.0087	-0.0067***
	(0.0152)	(0.0019)
GDD above threshold	-0.0178	0.0221*
	(0.0318)	(0.0109)
Precip below threshold	0.2114	0.0608
	(0.1470)	(0.0499)
Precip above threshold	0.0524	0.0760***
	(0.1147)	(0.0250)
Constant	9.9251***	74.5041***
	(0.9928)	(6.0013)
Observations	1531	7645
Mean of Dep Variable	10.50	59.01
R squared	0.321	0.324
Fixed Effects	State	Cty, Yr
T threshold	28	28
P threshold	50	50

Dependent variable in Column 1 is the change in the percentage of farms with more than 20K USD in equipment from 1978 to 1997. Dependent variable in Column 2 is the percentage of farms owning equipment valued at more than 20,000 USD. Long differences regressions are weighted by average farm acres between 1978 and 1982. Panel regressions weighted by average farm acres from 1978-2002. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

period with changes in the previous 1955-1975 period (with 5-year averaging at endpoints, we are then using two non-overlapping but contiguous periods from 1953-1977 and from 1978-2002). The data are shown in the right panel of updated Figure A1, and estimating the following regression:

$$\Delta GDD_{is}^{1980-2000} = \alpha + \beta GDD_{is}^{1955-1975} + \eta_s + \varepsilon_{is} \tag{11}$$

(where i is county and s indicates state) gives an estimate of β which is *positive* but small ($\beta = 0.10$) and statistically insignificant with state-level clustering. This is inconsistent with a mean reversion story: although many areas did cool during 1955-1975, these were not on average the areas that differentially warmed over the subsequent 20 years. This suggests that, based on the historical record, farmers would have no reason to believe the 1980-2000 changes were impermanent.

Furthermore, we note that the scientific literature provides very strong evidence that future temperatures across the US are going to continue to increase for centuries – a conclusion that was already understood and publicized by the 1980s, and solidified with the release of the IPCC's First Assessment Report in 1990. This report concluded that mean temperatures were likely to increase by 0.3C/decade over the next century, with land areas heating up faster than oceans. To the extent that farmers were aware of what scientists were saying (and other papers in this literature, e.g. Kelly, Kolstad and Mitchell (2005), assume that they were), this again suggests that farmers who experienced warming during 1980-2000 would have no reason to believe that these changes were impermanent.

A.12.1 Insurance take-up

Take-up of government insurance programs in response to warming could provide evidence that farmers recognized that climate was changing. As described in the main text, Table A.11 provides some evidence that participation in the government insurance program by 2000 was higher in counties who saw large increases in exposure to harmful temperatures (GDD>29C) over the previous two decades, and lower in counties that saw increase in exposure to generally helpful temperatures (GDD0-29C) over the same period.

As an alternate approach, Figure A.15 looks at insurance uptake in a distributed lag panel framework, showing coefficients from distributed lag fixed effects regressions of various measures of insurance takeup on lags of GDD>29. These results provide some additional suggestive evidence that farmers update their expectations about future temperature exposure as a function of recent past temperature exposure: total acreage insured, share of acreage insured, and the number of policies sold all increases significantly with $GDD > 29$ for the previous two seasons, although they do not respond significantly to changes in extreme heat beyond this.

Table A.11: Insurance take-up in 1998-2002 as a function of changes in GDD and precipitation over 1980-2000.

	(1) % Acreage Enrolled	(2) log Acres Enrolled	(3) Policies Sold	(4) Total Premiums
GDD below threshold	-0.0006*	-0.0005	-1.6411**	-3.7636***
	(0.0003)	(0.0006)	(0.7021)	(1.2821)
GDD above threshold	0.0026	0.0022	8.2093*	16.8366**
	(0.0018)	(0.0025)	(4.8110)	(7.6828)
Precip below threshold	0.0354**	0.0309**	-3.3723	57.6410
	(0.0162)	(0.0138)	(21.7244)	(57.8886)
Precip above threshold	-0.0050**	-0.0052*	-9.3570	-13.4772
	(0.0019)	(0.0025)	(10.5879)	(17.6880)
log corn area		1.0057***		
		(0.0267)		
Constant	0.7929***	-0.3736	704.2308***	1250.4442***
	(0.0237)	(0.3035)	(43.4052)	(86.8633)
Observations	1529	1529	1529	1529
R squared	0.354	0.955	0.480	0.489
Mean Dep. Var.	0.815	9.329	271.227	428.441

The outcome variables are given at the top of each column. Total premiums paid (column 4) are in thousands of dollars. All regressions include state fixed effects, with standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

Figure A.15: Figures shows coefficients (dots) and standard errors (whiskers) from distributed lag fixed effects regressions of various measures of insurance takeup on lags of GDD>29. All regressions include lags of other climate variables and are weighted by 1978-2002 average corn area, as in our main panel specification. Regression for total acreage insured includes total corn area as a control. Standard errors are clustered at the state level.

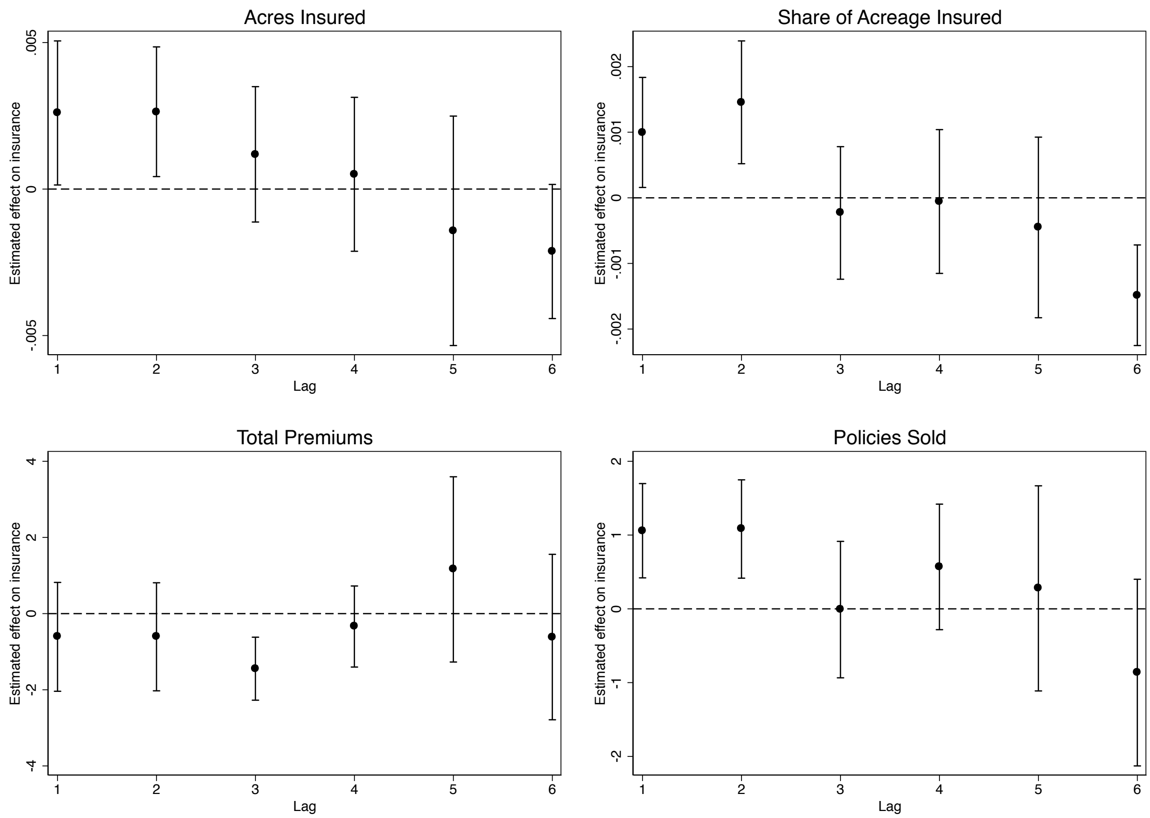

A.13 Climate change projections

We derive projected changes in corn productivity due to climate change by combining our long differences estimates of the the historical response of corn productivity to climate with climate projections from 18 general circulation models that have contributed to World Climate Research Programs Coupled Model Intercomparison Project phase 3 (WCRP CMIP3). Our main projections use the A1B emissions scenario, reported by 18 climate models in the CMIP3 database: CCMA, CNRM, CSIRO, GFDL0, GFDL1, GISS.AOM, GISS.EH, GISS.ER, IAP, INMCM3, IPSL, MIROC.HIRES, MIROC.MEDRES, ECHAM, MRI, CCSM, PCM, and HADCM3. For more on these models and their application, see Auffhammer et al. (2013) and Burke et al. (2013). The A1B scenario is considered a "medium" emission scenario, and represents a world experiencing "rapid and successful economic development" and a "balanced mix of energy technologies" (Nakicenovic et al., 2000). We choose to explore outcomes under only one emissions scenario both to simplify the results, and because emissions scenarios diverge much less by mid-century than they do by the end of the century, meaning our results are less sensitive to the choice of emissions scenario than end-of-century projections. Finally, following the climate literature, we adopt a "model democracy" approach and assume projections from all models are equally valid and should be weighted equally (Burke et al., 2013).

The resolution of these general circulation models is roughly 2.8°x2.8° (about 300km at the equator), and we map each county in our sample to its corresponding grid cell in the climate model grids. We derive estimates of climate change by mid-century by calculating model-projected changes in temperature (C) and precipitation (%) between 2040-2059 and 1980-1999, and then adding (for temperature) or multiplying (for precipitation) these changes to the observed record of temperature and precipitation in a given county. For temperature, because our main variable of interest is growing degree days, this requires adding monthly predicted changes in temperature in a given county to the daily time series series in that county, recomputing growing degree days under this new climate, and calculating the difference between baseline and future growing degree days.

Projections assume a fixed growing season (Apr 1 - Sept 30) and no large shifts in the area where corn is grown within the US. Area-weighted changes in temperature and precipitation over US corn area are shown in Figure A.16. The variation in temperature changes over our 1980-2000 study period span the lower third of the range of model-projected average temperature changes by 2050, and the variation in changes in precipitation in our sample fully span the range of projected average precipitation changes by 2050.

Figure A.16: *Projected changes in growing season temperature and precipitation across US corn growing area by 2050. Each dot represents a projection from a particular global climate model running the A1B emissions scenario.*

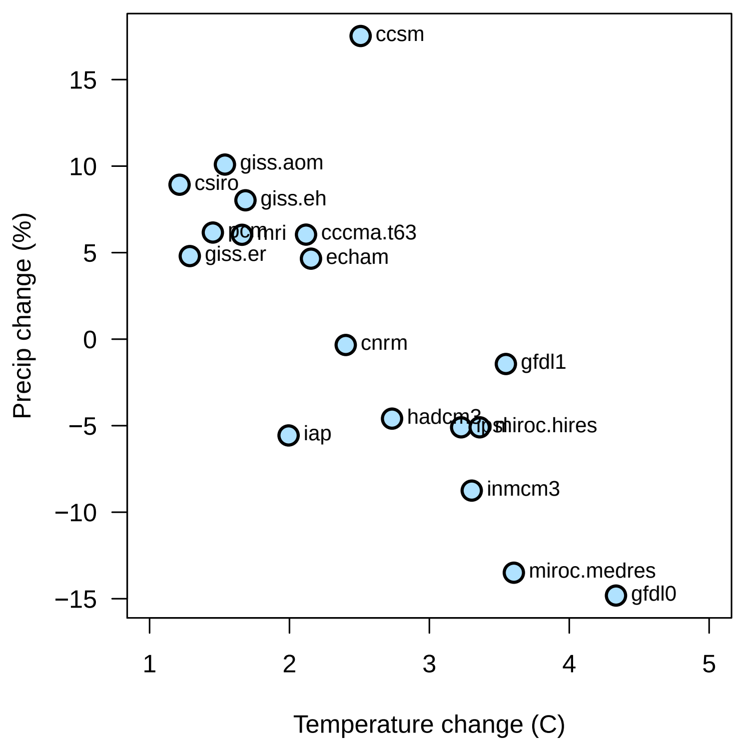

References

Auffhammer, M., S. Hsiang, W. Schlenker, and A. Sobel. 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy*, 7(2).

Burke, M., J. Dykema, D. Lobell, E. Miguel, and S. Satyanath. 2013. "Incorporating climate uncertainty into estimates of climate change impacts." *Working paper*.

Deschênes, O., and M. Greenstone. 2007. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather." *AER*, 97(1): 354–385.

Griliches, Z., and J.A. Hausman. 1986. "Errors in variables in panel data." *J. Econometrics*, 31(1): 93–118.

Gwynne, P. 1975. "The cooling world." *Newsweek*, 28: 64.

Hornbeck, Richard. 2012. "The Enduring Impact of the American Dust Bowl: Short and Long-Run Adjustments to Environmental Catastrophe." *American Economic Review*, 102(4).

Nakicenovic, N., J. Alcamo, G. Davis, B. De Vries, J. Fenhann, et al. 2000. "Special Report on Emissions Scenarios (SRES). Working Group III of the Intergovernmental Panel on Climate Change (IPCC), Ed IPCC." Cambridge University Press, Cambridge, UK.

Schlenker, W., and M.J. Roberts. 2009. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *PNAS*, 106(37): 15594–15598.