Adaptation to Climate Change: Evidence from US Agriculture

Marshall Burke and Kyle Emerick^{*}

January 14, 2015

Abstract

Understanding the potential impacts of climate change on economic outcomes requires knowing how agents might adapt to a changing climate. We exploit large variation in recent temperature and precipitation trends to identify adaptation to climate change in US agriculture, and use this information to generate new estimates of the potential impact of future climate change on agricultural outcomes. Longer-run adaptations appear to have mitigated less than half – and more likely none – of the large negative short-run impacts of extreme heat on productivity. Limited recent adaptation implies substantial losses under future climate change in the absence of countervailing investments.

JEL codes: N5, O13, Q1, Q54 **Keywords:** Climate change; adaptation; agriculture; climate impacts

^{*}Burke: Department of Environmental Earth System Science, and Center on Food Security and the Environment, Stanford University; mburke@stanford.edu. Emerick: Department of Economics, Tufts University; kyle.emerick@tufts.edu. We thank Max Auffhammer, Peter Berck, Melissa Dell, Richard Hornbeck, Sol Hsiang, David Lobell, Wolfram Schlenker, and seminar participants at Berkeley, Cornell, Davis, Minnesota, Stanford, UBC, the World Bank, the 2014 AAEA Annual Meeting, and the 2015 ASSA meeting for helpful comments and discussions. We thank Michael Roberts and Wolfram Schlenker for sharing data. All errors are our own.

1 Introduction

How quickly economic agents adjust to changes in their environment is a central question in economics, and is consequential for policy design across many domains (Samuelson, 1947; Viner, 1958; Davis and Weinstein, 2002; Cutler, Miller, and Norton, 2007; Hornbeck, 2012). The question has been a theoretical focus since at least Samuelson (1947), but has gained particular recent salience in the study of the economics of global climate change. Mounting evidence that the global climate is changing (Meehl et al., 2007) has motivated a growing body of work seeking to understand the likely impacts of these changes on economic outcomes of interest. Because many of the key climatic changes will evolve on a time-scale of decades, the key empirical challenge is in anticipating how economic agents will adjust in light of these longer-run changes. If adjustment is large and rapid, the resulting economic damages associated with climate change could be minimal. But if agents appear slow or unable to adjust on their own, overall damages from climate change could be much larger and of greater policy interest.

To understand how agents might adapt to a changing climate, an ideal but impossible experiment would observe two identical Earths, gradually change the climate on one, and observe whether outcomes diverged between the two. Empirical approximations of this experiment have typically either used cross-sectional variation to compare outcomes in hot versus cold areas (e.g. Mendelsohn, Nordhaus, and Shaw (1994); Schlenker, Hanemann, and Fisher (2005)), or have used variation over time to compare a given area's outcomes under hotter versus cooler conditions (e.g. Deschênes and Greenstone (2007); Schlenker and Roberts (2009); Deschênes and Greenstone (2011); Dell, Jones, and Olken (2012)). Due to omitted variables concerns in the cross-sectional approach, the recent literature has preferred the latter panel approach, noting that while average climate could be correlated with other time-invariant factors unobserved to the econometrician, short-run variation in climate within a given area (typically termed "weather") is plausibly random and thus better identifies the effect of changes in climate variables on economic outcomes.

While using variation in weather helps to solve identification problems, it perhaps more poorly approximates the ideal climate change experiment. In particular, if agents can adjust in the long run in ways that are unavailable to them in the short run¹, then impact estimates derived from shorter-run responses to weather might overstate damages from longer-run changes in climate. Alternatively, there could be short-run responses to inclement weather, such as pumping groundwater for irrigation in a drought year, that are not tenable in the

¹e.g. Samuelson's famed Le Chatelier principle, in which demand and supply elasticities are hypothesized to be smaller in the short run than in the long run due to fixed cost constraints.

long-run if the underlying resource is depletable (Fisher et al., 2012). Thus it is difficult to even sign the "bias" implicit in estimates of impacts derived from short-run responses to weather.

In this paper we exploit variation in longer-term changes in temperature and precipitation across the US to identify the effect of climate change on agricultural productivity, and to quantify whether longer-run adjustment to changes in climate has indeed exceeded shorterrun adjustment. Recent changes in climate have been large and vary substantially over space: as shown in Figure 1, temperatures in some counties fell by 0.5°C between 1980-2000 while rising 1.5°C in other counties, and precipitation across counties has fallen or risen by as much as 40% over the same period. We adopt a "long differences" approach and model county-level changes in agricultural outcomes over time as a function of these changes in temperature and precipitation, accounting for time-invariant unobservables at the county level and time-trending unobservables at the state level.

This approach offers three distinct advantages over existing work. First, unlike either the panel or cross-sectional approaches, it closely replicates the idealized climate change impact experiment, quantifying how farmer behavior responds to longer-run changes in climate while avoiding concerns about omitted variables bias. Second, observed variation in these recent climate changes largely spans the range of projected near-term changes in temperature and precipitation provided by global climate models, allowing us to make projections of future climate change impacts that do not rely on large out-of-sample extrapolations. Finally, by comparing how outcomes respond to longer-run changes in climate to how they respond to shorter run fluctuations as estimated in the typical panel model, we can test whether the shorter-run damages of climatic variation on agricultural outcomes are in fact mitigated in the longer-run. Quantifying this extent of recent climate adaptation in agriculture is of both academic and policy interest, and a topic about which there exists little direct evidence.

We find that productivity of the primary US field crops, corn and soy, is substantially affected by these long-run trends in climate. Our main estimate for corn suggests that spending a single day at 30°C (86°F) instead of the optimal 29°C reduces yields at the end of the season by about half a percent, which is a large effect.² The magnitude of this effect is net of any adaptations made by farmers over the 20 year estimation period, and is robust to using different time periods and differencing lengths.

To quantify the magnitude of any yield-stabilizing adaptations that have occurred, we then compare these long differences estimates to panel estimates of short-run responses to weather. Long run adaptations appear to have mitigated less than about half of the short-

²The within-county standard deviation of days of exposure to "extreme" temperatures above 29° C is 30, meaning a 1 standard deviation increase in exposure would reduce yields by 15%.

run effects of extreme heat exposure on corn yields, and point estimates across a range of specifications suggest that long run adaptions have more likely offset *none* of these short run impacts. We also show limited evidence for adaptation along other margins within agriculture: revenues are similarly harmed by extreme heat exposure, and farmers do not appear to be substantially altering the inputs they use nor the crops they grow in response to a changing climate.

We then examine different explanations for *why* adjustment to recent climate change has been minimal. For instance, adaptation could be limited because there are few adjustment opportunities to exploit, or alternatively because farmers don't recognize that climate has in fact changed and that adaptation is needed. Which explanation prevails is important for how we interpret our results, and in particular how they extrapolate to future warming scenarios. If farmers failed to adapt in the past because they did not recognize the climate was changing, but in the future they become aware of these changes and quickly adapt, then our findings would be a poor guide to future impacts of warming. On the other hand, if farmers had recognized the need for adaptation but were unable to do so, then their past responses to extreme heat exposure would provide a plausible "business-as-usual" benchmark for the impacts of future warming in the absence of unprecedented adaptation.

While we cannot directly observe farmer perceptions of climate change, there is both theoretical and empirical guidance on which locations should be more likely to have learned about the negative effects of extreme heat or to have recognized that the climate was changing: locations that faced larger exposure to extreme heat in an earlier period, locations where the underlying temperature variance is lower (making any warming "signal" stronger), locations with better educated farmers, or locations where voting behavior suggest that a belief in climate change is more likely. We find no evidence that farmers in such areas responded any differently to extreme heat exposure than farmers previously un-exposed, less educated, or in more climate-change-skeptical regions, providing suggestive evidence that adaptation was not limited by a failure of recognition. Nevertheless, our inability to directly observe farmer perceptions means that we cannot rule out that the observed lack of adaptation was driven by a difficulty in recognizing that climate was changing.

As a final exercise, we combine our long differences estimates with output from 18 global climate models to project the impacts of future climate change on the productivity of corn, a crop increasingly intertwined with the global food and fuel economy. Such projections are an important input to climate policy discussions, but bear the obvious caveat that they constrain future adjustment capabilities to what farmers were capable of in the recent past. Nevertheless, because our projections are less dependent on large out-of-sample extrapolation, and because they account for farmers' recent ability to adapt to longer-run changes in climate, we believe they are a substantial improvement over existing approaches. Our median estimate is that corn yields will be about 15% lower by mid-century relative to a world without climate change, with some climate models projecting losses as low as 7% and others as high as 64%. Valued at current prices and production quantities, this fall in corn productivity in our sample counties would generate annual losses of \$6.7 billion dollars by 2050. We note that a 15% yield loss is on par with the estimated 15-25% productivity losses resulting from the well-publicized "extreme" drought and heat wave that struck the US mid-west in the summer of 2012.³ Given the substantial role that corn plays in US agricultural production and the dominant role that the US plays in the global trade of corn, these results imply substantial damages to US producers and global consumers of corn if the more negative outcomes in this range are realized.

Our work contributes to the rapidly growing literature on climate impacts, and in particular to a host of recent work examining the potential impacts of climate change on US agriculture (Mendelsohn, Nordhaus, and Shaw, 1994; Schlenker, Hanemann, and Fisher, 2005; Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Fisher et al., 2012). We build on this work by directly quantifying how farmers have responded to longer-run changes in climate, and are able to construct projections of future climate impacts that account for this observed ability to adjust.

Methodologically our work is closest to Dell, Jones, and Olken (2012) and to Lobell and Asner (2003). Dell, Jones, and Olken (2012) focus on panel estimates of the impacts of country-level temperature variation on economic growth, but also use cross-country differences in recent warming to estimate whether there has been "medium-run" adaptation. Their point estimates suggest little difference between responses to short-run fluctuations and medium-run warming, but estimates for the latter are imprecise and not always significantly different from zero, meaning that large adaptation cannot be ruled out. Lobell and Asner (2003) study the effect of trends in average temperature on trends in US crop yields, finding that warmer average temperatures are correlated with declining yields. We build on this work by providing more precise estimates of recent adaptation, and by accounting more fully for time-trending unobservables that might otherwise bias estimates.

Our findings also relate to a broader literature on long-run economic adjustments. A body of historical research suggests that economic productivity often substantially recovers in the longer run after an initial negative shock (Davis and Weinstein, 2002; Miguel and Roland, 2011), and that in the long run farmers in particular are able to exploit conditions that

³For instance, see http://www.ers.usda.gov/topics/in-the-news/us-drought-2012-farm-and-food-impacts.aspx. Estimated losses in 2012 depend on whether the comparison is against previous season's yield or the yield projected at planting in 2012, and appear to range between roughly 15-25%.

originally appeared hostile (Olmstead and Rhode, 2011). Somewhat in contrast, Hornbeck (2012) exploits variation in soil erosion during the 1930's American Dust Bowl to show that negative environmental shocks can have substantial and lasting effects on productivity. Using data from a more recent period, we examine responsiveness to a slower-moving environmental "shock" that is very representative of what future climate change will likely bring. Similar to Hornbeck (2012), we find limited evidence that agricultural productivity has adapted to these environmental changes, with fairly negative implications for the future impacts of climate change on the agricultural sector.

The remainder of this paper is organized as follows. In Section 2 we develop a simple model of farmer adaptation and use it to motivate our empirical approach. Section 3 describes our main results on the extent of past adaptation, and Section 4 attempts to interpret the lack of adaptation that we observe. Section 5 uses data from global climate models to build projections of future yield impacts, and Section 6 concludes and discusses implications for policy.

2 Model and Empirical Approach

Agriculture is a key sector where future climate change is estimated to have large detrimental effects, and is a primary focus of the empirical literature on climate change impacts. To formalize the ways in which our identification of climate impacts differs from that of the past literature, we develop a simple model of farmer adaptation, building on earlier work by Kelly, Kolstad, and Mitchell (2005). The climate literature generally understands adaptation as any adjustment to a changing environment that exploits beneficial opportunities or moderates negative impacts.⁴ Adaptation thus requires an agent to recognize that something in her environment has changed, to believe that an alternative course of action is now preferable to her current course, and to have the capability to implement that alternative course.

We consider a farmer facing a choice about which of two crop varieties to grow, where one performs relatively better in cooler climates (variety 1) and the other in warmer climates (variety 2). We assume this relative performance is known to the farmer. Denote the choice of variety for farmer *i* as $x_{it} \in \{0, 1\}$, with $x_{it} = 1$ the choice to grow the relatively heattolerant variety 2. The output of farmer *i* in period *t* is $y_{it} = f(x_{it}, z_{it})$, where z_{it} is realized temperature in period *t* and is drawn from a normal distribution with mean ω_t and variance σ^2 . We assume a quadratic overall production technology with respect to temperature:

$$y_{it} = \beta_0 + \beta_1 z_{it} + \beta_2 z_{it}^2 + x_{it} (\alpha_0 + \alpha_1 z_{it} + \alpha_2 z_{it}^2)$$
(1)

 $^{^4\}mathrm{See}$ Zilberman, Zhao, and Heiman (2012) and Burke and Lobell (2010) for an overview.

with production for the conventional variety given by $\beta_0 + \beta_1 z_{it} + \beta_2 z_{it}^2$, and the differential productivity between the conventional and heat-tolerant varieties given by $\alpha_0 + \alpha_1 z_{it} + \alpha_2 z_{it}^2$.

The farmer in year *i* chooses x_{it} to maximize expected output prior to realizing weather. The heat-tolerant crop will be chosen if $E(\alpha_0 + \alpha_1 z_{it} + \alpha_2 z_{it}^2) > 0$, which can be rewritten as

$$\alpha_0 + \alpha_1 \omega_t + \alpha_2 (\omega_t^2 + \sigma^2) > 0.$$
⁽²⁾

We assume that the α and β parameters are known to the farmer but not to the econometrician. Figure 2 displays the productivity of the two varieties as a function of temperature. As drawn, the productivity frontiers have similar concavity⁵ ($\alpha_2 \approx 0$) such that the perfectly informed farmer adopts the heat-tolerant crop when the expected temperature exceeds $\tilde{\omega}$.

We incorporate climate change as a shift in mean temperature from $\omega \to \omega'$, with $\omega < \tilde{\omega} < \omega'$. In keeping with evidence from climate science (see Meehl et al. (2007)), we assume that this increase in mean is not accompanied by a change in variance, such that after climate change the farmer experiences $z_{it} \sim N(\omega', \sigma^2)$ in each year. A fully informed farmer recognizes this change and immediately adopts the heat-tolerant crop, which we consider "adaptation". In reality, farmers likely learn about changes in climate over time and only adjust behavior after acquiring strong enough information that climate has changed. Following Kelly, Kolstad, and Mitchell (2005), we assume this learning follows a simple Bayesian process where the farmer has an prior belief about ω_t but knows that this belief is imperfect. We denote the belief as μ_t and its variance as $1/\tau_t$, such that in period t the farmer believes $\omega \sim N(\mu_t, 1/\tau_t)$. In each period she observes z_{it} and updates her belief about the average temperature to μ_{t+1} using a weighted combination of her prior belief about mean climate realization she experiences. Letting $\rho = 1/\sigma^2$, the farmer's belief about mean climate after T years is given by (DeGroot, 1970):

$$\mu_T = \frac{\tau_t \mu_t + T\rho z_{it}}{\tau_t + T\rho} \tag{3}$$

With $\tau_{t+1} = \tau_t + \rho$, then in expectation it follows that:

$$\mu_T - \omega' = \frac{\tau_0(\mu_0 - \omega')}{\tau_0 + T\rho} \tag{4}$$

⁵A negative value of α_2 would indicate that productivity of the heat-tolerant crop is more responsive to temperature changes (i.e. the productivity or profit frontier for the heat-tolerant crop is "more concave"). In this case, if climate variability is large, then the expected gain from adaptation at average climate must be large enough to offset expected losses in bad years. With $\alpha_2 > 0$, the response function for the heat tolerant crop is "flatter" such that the farmer is willing to adopt the heat tolerant crop before the intersection of the two curves because of the increased certainty that it provides.

Equation (4) has two important implications: beliefs about mean temperature converge to the true value as the number of time periods increases $(T \uparrow)$, and converge more quickly when there is less variance in annual temperature (i.e. when ρ is larger). This suggests that farmers should be more likely to recognize changes in climate – and thus adapt to those changes, if information is a constraint to adaptation – in areas where the temperature variance is low, and when they are given more time to observe realizations of the new climate. We use these predictions to help us interpret our main findings in what follows.

2.1 Existing approaches

Returning to Figure 2, the long-term damages imposed by a shift in climate will be $v_0 - v_1$ if adaptation takes place.⁶ Past literature has taken two approaches to estimating this quantity. In pioneering work, Mendelsohn, Nordhaus, and Shaw (1994) use cross-sectional variation in average temperature and precipitation (and their squares) to explain variation in agricultural outcomes across US counties. The cross sectional specification is

$$y_i = \alpha + \beta_1 w_i + \beta_2 w_i^2 + c_i + \varepsilon_i, \tag{5}$$

where y_i is some outcome of interest in county *i*, w_i is again the average temperature, and c_i other time invariant factors affecting outcomes (such as soil quality). Mendelsohn et al's preferred dependent variable is land values, which represent the present discounted value of the future stream of profits that could be generated with a given parcel of land, and thus in principle embody any possible long-run adaptation to average climate. Therefore, a county with average temperature of ω will achieve v_0 on average, a county with average temperature of ω' will achieve v_1 , and the estimates of β_1 and β_2 along with a projected rise in average temperatures from ω to ω' would seem to identify the desired quantity of $v_0 - v_1$.

Cross sectional models in this setting make an oft-criticized assumption: that average climate is not correlated with other unobserved factors (the c_i – soil quality, labor productivity, technology availability etc) that also affect outcomes of interest (Schlenker, Hanemann, and Fisher, 2005; Deschênes and Greenstone, 2007). Given these omitted variables concerns, more recent work has used panel data to explore the relationship between agricultural outcomes and variation in temperature and precipitation (Deschênes and Greenstone (2007); Schlenker and Roberts (2009); Welch et al. (2010); Lobell, Schlenker, and Costa-Roberts

 $^{^{6}}$ Kelly, Kolstad, and Mitchell (2005) call this the "equilibrium response", in contrast to the costs incurred when undertaking adaptation (e.g. the purchase of a more expensive heat-tolerant variety), which they term "adjustment costs".

(2011)).⁷ The data generating process in this approach is:

$$y_{it} = \alpha + \beta_1 z_{it} + \beta_2 z_{it}^2 + c_i + \varepsilon_{it} \tag{6}$$

All time invariant factors are absorbed by the location fixed effects c_i , and impacts of temperature and precipitation on (typically annual) outcomes are thus identified from deviations from location-specific means.⁸ Because this year-to-year variation in temperature and precipitation (typically termed "weather") is plausibly exogenous, fixed effects regressions overcome omitted variables concerns with cross-sectional models, and the effect of temperature on outcomes such as yield or profits can be interpreted causally.

Many studies then combine the estimated short-run responses from panel regressions with output from global climate models to project potential impacts under future climate change.⁹ In making these projections, the implicit assumption is again that short-run responses to variation in weather are representative of how farmers will respond to longer-run changes in average climate. It is not obvious this will be the case. Consider a panel covering many years, with a temperature rise from ω to ω' occurring somewhere within these years. The panel model would identify movement along either one of the two curves shown in Figure 2, with the point estimate being a weighted average of the slopes of the two curves, with weights depending on if and when the varietal switch occurred. If the heat-tolerant crop is adopted at the end of the period then fixed effects estimates will be heavily weighted towards the curve for the conventional crop, overstating equilibrium losses. If adaptation is instantaneous then fixed effects estimates trace out the curve for the heat-tolerant crop, which could understate impacts if (as drawn) the slope of the response function is positive at ω' . Thus estimates of short-run responses to weather will not even bound estimates of longer-run response to climate. Panel models therefore solve identification problems in the cross-sectional approach, at the cost of more poorly approximating the idealized climate

⁷Examples in the climate literature outside of agriculture include Burke et al. (2009); Deschênes and Greenstone (2011); Auffhammer and Aroonruengsawat (2011); Dell, Jones, and Olken (2012).

⁸McIntosh and Schlenker (2006) show that including a quadratic term in the standard panel fixed effects model allows unit means to re-enter the estimation. Inclusion of a squared term therefore results in impacts of the independent variable of interest being derived not only from within-unit variation over time but also from between-unit variation in means. In principle, this would allow for estimation of the outer as well as the inner envelope, a strategy explored by Schlenker (2006), although it is not clear that omitted variables concerns have not also re-entered the estimation along with the unit means. In any case, growing degree days allow temperature to enter non-linearly without the complication of the quadratic term, and we exploit this fact to generate estimates of adaptation. Furthermore, using trends in climate to identify climate sensitivities remains an arguably more "direct" approach to understanding near-term impacts of future climate change, and is thus the approach we take here.

 $^{^{9}}$ See Burke et al. (2013) for a review of these studies and for the use of global climate models in this context.

change experiment.

2.2 The long differences approach

We attempt to simultaneously overcome the limitations of both the cross-sectional and panel approaches by long differencing. We construct longer run yield and temperature averages at two different points in time for a given location, and calculate changes in average yields as a function of changes in average temperature. Consider two multi-year periods denoted "a" and "b", each spanning n years. Our approach is to separately sum over all the years in each period, e.g. with the average yield in period a given by $\overline{y_{ia}} = \frac{1}{n} \sum_{t \in a} y_{it}$ and average temperature $\overline{z_{ia}}$ representing the averaged z_{it} 's over the same period. The resulting equation for period a is:

$$\overline{y_{ia}} = \alpha + \beta_1 \overline{z_{ia}} + \beta_2 \overline{z_{ia}}^2 + c_i + \overline{\varepsilon_{ia}}.$$
(7)

Defining period b similarly, we can "long difference" over the two periods to get:

$$\overline{y_{ib}} - \overline{y_{ia}} = \beta_1 (\overline{z_{ib}} - \overline{z_{ia}}) + \beta_2 (\overline{z_{ib}}^2 - \overline{z_{ia}}^2) + (c_i - c_i) + (\overline{\varepsilon_{ib}} - \overline{\varepsilon_{ia}})$$
(8)

The time-invariant factors drop out, and we can rewrite as:

$$\Delta \overline{y_i} = \beta_1 \Delta \overline{z_i} + \beta_2 \Delta (\overline{z_i})^2 + \Delta \varepsilon_i, \tag{9}$$

Generating unbiased estimates of β_1 and β_2 requires that changes in temperature between the two periods are not correlated with time-varying unobservables that also affect outcomes of interest. Below we provide evidence that differential climate trends across our sample of US counties are likely exogenous and surprisingly large.

Estimating the impact of climate on agricultural productivity with the long differences approach in (9) offers substantial advantages over both the cross-sectional and panel approaches. First, it arguably better approximates the ideal "parallel worlds" experiment. That experiment randomly assigns climate trends to different earths, and the long differences approximation utilizes variation in longer-run climate change that are unlikely to be correlated with variables that explain changes in yield. Second, unlike the cross-sectional approach, the long differences estimates are immune to time-invariant omitted variables, and unlike the panel approach the relationship between climate and agricultural productivity is estimated from long-term changes in average conditions instead of short-run year-to-year variation. Finally, because long differences estimates will embody any adaptations that farmers have undertaken to recent trends, and because the range in these trends falls within the range of projected climate change over at least the next three decades, then projections of future climate change impacts on agricultural productivity based on long differences estimates would appear more trustworthy than those based on either panel or cross-sectional methods.

We then use this strategy to quantify the extent of recent adaptation in US agriculture, comparing our long differences estimates to those from an annual panel model. We would interpret more positive long difference estimates as evidence of adaptation: that farmers are better able to adjust to longer-run changes in climate than they are to shorter-run changes in weather. In Figure 2, if any adaptation takes place, the long differences approach should identify $v_0 - v_1$. If no adaptation occurs, then long difference regressions will identify $v_0 - v_2$, i.e. the same damages identified by fixed effects. We attempt to rule out other explanations for divergence between panel and long-differences estimates – e.g. measurement error, or adaptation outside of agriculture – in Section 4.

2.3 Data and estimation

Our agricultural data come from the United States Department of Agriculture's National Agricultural Statistics Service. Crop area and yield data are available at the county-year level, and economic measures of productivity such as total revenues and agricultural land values are available every five years when the Agricultural Census is conducted.¹⁰ Our unit of observation is thus the county, and in keeping with the literature we focus the main part of the analysis on counties that are east of the 100th meridian. The reason for this is that cropland in the American West typically relies on highly subsidized irrigation systems, and the degree of adaptation embodied in the use and expansion of these systems might poorly extrapolate to future scenarios as the federal government is unlikely to subsidize new water projects as extensively as it has in the past (Schlenker, Hanemann, and Fisher, 2005). Over the last decade, the counties east of the 100th meridian accounted for 93% of US corn production and 99% of US soy production.

Our climate data are drawn from Schlenker and Roberts (2009) and consist of daily interpolated values of precipitation totals and maximum and minimum temperatures for 4 km grid cells covering the entire United States over the period 1950-2005. These data are aggregated to the county-day level by averaging daily values over the grid cells in each county where crops are grown, as estimated from satellite data.¹¹

Past literature has demonstrated strong non-linearities in the relationship between temperature and agricultural outcomes (e.g. Schlenker and Roberts (2009)). Such non-linearities are generally captured using the concept of growing degree days (GDD), which measure the

¹⁰We thank Michael Roberts for sharing additional census data that are not yet archived online.

¹¹We thank Wolfram Schlenker for sharing the weather data and the code to process them.

amount of time a crop is exposed to temperatures between a given lower and upper bound. Following Schlenker and Roberts (2009), we use the within-day distribution of temperatures to calculate the percent of each day that cropped area in each county is exposed to temperatures between given lower and upper bounds, and then sum these daily exposures over a fixed growing season (April 1 to September 30th) to get a measure of annual growing degree days for those bounds.

Using this notion of GDD, and using the county agricultural data described above, we model agricultural outcomes as a simple piecewise linear function of temperature and precipitation.¹² We estimate the long differences model:

$$\Delta y_{is} = \beta_1 \Delta GDD_{is;l_0:l_1} + \beta_2 \Delta GDD_{is;l_1:\infty} + \beta_3 \Delta Prec_{is;p < p_0} + \beta_4 \Delta Prec_{is;p > p_0} + \alpha_s + \Delta \varepsilon_{is}, \quad (10)$$

where Δy_{is} is the change in some outcome y in county i in state s between two periods. In our main specification these two periods are 1980 and 2000, and we calculate endpoints as 5-year averages to more effectively capture the change in average climate or outcomes over time. That is, for the 1980-2000 period we take averages for each variable over 1978-1982 and over 1998-2002, and difference these two averages.

The lower temperature "piece" in (10) is the sum of GDD between the bounds l_0 and l_1 , and $\Delta GDD_{i;l_0:l_1}$ term gives the change in GDD between these bounds over the two periods. The upper temperature "piece" has a lower bound of l_1 and is unbounded at the upper end, and the $\Delta GDD_{i;l_1:\infty}$ term measures the change in these GDD between the two periods.¹³ We also measure precipitation in a county as a piecewise linear function with a kink at p_0 . The variable $Prec_{i;p<p_0}$ is therefore the difference between precipitation and p_0 interacted with an indicator variable for precipitation being below the threshold p_0 . $Prec_{i;p>p_0}$ is similarly defined for precipitation above the threshold.¹⁴ In the estimation we set $l_0 = 0$ and allow the data to determine l_1 and p_0 by looping over all possible thresholds and selecting the model

¹²We choose the piecewise linear approach for two reasons. First, existing work on US agricultural response to climate suggests that a simple piecewise linear function delivers results very similar to those estimated with much more complicated functional forms (Schlenker and Roberts, 2009). Second, these other functional forms typically feature higher order terms, which in a panel setting means that unit-specific means re-enter the estimation (McIntosh and Schlenker, 2006). This not only raises omitted variables concerns, but it complicates our strategy for estimating the extent of past adaptation by comparing long differences with panel estimates; in essence, identification in the panel models in no longer limited to location-specific variation over time.

¹³As an example, if $l_0 = 0$ and $l_1 = 30$, then a given set of daily average temperatures of -1, 0, 1, 10, 29, 31, and 35 would result in $GDD_{is;l_0:l_1}$ equal to 0, 0, 1, 10, 29, 30, and 30, and $GDD_{is;l_1:\infty}$ equal to 0, 0, 0, 0, 1, and 5. In practice we use the within-day distribution of temperatures to allow fractions of days to be spent above or below a given threshold, but the principle is the same.

¹⁴A simple example is useful to illustrate the differencing of precipitation variables when the threshold is crossed between periods. Consider a county with an increase in average precipitation from 35 mm in 1980 to 50 mm in 2000. If the precipitation threshold is 40 mm, then $\Delta Prec_{is;p>p_0} = 5$ and $\Delta Prec_{is;p>p_0} = 10$.

with the lowest sum of squared residuals.

Importantly, we also include in (10) a state fixed effect α_s which controls for any unobserved state-level trends. This means that identification comes only from within-state variation, eliminating any concerns of time-trending unobservables at the state level. Finally, to quantify the extent of recent adaptation, we estimate a panel version of (10), where observations are at the county-year level and the regression includes county and year fixed effects. As suggested by earlier studies (e.g. Schlenker and Roberts (2009)), the key coefficient in both models is likely to be β_2 , which measures how corn yields are affected by exposure to extreme heat. If farmers adapt significantly to climate change then we would expect the coefficient β_2 to be significantly larger in absolute value when estimated with panel fixed effects as compared to our long differences approach. The value $1 - \beta_2^{LD}/\beta_2^{FE}$, gives the percentage of the negative short-run impact that is offset in the longer run, and is our measure of adaptation to extreme heat.

There are two potential concerns with our empirical approach. The first is that the inclusion of state fixed effects could absorb most of the variation of interest in our temperature variables. Second, our supposed differential trends in temperature across counties could just be driven by short-run variation in weather around the chosen endpoint years. This has two potential implications. First, there could have been little "true" long run change in temperature to adapt to. Second, even if temperature was trending differentially across counties, our long differences estimator might mechanically deliver estimates that are similar to those from a county-year panel if inter-annual variation in temperature around this trend (i.e. weather) was large. In Section A.1, we explore each of these issues in more detail. We first demonstrate that the residual variation in our temperature changes of interest remains large (even relative to projected future changes) after accounting for state fixed effects, and that this variation very likely represents true long-run increases in temperature rather than large variation in chosen endpoint years. Second, we clarify under what conditions the panel and long differences approaches will differentially identify "true" long-run and short-run responses to changes in temperature, and demonstrate via simulation that our long differences estimates are likely based on agents' responses to longer-run changes in temperature.

Figure 1 displays the variation that is used in our identification strategy. Some US counties have cooled slightly over the past 3 decades, while others have experienced warming equivalent to over 1.5 times the standard deviation of local temperature – roughly equivalent to the mean warming projected by global climate models to occur over US corn area by 2030. Differential trends in precipitation over the 1980-2000 have been similarly large, with precipitation decreasing by more than 30% in some counties and increasing by 30% in others

- a range that again almost fully contains the range in climate model projections of future precipitation change over the same area by the mid-21st century. Substantial variation is apparent even within states. For instance, Lee County in the southeastern Iowa experienced an increase in average daily temperature during the main corn growing season of 0.46°C, and Mahaska county – approximately 80 miles to the northwest – experienced a decrease in temperature of 0.3°C over the same period. Corn yields in parts of northern Kentucky declined slightly while rising by 20-30% only 100 miles to the south.

Importantly, there remains large variation in our main regressor of interest (exposure to extreme heat) even after conditioning on other climate variables and state fixed effects (see Table A.1), and as shown in Figure A.3, this variation substantially overlaps projections from global climate models of future changes in extreme heat exposure. This means that applying our estimates to predictions from climate models is not asking our model to extrapolate far out of sample.

While we explore robustness of our results to different time periods and differencing lengths, we focus on the post-1980 period for a number of reasons. First, warming trends since 1980 were much larger than in earlier periods. For instance, over the 1960-1980 period, only half of the counties in our sample experienced average warming, and none experience warming of more than 1C (see Figure A.10). Second, recognition of climate change was much higher in this later period, which helps alleviate some concerns that a lack of recognition of climate change is what is driving our results. In particular, prior to 1980 there was even significant scientific and popular concern about the risks from "global cooling" (e.g. Gwynne (1975)), and only during the 1980's and 1990's was there growing recognition that the climate was warming and that increasing greenhouse gas emissions meant there would very likely be further warming in the future.

2.4 Are recent climate trends exogenous?

There are a few potential violations to the identifying assumption in (10). The first is that trends in local emissions could affect both climate and agricultural outcomes. In particular, although greenhouse gases such as carbon dioxide typically become "well mixed" in the atmosphere soon after they are emitted, other species such as aerosols are taken out of the atmosphere by precipitation on a time scale of days, meaning that any effect they have will be local. Aerosols both decrease the amount of incoming solar radiation, which cools surface temperatures and lowers soil evaporation, and they tend to increase cloud formation, although it is somewhat ambiguous whether this leads to an increase in precipitation. For instance, Leibensperger et al. (2011) found that peak aerosol emissions in the US during the 1970s and 1980s reduced surface temperatures over the central and Mid-Atlantic US by up to 1° C, and led to modest increases in precipitation over the same region.

The effect of aerosols on crops is less well understood (Auffhammer, Ramanathan, and Vincent, 2006). While any indirect effect through temperature or precipitation will already be picked up in the data, aerosols become an omitted variables concern if their other influence on crops – namely their effect on solar radiation – have important effects on crop productivity. Because crop productivity is generally thought to be increasing and concave in solar radiation, reductions in solar radiation are likely to be harmful, particularly to C_4 photosynthesis plants like corn that do not become light saturated under typical conditions.¹⁵ However, aerosols also increase the "diffuse" portion of light (think of the relatively even light on a cloudy day), which allows additional light to reach below the canopy, increasing productivity. A recent modeling study finds negative net effects for corn, with aerosol concentrations (circa the year 2000) reducing corn yields over the US midwest by about 10%, albeit with relatively large error bars. This would make it likely that, if anything, aerosols will cause us to understate any negative effect of warming on crop yields: aerosols lead to both cooling (which is generally beneficial in our sample) and to a reduction in solar radiation (which on net appears harmful for corn). In any case, the inclusion of state fixed effects means that we would need significant within-state variation in aerosol emissions for this to be a concern.

The second main omitted variables concern is changes in local land use. Evidence from the physical sciences suggests that conversion between types of land (e.g. conversion of forest to pasture, or pasture to cropland), or changes in management practices within preexisting farmland (e.g. expansion of irrigation) can have significant effects on local climate. For instance, expansion in irrigation has been shown to cause local cooling (Lobell, Bala, and Duffy, 2006), which would increase yields both directly (by reducing water stress) and indirectly (via cooling), leading to a potential omitted variables problem. The main empirical difficulty is that local land use change could also be an adaptation to changing climate – i.e. a consequence of a changing climate as well as a cause. In the case of irrigation, adaptation and irrigation-induced climate change are likely to go in opposite directions: if irrigation is an omitted variable problem, we would need to see greater irrigation expansion in cooler areas, whereas if irrigation is an adaptation, we would expect relatively more expansion in warm areas. Overall, though, because we see little change in either land area or land management practices, we believe these omitted variables concerns to be limited as well.

¹⁵Crops that photosynthesize via the C_3 pathway, which include wheat, rice, and soybeans, become "light saturated" at one-third to one-half of natural sunlight, meaning that reductions in solar radiation above that threshold would have minimal effects on productivity. C_4 plants such as corn do not light saturate under normal sunlight, so are immediately harmed by reductions in solar radiation (Greenwald et al. (2006)).

The most recent evidence from the physical sciences suggests that the large differential warming trends observed over the US over the past few decades are likely due to natural climate variability - in particular, to variation in ocean temperatures and their consequent effect on climate over land through increased localized precipitation (which leads to local cooling) or through cold air flowing in from the north (Meehl, Arblaster, and Branstator, 2012), effects which need not be homogenous within US states. As such, these trends appear to represent a true "natural experiment", and are likely exogenous with respect to the outcomes we wish to measure. Nevertheless, as a final check on exogeneity, we show in Table A.2 that the within-state change in exposure to extreme heat during the 1980-2000 period are not strongly correlated with several county-level covariates.

3 Empirical Results

Our primary analysis focuses on the effect of longer-run changes in climate on the productivity of corn and soy, the two most important crops in the US in terms of both area sown and production value. The yield (production per acre) of these two crops is the most basic measure of agricultural productivity, and is well measured annually at the county level. However, because a focus on yields alone will not cover the full suite of adaptations that farmers might have employed, we then examine adjustments along other possible margins.

3.1 Corn productivity

The results from our main specifications for corn yields are given in Table 1 and shown graphically in Figure 3. In our piecewise linear approach, productivity is expected to increase linearly up to an endogenous threshold and then decrease linearly above that threshold, and the long differences and panel models reassuringly deliver very similar temperature thresholds (29°C and 28°C, respectively) and precipitation thresholds (42cm and 50cm). In Columns 1-4 we run both models under the thresholds selected by the long differences, and in Columns 5-8 we fix thresholds at values chosen by the panel.

The panel and long differences models deliver very similar estimates of the responsiveness of corn yields to temperature. Exposure to GDD below 29°C (row 1) have small and generally insignificant effects on yields, but increases in exposure of corn to temperatures above 29° result in sharp declines in yields, as is seen in the second row of the table and in Figure 3. In our most conservative specification with state fixed effects, exposure to each additional degree-day of heat above 29°C results in a decrease in overall corn yield of 0.44%.¹⁶ The

¹⁶While we prefer the more conservative specifications with state fixed effects in Table 1, one concern with

panel model delivers a slightly more negative point estimate, a -0.56% yield decline for every one degree increase above 29°C, but (as quantified below) we cannot reject that the estimates are the same. We obtain similar results when the two models are run under the temperature and precipitation thresholds chosen by the panel model (Columns 5-8), and similar results when the panel model is estimated with state-by-year fixed effects rather than year fixed effects.

The estimates of the effects of precipitation on corn productivity are somewhat more variable. The piecewise linear approach selected precipitation thresholds at 42 cm (long differences) or 50 cm (panel), but most of the variation in precipitation is at values above 42 cm - e.g. the 10th percentile of annual county precipitation is 41.3 cm. Long differences point estimates suggest an approximate increase in yields of 0.33% for each additional centimeter of rainfall above 42 cm, which are of the opposite sign and substantially larger than panel estimates. Nevertheless, we note that even the long differences precipitation estimates remain quite small relative to temperature effects: on a growing season precipitation sample mean of 57cm, a 20% decrease (roughly the most negative climate model projection for US corn area by the end of century) would reduce overall yields by less than 4%. Given that precipitation is likely measured with greater error than temperature, we cannot rule out that our results understate the role of precipitation changes in corn yields (Lobell, 2013). But as we show in Section 5, and consistent with other recent findings (Schlenker and Roberts, 2009; Schlenker and Lobell, 2010), any future impacts of climate change via changes in precipitation are likely to be dominated by changes in yields induced by increased exposure to extreme heat, even if precipitation is measured with some error.

To test robustness of the corn results, we show in the remainder of this subsection that our results are relatively insensitive to the choice of endpoint years, to the number of years used to calculate endpoints, and to an alternate estimation strategy which further weakens our identification assumptions. In Appendix A.3, we provide further evidence that our results are insensitive to the exclusion of yield and temperature outliers, and to the inclusion of baseline covariates in the regression.

We first show that our results are largely unchanged when we change the time period under study. In particular, we estimate Equation (10) varying T_0 from 1955 to 1995 in 5

the inclusion of state fixed effects is that farmer responses to increasing temperature might vary meaningfully at the state level, for instance if governments in states that experienced substantial warming helped their farmers invest in adaptation measures. These policies would be absorbed by the state fixed effects, and could obscure meaningful adaptation measures undertaken by farmers. Our results appear inconsistent with a story of state-level adaptation to extreme heat. The effects of extreme heat in the specifications without state fixed effects (columns 1 and 4) are substantially more negative than in the comparable specifications with state fixed effects, which is the opposite of what would be expected if state-level adaptation policies were an omitted variable in these regressions.

year increments, and for each value of T_0 we estimate 5, 10, 15, 20, 25, and 30 year difference models.¹⁷ Results are shown graphically in Figure 4. We display the difference between the estimate of β_2 for 1980-2000 (our baseline estimate) and the estimate of β_2 for the period determined by the starting year and differencing length. The 95% confidence intervals of the differences are calculated by bootstrapping.¹⁸ The average estimate of β_2 across these 39 models is -0.0058, with only 8 of the estimates of β_2 being statistically different from our main 1980-2000 estimate and *none* statistically different in the positive direction. This suggests if anything that our baseline point estimate on the effect of extreme heat is conservative.¹⁹ We conduct an analogous exercise for the panel model to make sure that the effect of extreme heat in the panel does not vary with the chosen time period. Results are plotted in Figure A.11, and agree with earlier findings in Schlenker and Roberts (2009) that the effects of inter-annual deviations in extreme heat have not declined significantly over time.²⁰

Section 2.3 provided initial evidence that our "long-run" differences over time reflect substantial longer-run changes in climate rather than large short-run variation around the endpoint years. To provide additional evidence that this is true, we re-construct our long differences with endpoints averaged over 10 years rather than 5, which should help average out idiosyncratic noise. As a further test, we utilize the entire 1950-2005 sample, split it into 28-year periods (1950-1977 and 1978-2005), average yield and climate within each period, and then difference the period and perform our long differences estimation. We vary the sample to include any county growing corn in either period, or all counties growing corn in either period (or something in-between). As shown in Table A.6, the effect of extreme heat is large, negative, and highly significant across all specifications, and these results again

¹⁷Some models of course could not be estimated since our data end at 2005, meaning our 5-year smoothed estimates are only available through 2003. In each model we limit the sample to the set of counties from Table 1. Each regression is weighted by 5 year average corn area during the starting year. The temperature and precipitation thresholds are fixed at 29° and 42 cm across models.

¹⁸We drew 1000 samples of 31 states with replacement and estimated all regressions for each sample. The differences between the 1980-2000 estimate and all other possible estimates were calculated for each sample. The bootstrapped standard errors are the standard deviations of the differences in estimates.

¹⁹In Appendix Figure A.9, we display the raw coefficients and their confidence intervals for each period: all estimates are negative, and in only 8 out of 39 cases to we fail to reject a significant negative effect of extreme heat on corn productivity.

²⁰While 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. As emphasized in our conceptual framework, 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.

suggest if anything that our baseline results are conservative.

Finally, our estimates in Equation (9) would be biased in the presence of within-state time-varying unobservables correlated with both climate and yields. To address this possibility, we use our many decades of data to construct a two period panel of long differences, which further weakens our identification assumption. We estimate the following model:

$$\Delta y_{it} = \beta_1 \Delta GDD_{it;l_0:l_1} + \beta_2 \Delta GDD_{it;l_1:\infty} + \beta_3 \Delta Prec_{it;p < p_0} + \beta_4 \Delta Prec_{it;p > p_0} + \alpha_i + \delta_t + \varepsilon_{it}, \quad (11)$$

where all variables are measured in 20 year differences with t indicating the time period over which the difference is taken. Unobserved differences in average county-level trends are accounted for by the α_i , and δ_t accounts for any common trends across counties within a given period. The β 's are now identified off *within-county* differences in climate changes over time, after having accounted for any differences in trends common to all counties. An omitted variable in this setting would need to be a county-level variable whose trend over time differs across the two periods in a way correlated with the county-level difference in climate changes across the two periods, and it is difficult to construct stories for omitted variables that meet these conditions.

In Table 2 we report estimates from both the 1955-1995 period and the 1960-2000 period. In all models the effect of temperature above 29° remains negative and significant even after the inclusion of county fixed effects. The main coefficients for GDD>29 are also similar to our baseline estimates in Table 1. The main long differences estimates are therefore robust to controlling for a richer set of county-specific time-varying factors.

3.2 Adaptation in corn

Comparing panel and long differences coefficients provides an estimate of recent adaptation to temperature and precipitation changes, with $1 - \beta_2^{LD}/\beta_2^{FE}$ giving the share of the shortrun impacts of extreme heat that are offset in the longer run. Point estimates from Table 1 suggest that 22-23% of short-term yield losses from exposure to extreme heat have been alleviated through longer run adaptations. To quantify the uncertainty in this adaptation estimate, we bootstrap our data 1000 times (sampling U.S. states with replacement to account for spatial correlation) and recalculate $1 - \beta_2^{LD}/\beta_2^{FE}$ for each iteration.²¹ We run this procedure for the 1980-2000 period reported in Table 1, and repeat it for the each of the 20, 25, and 30-year intervals shown in Figure 4 that start in 1970 or later. The distribution of

 $^{^{21}}$ That is, we take a draw of states with replacement, estimate both long differences and panel model for those states, compute the ratio of extreme heat coefficients between the two models, save this ratio, and repeat 1000 times for a given time period.

bootstrapped adaptation estimates then allow us to test, for each time period of interest, the null hypothesis of "no adaptation" to extreme heat – i.e. that $1 - \beta_2^{LD} / \beta_2^{FE} = 0$.

Results are shown in Figure 5 and suggest that, on the whole, longer-run adaptation to extreme heat in corn has been limited. Median estimates from each distribution all indicate that adaption has offset less than 25% of short run impacts – and point estimates are actually slightly *negative* in two-thirds of the cases. In almost all cases we can conclude that adaptation has offset at most half of the negative shorter-run impacts of extreme heat on corn yields. Finally, all confidence intervals span zero, meaning we can never reject that there has been no more adaptation to extreme heat in the long run than has been in the short run (one-sided p-values on the test of the null against the alternative hypothesis that $1 - \beta_2^{LD}/\beta_2^{FE} > 0$ are p=0.12 or greater, as shown on the left of the figure).

3.3 Soy productivity

All of our analysis up to this point has focused on corn, the dominant field crop in the US by both area and value. It is possible, however, that the set of available adaptations differs by crop and there could be additional scope for adaptation with other crops. Soy is the country's second most important crop in terms of both land area and value of output. In Figure A.13 we show the various estimates of the effect of extreme heat on log soy yields as derived from the long differences model. The horizontal line in each panel is the 1978-2002 panel estimate of β_2 for soy which is -0.0047, almost identical to the corn estimate. The thresholds for temperature and precipitation are 29° and 50 cm, which are those that produce the best fit for the panel model. While the soy results are somewhat noisier than the corn results, the average response to extreme heat across the 39 estimates is -0.0032, giving us a point estimate of longer run adaptation to extreme heat of about 32%. 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, there appears to have been limited adaptation to extreme temperatures amongst soy farmers.

4 Alternate Explanations

Results so far suggest that corn and soy farmers are no more able to deal with increased extreme heat exposure over the long run than they are in the short run. We now explore the extent to which this limited apparent adaptation we observe in crop yields is due to (i) measurement error, (ii) selection into or out of agriculture, (iii) adaptation along other margins, (iv) disincentives induced by existing US government policy, (v) and/or a lack of recognition that climate is changing. Evidence in favor of the first two hypotheses would challenge the validity of results; evidence in favor of any of the last three would alter their interpretation, and could make our long difference estimates a potentially poor basis for projecting future impacts if policies or information were to change.

4.1 Measurement error

A key concern with fixed effects estimates of the impact of climate variation is attenuation bias caused by measurement error in climate variables. Fixed effects estimates are particularly susceptible to attenuation since they rely on short-term deviations from average climate to identify coefficients. This makes it more difficult to separate noise from true variation in temperature and precipitation compared to a setting where identification comes from relatively better-measured averages over space or time (such as in our long differences results). Therefore one explanation for the limited observed yield adaptation is simply that panel estimates are attenuated relative to long differences estimates, and thus that that comparing the two estimates will mechanically understate any adaptation that has occurred.

We first note that because temperature and precipitation are generally negatively correlated, measurement error in *both* climate variables is likely to partially offset the attenuation caused by mis-measurement of temperature (Bound, Brown, and Mathiowetz, 2001). With more rainfall helping yields and warmer temperatures harming them, classical measurement error in precipitation could bias the temperature effect *away* from zero: the negative correlation between temperature and rainfall results in warmer years having artificially low yields due to attenuation in the precipitation variable. It is therefore not likely the case that the only effect of measurement error on the temperature coefficients is attenuation.²²

We also follow Griliches and Hausman (1986) and investigate the potential for large attenuation in our fixed effects estimates by comparing different panel estimators. If climate in a given county is highly correlated across time periods and measurement error is uncorrelated between successive time periods, then as Griliches and Hausman (1986) show, random effects estimates should be larger in absolute value than the fixed effects estimates which in turn should be larger than estimates using first differences. The intuition is that random effects estimates are identified using a combination of within and between variation and therefore are less prone to measurement error than fixed effects estimates and first differences which

 $^{^{22}}$ This result holds so long as the measurement error for temperature and precipitation is uncorrelated with the "true" temperature and precipitation values - i.e. that both exhibit classical measurement error but does not require the temperature and precipitation errors to be uncorrelated. We have verified this via simulation, with results available upon request.

rely entirely on within-county variation. Table A.7 shows that estimates from all three estimators are remarkably similar, providing suggestive evidence that measurement error is not responsible for the similarity between fixed effects and long differences estimates.

4.2 Selection

A second explanation for the observed lack of adaptation is a selection story in which better performing farmers exit agriculture in response to warming temperatures. This would leave the remaining population with lower average yields and thus create a mechanical negative relationship between warming temperatures and yields. Although the alternate selection story appears just as plausible – that better performing farmers are more able to maintain yields in the face of climate change, and the worse performers are the ones who exit – we can check in the data whether characteristics that are correlated with productivity also changed differentially between places that heated and those that did not. Table A.10 provides suggestive evidence that this is not the case. The percentage of farms owning more than \$20,000 equipment, which is positively correlated with productivity, is only weakly correlated to extreme heat exposure. While this cannot fully rule out selective exit from agriculture, it provides some evidence that selection is not driving our yield results.

4.3 Adaptation along other margins

A third explanation is that a focus on corn and soy yields, while capturing many of the off-mentioned modes of adaptation (e.g. switching seed varieties), might not capture all possible margins of adjustment available to farmers and thus could understate the extent of overall adaptation to climate change.

One way to capture broader economic adjustment to changes in climate is to explore climate impacts on farm revenues or profits, an approach adopted in some of the recent literature (e.g. Deschênes and Greenstone (2007)). There are at least two empirical challenges with using profits in particular. The first is that measures of revenues and expenditures are only available every 5 years when the US Agricultural Census is conducted. Given that our differencing approach seeks to capture change in average farm outcomes over time, if both revenues and costs respond to annual fluctuations in climate, then differencing two "snapshots" from particular years might provide a very noisy measure of the longer term change in profitability. A second concern is that available data on expenses do not measure all relevant costs (e.g. the value of own or family labor on the farm), which might further bias profit estimates if these expenses also respond to changes in climate. As shown in Appendix Section A.8, long differences regressions with such a measure of "profit" as the dependent variable are indeed very noisy, and we cannot reject that there is no effect on profits, and similarly cannot 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 impacts on economic profitability. The first is to construct an annual measure of revenue per acre, which we do by combining annual county-level yield data with annual data on state-level prices.²³ We then sum these revenues across the six major crops grown during the main Spring-Summer-Fall season in our sample counties: corn, soy, cotton, spring wheat, hay, and rice. This revenue measure will underestimate total revenue to the extent that not all contributing crops are included, but should capture any gain from switching among these primary crops in response to a changing climate. It will also capture any offsetting effect of price movements caused by yield declines, which while not an adaptation measure per se might reduce the need for other adaptation. Our second approach proceeds with the available expenses data from the Census to examine the impact of longer-run changes in climate on different input expenditures.²⁴

Table 3 shows results for our revenue measures. Consistent with some offsetting price movements, point estimates on how corn revenues per acre respond to extreme heat are slightly less negative than yield estimates under both panel and long differences models (Columns 1 and 2), but at least for the differences model we cannot reject that the coefficients are the same as the yield estimates. Revenues for the six main crops appear roughly equally sensitive to extreme heat in a panel and long differences setting (Columns 3 and 4), again suggesting that longer run adaptation has been minimal.²⁵ Furthermore, we show in Table A.8 that trends in climate have had minimal effects on expenditures on fertilizer, seed, chemical, and petroleum. We interpret this as further evidence that yield declines are

 $^{^{23}}$ Prices are only available at the state level and to our knowledge do not vary much within states within a given year.

²⁴We attempt to capture changes in *average* expenditures by averaging two census outcomes near each endpoint and then differencing these averaged values. For example, ag census data are available in 1978, 1982, 1987, 1992, 1997, and 2002. The change in fertilizer expenditures over the period are constructed as: Δ fertilizer expenditure₁₉₈₀₋₂₀₀₀ = (fert₁₉₉₇ + fert₂₀₀₂)/2 - (fert₁₉₇₈ + fert₁₉₈₂)/2

²⁵Coefficient estimates on the six-crop revenue measure are nevertheless about half the size of estimates for corn. We do not interpret this as evidence for adaptation for two reasons. First, panel and long differences estimates for how crop revenues respond to extreme heat are the same. Second, adaptationrelated explanations for why crop revenues should be less sensitive than corn revenue – e.g. farmers switch among crops to optimize revenues – would require that farmers are able to adjust their crop mix on an *annual* level *before* any extreme heat for that season is realized. This seems unlikely. We believe a more likely explanation is that we are more poorly measuring the climate variables and thresholds that are relevant to these other crops; regressions are run under the corn temperature and precipitation thresholds, and using data based on the corn growing season. If climate is measured with noise, then coefficient estimates will be attenuated.

not masking other adjustments that somehow reduce the economic losses associated with exposure to extreme heat.

To further explore whether our yield estimates hide beneficial switching out of corn and to other crops, we repeat our long differences estimation with changes in (log) corn area and changes in the percentage of total farmland planted to corn as dependent variables. Results are given in Table 4, and we focus on the sample of counties with extreme heat outliers trimmed.²⁶ There appears to have been minimal impact of increased exposure to extreme heat on total area planted to corn (Column 1), but we do find some evidence that the percentage of total farm area planted to corn declined in areas where extreme heat exposure grew. This effect appears small. In counties where increases in extreme heat were the most severe, observed increases in GDD above 29°C would have reduced the percentage of area planted to corn by roughly 3.5%.²⁷

A final adaptation available to farmers would be to exit agriculture altogether, an option that recent literature has suggested is a possibility. For instance, Hornbeck (2012) shows that population decline was the main margin of adjustment across the Great Plains after the American Dust Bowl. Feng, Oppenheimer, and Schlenker (2012) use weather as an instrument for yields to show that declines in agricultural productivity in more recent times result in more outmigration from rural areas of the Corn Belt. To quantify adaptation along this margin, we repeat our long differences estimation with total farm area, total number of farms, and county population as dependent variables. If there is a net reduction in the number of people farming due to increased exposure to extreme heat, we should see a decline in the number of farms; if this additional farmland is not purchased and farmed by remaining farmers, we should also see a decline in total farmland.

Results are in Columns 3-5 of Table 4. Point estimates of the effect of extreme heat on both (log) farm area and number of farms are negative but small and statistically insignificant. Nevertheless, the standard error on the number of farms measure is such that we

 $^{^{26}}$ As shown in Table A.5 - and unlike for our yield outcomes - a few outcomes in this table are altered fairly substantially when these five outliers (0.3% of the sample) are included. Given that these counties are all geographically distinct (along the Mexico border in southern Texas), and experienced up to 20 times the average increase in exposure to extreme heat than our median county in the sample, it seems reasonable to exclude them from the analysis.

²⁷We also explore whether corn planting dates, which are available at the state level for 14 states in our sample for the full study period, responded differentially in areas that warmed – which could be additional evidence of adaptation which we might miss if we've fixed the agricultural growing season from April 1 to Sept 30th. Using NASS data, we define the planting date as the week of the year in which 50% of that year's crop has been planted, and find that this date moved an average of 8 days earlier between 1980-2000 in our 14-state sample, consistent with overall warming allowing earlier planting. However, we estimate a slight *positive* but insignificant relationship between change in extreme heat (GDD>29) and planting date in the sample, providing little conclusive evidence that planting dates shifted in response to changes in extreme heat exposure. Results available upon request.

cannot rule out a 5-10% decline in the number of farms for the counties experiencing the greatest increase in exposure to extreme heat over our main sample period.²⁸ Point estimates on the response of population to extreme heat exposure are similar to estimates for number of farms, and again although estimates are not statistically significant we cannot rule out population declines of 5-10% for the counties that warmed the most. Taken together, and consistent with the recent literature, these results suggest that simply not farming may be an immediate adaptation to climate change for some farmers – although we have little to say on the welfare effects of such migration.

4.4 Policy disincentives to adapt

A fourth explanation for limited adaptation is that certain governmental agricultural support programs – subsidized crop insurance in particular – could have reduced farmers' incentives to adapt. In the crop insurance program, the federal government insures farmers against substantial losses while also paying most or all of their insurance premiums, and this plausibly could have reduced farmers' incentives to undertake costly adaptations.²⁹

As one check on whether observed lack of adaptation is being driven by the existence of subsidized insurance, we utilize the large-scale expansion of the federal crop insurance program in the mid-1990s and compare the impact of long-run changes in temperature before and after the expansion. This expansion, related to a set of revised government policies that were instated beginning in 1994, roughly tripled participation in the crop insurance program relative to the late 1980s, and by the end of our study period over 80% of farmers were participating in the program. We find that the effects of temperature in the post-expansion period were the same or even slightly smaller (in absolute value) than the effects in the preexpansion period, which is the opposite of what would be expected if subsidized insurance had reduced farmers' incentive to adapt.³⁰ While this is not a perfect test – other things could have changed over time that affected farmers ability to adapt – it provides suggestive

 $^{^{28}}$ As an alternate approach, and to address any concern that exiting agriculture is a particularly slow process, we adopt a strategy similar to Hornbeck (2012) and examine how the number of farmers in the 1980's and 1990's responded to variation in warming during the 1970s. Point estimates indicate small but statistically significant reductions in the number of farms following earlier exposure to extreme heat, again suggesting that simply not farming may be an immediate adaptation to climate change for some farmers.

²⁹For more details on the program, see http://www.rma.usda.gov/. We note that direct income support from the government constitutes a rather small percentage of cash income during our main study period – an average of 7% in the Corn Belt during the 1980-2000 period – suggesting that the distortionary effects of these programs on the adaptation decision were likely small. Additional data on farm income over time are available here: http://www.ers.usda.gov/data-products/farm-income-and-wealth-statistics.aspx.

³⁰Running the long differences model for 1997-2003 (thus, with 5-yr average endpoints, utilizing data from 1995-2005) gives a $\beta_{GDD>29} = -0.00438$ (SE = 0.00179), which is almost exactly equal to our baseline estimate for the 1980-2000 period, and *less* negative than the coefficient for the long differences run over 1980-1993 ($\beta_{GDD>29} = -0.0056$)

evidence that our results are not being wholly driven by government programs.

4.5 Lack of recognition of climate change

Finally, it could be the case that farmers didn't adapt because they didn't realize the climate was changing and that adaptation was needed. Although this doesn't affect the internal validity of our results, it could mean that our results might provide a poor guide to impacts under future climate change if the need for adaptation becomes apparent. Unfortunately we do not directly observe farmer perceptions of temperature increases, nor their knowledge of the relationship between temperature and crop yields.³¹ To make progress, and building directly on the model presented in Section 2, we first explore whether farmers' responsiveness is a function of characteristics that likely shape their ability to learn about a changing climate. In particular, if adaptation is limited by a difficulty in learning about climate change, then we should observe more adaptation when farmers are given more time to learn about a given change in climate, and more adaptation if they are in an area with a lower temperature variance and thus a clearer "signal" of a given change in climate.

Our data are inconsistent with either of these predictions. First, as shown in Figure 4, point estimates for longer long-difference periods (e.g. the 25- and 30-year estimates in the bottom right panels) are almost uniformly more negative than estimates for the 1980-2000 period, although we cannot reject that they are the same in most cases. Second, we find little evidence that a lower temperature variance at baseline increased adaptation to a subsequent temperature increase. In the first column of Table 5, we re-estimate our main equation, interacting the 1980-2000 extreme temperature change in a given county with the baseline (1950-1980) variance in extreme heat exposure in that county. The estimate on the interaction term is small and statistically insignificant, providing little evidence that a lower underlying variance helped farmers separate signal from noise. As a third check, and following on recent survey evidence suggesting that past experience informs current beliefs about climate change³², we explore whether counties that were rapidly warming prior to our study period were more adaptive during our study period. In particular, we allow the effect

 $^{^{31}}$ A few existing surveys do ask farmers about their perceptions of different aspects of climate change, but the results are difficult to interpret. For instance, although Iowa is one of the states where temperature has changed the *least* in recent years, 68% of Iowa farmers in a recent survey indicated that they believe that "climate change is occurring" (Iowa State Extension Service, 2011), but only 35% of them were concerned about the impacts of climate change on their farm operation. Similarly, only 18% of North Carolina farmers believed that climate change will decrease average yields by at least 5% over the next 25 years (Rejesus, 2012), but slightly less than a 5% decline by 2030 could be consistent with projected impacts under more conservative warming scenarios, meaning these responses do not necessarily suggest a distorted perception of climate change.

 $^{^{32}}$ For instance, Myers et al. (2012) and Howe et al. (2012) show that persons residing in areas that have warmed in recent history are more likely to believe in future climate changes.

of extreme heat over the 1980-2000 period in a given county to depend on the change in extreme heat in that county during the period from 1960-1980, or during 1970-80 (if farmers weight recent evidence more heavily). As shown in Columns 2 and 3 of Table 5, coefficients on either interaction are small and insignificant, and only the coefficient on the 1960-80 interaction has the expected sign.

As a fourth check on the role of beliefs in shaping adaptation, we exploit the fact that beliefs about climate change display well known heterogeneity by political party affiliation, with Republicans consistently less likely than Democrats to believe that climate change is occurring (e.g. Dunlap and McCright (2008)). We re-estimate our main equation and include an interaction between our climate variables and George W. Bush's county-level vote share in the 2000 presidential election. Because public debate and awareness about climate change begin in earnest in the late 1980's and early '90's, this is a reasonable – if highly imperfect – proxy for beliefs about climate change. Results are given in Column 5 of Table 5, and again suggest that expectations about climate change, as proxied by political beliefs, had a minimal effect on the responsiveness of farmers to extreme heat exposure: more Republican counties were if anything less sensitive to extreme heat exposure over the study period.

A final possibility is that adaptation is limited not by farmers' difficulty in learning about changing climate, but instead by difficulty in learning about the production function with respect to climate – in particular, learning that extreme heat can be damaging to productivity. Although this is a different type of learning, it suggests similar empirical tests as before: farmers should have been more likely to learn about the production function had they been given more time to do so, or had they been exposed to extreme heat in a previous period. As just discussed, we find little evidence that either of these predictions is true. This remains the case when we expand the latter prediction to include the possibility that counties could learn from other nearby counties' experiences, interacting county-level changes in extreme heat over 1980-2000 with state-level changes in extreme heat over the previous period (column 4 of Table 5). As an alternate check, and building on existing evidence that higher educational achievement accelerates learning about agricultural technologies (Feder, Just, and Zilberman, 1985), we allow the effect of extreme heat to vary by county-level educational attainment using data on county-level high school graduation rates from the 1980 US Census. As shown in Column 6 of Table 5, we find little evidence that county-level educational attainment affected subsequent adaptation.

As indirect evidence that farmers did recognize that changes in climate were shaping productivity during our study period, we study whether uptake of government crop insurance varied as a function of changing exposure to extreme heat. Although premiums in the crop insurance program are very highly subsidized, meaning that farmers might purchase insurance regardless of the amount of risk they face, the average percent of corn acreage covered by these insurance programs by the end of our study period was "only" 80% (with some counties below 40%), suggesting that there remained some variation in insurance purchases.

To see whether insurance take-up responded to our observed climate trends, we reestimate our long-differences model using insurance adoption at the end of our study period (i.e. a 5-year average over 1998-2002) as the dependent variable. We explore four measures of take-up: the percent of corn acreage in a county enrolled in any of the multiple crop government crop insurance programs, the log of acres enrolled in a county, the number of policies sold in each county, and the total premiums paid (including subsidies) in each county. Results from this exercise are shown in Table A.11. While the coefficients on the temperature variables are only sometimes significant with state-level clustering, results suggest 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. Moving from the 10th to the 90th percentile of the distribution of GDD>29C changes implies roughly a 5 percentage point increase in the acreage insured, a 23% increase in the number of policies sold, and a 20% increase in the total premiums paid. Again, however, only one of these estimates is significant at conventional levels with state-level clustering, so we do not wish to over-sell these results.

Combining this result on insurance take-up with the unchanging productivity effect after the expansion of subsidized crop insurance, the data suggest that increased take-up of insurance served as one response to the changing climate. However, there is no evidence that this response led to increased sensitivity of corn productivity to extreme heat (although we do not observe the counterfactual). It therefore appears that participation in the federal crop insurance program did not lead farmers to take additional risks that would have resulted in increased sensitivity of productivity to extreme heat.

Taken as a whole, then, we find little evidence that farmers who were more likely to learn about the effects of extreme heat on yields, or farmers who were more likely to update their expectations about future exposure to extreme heat, were more able to adapt to subsequent extreme heat exposure. This implies that the lack of observed adaptation is not fully explained by a lack of recognition that the climate was changing for the worse, and indeed we do find some evidence that changes in climate were in fact being recognized. Thus insofar as farmers recognized the warming trend for what it was but had few adaptation options to exploit, then using these observed responses to warming to project future climate change impacts appears a reasonable "business-as-usual" approach. Nevertheless, because we cannot definitively rule out that past responses were affected by imperfect recognition of climate and its effects, and because farmers might more effectively learn about these things in the future, these caveats must be kept in mind when interpreting our projections.

5 Projections of impacts under future climate change

Our final empirical exercise is to build projections of the impacts of future climate change on agricultural outcomes in the US. To do this we combine estimates of climate sensitivities from our long differences approach with projections of future changes in temperature and precipitation derived from 18 global climate models running the A1B emissions scenario. Using data from the full ensemble of available climate models is important for capturing the range of uncertainty inherent in future climate change (Burke et al., 2013). Details of the emissions scenario, the climate models, and their application are provided in the Appendix.

The overall purpose of these projections is to provide insight into potential impacts under a "business-as-usual" scenario in which the future world responds to changes in climate similarly to how it has responded in the past. While it is unknowable whether future responses to climate will in fact resemble past responses – farmers could adapt production practices in previously unobserved ways, or could move crop production to entirely new areas – our long differences approach offers two advantages over existing projections. First, the range of long-run changes in climate projected by climate models through mid-century is largely contained in the range of long-run changes in climate in our historical sample, meaning our projections are not large extrapolations beyond past changes. Second, our estimates better account for farmers' recent ability to adapt to longer-run changes in climate, relative to typical panel-based projections that use shorter-run responses in the past to inform estimates of longer-run responses in the future.

In Figure 6 we present projections of average annual changes in corn yield by 2050 across the 18 climate models. In the top panel we use long differences estimates to generate predictions from precipitation changes, temperature changes, and combined effects of changing both temperature and precipitation. The most substantial negative effects of climate change are driven by increases in temperature, and while the magnitude of the negative effects of temperature vary across climate models, all predict fairly substantial negative effects of future warming on corn productivity. For instance, under climate change projections from the commonly used Hadley CM3 climate model, our long differences estimates deliver a predicted decrease in yields of approximately 27.3% relative to a world that did not experience climate change. The magnitude of this projection is similar to the projections from fixed effects estimates in Schlenker and Roberts (2009).

The bottom panel of the figure compares projections from long difference and panel

models for each of the 18 different climate models. The similarity of regression estimates in the historical data results in projections that are comparable for both long differences and fixed effects, although the long differences estimates are somewhat noisier. We note that this noise is almost entirely due to the coefficient and standard error on GDD below 29C, which is much less precisely estimated in the long differences than in the panel. Since a given temperature rise increases exposure to both harmful and beneficial GDD for almost all counties in our sample, the noise in the GDD below 29C estimate greatly expands the confidence interval on the long differences projections.

Nevertheless, net of any adaptations that farmers have employed in the past, the median climate model projects average yield declines of 15% by mid-century, with some models projecting yield losses as low as 7% and others losses as high as 64%. To put these projected losses in perspective, the 2012 drought and heat wave that was considered one of the worst on record and that received extensive attention in the press decreased average corn yields for the year by 15-25% relative to the prior few years.³³ Our median projection suggests that by 2050, every year will experience losses roughly this large. Valued at production quantities and prices averaged over 2006-2010 for our sample counties, 15% yield losses would generate annual dollar losses of \$6.7 billion by 2050.

6 Conclusions

Quantitative estimates of the impacts of climate change on various economic outcomes are an important input to public policy, informing decisions about investments in both emissions reductions and in measures to help economies adapt to a changing climate. A common concern with many existing impact estimates is that they do not account for longer-run adjustments that economic agents might make in the face of a changing climate. These studies typically rely on short-run variation in weather to estimate how outcomes respond to temperature and precipitation changes, an approach that helps solve identification problems but that might fail to capture important adjustments that agents can make in the longer-run.

We exploit large variation in multi-decade changes in temperature and precipitation across US counties to estimate how farmers have responded to longer-run changes in climate. We argue that these changes are plausibly exogenous and show that their magnitude is on par with future changes in climate projected by global climate models, making them an ideal source of variation to identify historical responses to longer-run changes in climate and in

³³For instance, see http://www.ers.usda.gov/topics/in-the-news/us-drought-2012-farm-and-food-impacts.aspx. The range in estimated losses depend on whether the comparison is against previous season's yield or the yield projected at planting in 2012.

turn to project future impacts.

We show that the productivity of the two main US crops, corn and soy, responded very negatively to multi-decadal changes in exposure to extreme heat. These estimates of longerrun responses are indistinguishable from estimates of how the same crops responded to short-run (annual) variation in extreme heat over the same period, suggesting that farmers were no more able to mitigate the negative effects of climate in the long run than they were in the short run. This apparent lack of adaptation does not appear to be driven by any of a variety of alternative explanations: fixed effect estimates do not appear substantially attenuated relative to long differences estimates, results do not appear to be driven by timetrending unobservables, and farmers do not appear to be adapting along other margins within agriculture. We also provide evidence that this lack of adaptation was not driven by a lack of recognition that climate was changing, perhaps suggesting that farmers either lacked adaptation options or found them too expensive to exploit.

Using climate change projections from 18 global climate models, we project potential impacts on corn productivity by mid-century. If future adaptations are as effective as past adaptations in mitigating the effects of exposure to extreme heat, our median estimate is that future climate change will reduce annual corn productivity in 2050 by roughly 15%, which is on par with the effect of the highly-publicized "extreme" drought and heat wave experienced across the US corn belt in the summer of 2012. Given that these projections account for farmers' present adaptive abilities, our results imply substantial losses under future climate change in the absence of unprecedented adaptation.

References

- Auffhammer, M. and A. Aroonruengsawat. 2011. "Simulating the impacts of climate change, prices and population on California's residential electricity consumption." *Climatic Change* 109:191– 210.
- 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).
- Auffhammer, M., V. Ramanathan, and J.R. Vincent. 2006. "Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in India." *PNAS* 103 (52):19668–19672.
- Bound, J., C. Brown, and N. Mathiowetz. 2001. "Measurement error in survey data." *Handbook* of *Econometrics* 5:3705–3843.
- Burke, M., J. Dykema, D. Lobell, E. Miguel, and S. Satyanath. 2013. "Incorporating climate uncertainty into estimates of climate change impacts." *Working paper*.

- Burke, M. and D. Lobell. 2010. "Food Security and Adaptation to Climate Change: What Do We Know?" In *Climate Change and Food Security*. Springer, 133–153.
- Burke, M.B., E. Miguel, S. Satyanath, J.A. Dykema, and D.B. Lobell. 2009. "Warming increases the risk of civil war in Africa." PNAS 106 (49):20670–20674.
- Cutler, D.M., G. Miller, and D.M. Norton. 2007. "Evidence on early-life income and late-life health from America's Dust Bowl era." *PNAS* 104 (33):13244.
- Davis, D.R. and D.E. Weinstein. 2002. "Bones, Bombs, and Break Points: The Geography of Economic Activity." *American Economic Review* 92 (5):1269–1289.
- DeGroot, Morris H. 1970. Optimal statistical decisions, vol. 82. Wiley-Interscience.
- Dell, M., B.F. Jones, and B.A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." AEJ: Macro 4 (3):66–95.
- 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.
- ——. 2011. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US." American Economic Journal: Applied Economics 3 (4):152–185.
- Dunlap, R.E. and A.M. McCright. 2008. "A widening gap: Republican and Democratic views on climate change." *Environment* 50 (5):26–35.
- Feder, G., R.E. Just, and D. Zilberman. 1985. "Adoption of agricultural innovations in developing countries: A survey." *Economic development and cultural change* 33 (2):255–298.
- Feng, S., M. Oppenheimer, and W. Schlenker. 2012. "Climate Change, Crop Yields, and Internal Migration in the United States." NBER Working paper.
- Fisher, A.C., W.M. Hanemann, M.J. Roberts, and W. Schlenker. 2012. "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment." American Economic Review 102 (7):3749–3760.
- Greenwald, R., MH Bergin, J. Xu, D. Cohan, G. Hoogenboom, and WL Chameides. 2006. "The influence of aerosols on crop production: A study using the CERES crop model." Agricultural Systems 89 (2-3):390–413.
- 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).
- Howe, P.D., E.M. Markowitz, T.M. Lee, C.Y. Ko, and A. Leiserowitz. 2012. "Global perceptions of local temperature change." *Nature Climate Change*.
- Iowa State Extension Service. 2011. "Iowa Farm and Rural Live Poll." Tech. rep., Iowa State University, Extension and Outreach.

- Kelly, D.L., C.D. Kolstad, and G.T. Mitchell. 2005. "Adjustment costs from environmental change." Journal of Environmental Economics and Management 50 (3):468–495.
- Leibensperger, EM, LJ Mickley, DJ Jacob, WT Chen, JH Seinfeld, A. Nenes, PJ Adams, DG Streets, N. Kumar, and D. Rind. 2011. "Climatic effects of 1950–2050 changes in US anthropogenic aerosols–Part 2: Climate response." Atmos. Chem. Phys. Discuss 11:24127–24164.
- Lobell, David B. 2013. "Errors in climate datasets and their effects on statistical crop models." Agricultural and Forest Meteorology 170:58–66.
- Lobell, D.B. and G.P. Asner. 2003. "Climate and management contributions to recent trends in US agricultural yields." *Science* 299 (5609):1032–1032.
- Lobell, DB, G. Bala, and PB Duffy. 2006. "Biogeophysical impacts of cropland management changes on climate." *Geophys. Res. Lett* 33 (6):L06708.
- Lobell, D.B., W. Schlenker, and J. Costa-Roberts. 2011. "Climate trends and global crop production since 1980." Science 333 (6042):616.
- McIntosh, C.T. and W. Schlenker. 2006. "Identifying non-linearities in fixed effects models." UC-San Diego Working Paper.
- Meehl, G.A., T.F. Stocker, W.D. Collins, AT Friedlingstein, A.T. Gaye, J.M. Gregory, A. Kitoh, R. Knutti, J.M. Murphy, A. Noda et al. 2007. "Global climate projections." In *IPCC Fourth* Assessment Report, Working Group I, Chapter 10. Cambridge University Press.
- Meehl, Gerald A, Julie M Arblaster, and Grant Branstator. 2012. "Mechanisms contributing to the warming hole and the consequent US east-west differential of heat extremes." Journal of Climate 25 (18):6394–6408.
- Mendelsohn, R., W.D. Nordhaus, and D. Shaw. 1994. "The impact of global warming on agriculture: a Ricardian analysis." *The American Economic Review* :753–771.
- Miguel, E. and G. Roland. 2011. "The long-run impact of bombing Vietnam." Journal of Development Economics 96 (1):1–15.
- Myers, T.A., E.W. Maibach, C. Roser-Renouf, K. Akerlof, and A.A. Leiserowitz. 2012. "The relationship between personal experience and belief in the reality of global warming." *Nature Climate Change*.
- 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.
- Olmstead, A.L. and P.W. Rhode. 2011. "Adapting North American wheat production to climatic challenges, 1839–2009." PNAS 108 (2):480–485.
- Rejesus, Roderick. 2012. "Farmer perceptions and beliefs about climate change: A North Carolina perspective." Tech. rep., North Carolina Cooperative Extension Service.
- Samuelson, P.A. 1947. Foundations of Economic Analysis. Harvard University Press.

- Schlenker, W. 2006. "Interannual weather variation and crop yields." Working paper .
- Schlenker, W., W.M. Hanemann, and A.C. Fisher. 2005. "Will US agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach." *AER* 95 (1):395–406.
- Schlenker, W. and D.B. Lobell. 2010. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5:014010.
- 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.
- Viner, J. 1958. The long view and the short: Studies in economic theory and policy. Free Press.
- Welch, J.R., J.R. Vincent, M. Auffhammer, P.F. Moya, A. Dobermann, and D. Dawe. 2010. "Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures." PNAS 107 (33):14562–14567.
- Zilberman, D., J. Zhao, and A. Heiman. 2012. "Adoption Versus Adaptation, with Emphasis on Climate Change." Annual Review of Resource Economics 4:13.1–13.227.

Figure 1: Change in temperature $(^{\circ}C)$, precipitation (%), and log corn yields over the period 1980-2000 for counties east of the 100th meridian. Temperature and precipitation are measured over the main April - September growing season. Colors for each map correspond to the colored bins in the histogram beneath the plot.



Figure 2: Productivity of two different corn varieties as a function of temperature.


Figure 3: Relationship between temperature and corn yields. 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°C, as estimated by long differences (dark black line) and panel models (dashed grey line). The shaded area gives the confidence interval around the long differences estimates.



Figure 4: Differences between main estimate from 1980-2000 (specification 2 in Table 1) and other estimates under various starting years and differencing lengths. Dots are differences in estimates and whiskers are 95% confidence intervals of the differences. Standard errors are calculated by bootstrapping, where 1000 samples of 31 states were drawn (with replacement) and the difference between estimates was calculated for each bootstrapped sample.



Figure 5: Percentage of the short run impacts of extreme heat on corn productivity that are mitigated in the longer run. Each boxplot corresponds to a particular time period as labeled, and represent 1000 bootstrap estimates of $1-\beta_2^{LD}/\beta_2^{FE}$ for that time period. See text for details. The dark line in each distribution is the median, the grey box the interquartile range, and the whiskers represent the 5th-95th percentile. The distribution plotted at bottom represents the combination of all the estimates in the above distributions. The values at left give the one-sided p-value on the test that $1-\beta_2^{LD}/\beta_2^{FE}=0$.



% of impact offset

Figure 6: Projected impacts of climate change on corn yields by 2050. Top panel: impacts as projected by the long differences model, for each of the 18 climate models reporting the A1B ("business as usual") climate scenario. Circles represent projection point estimates, whiskers the 95% CI, and colors represent projections using only precipitation changes (blue), temperature changes (black), or both combined (red). Projections are separately for each climate model, as labeled. Bottom panel: projected impacts of combined temperature and precipitation changes across the same climate models, based on long differences (red) on panel estimates (black) of historical sensitivities to climate. The median projection is shown as a dashed line.



	یں۔ - د	-			- - -	-		
Table 1: Comparison o	t long differe (1) Diffs	ences and par (2) Diffs	nel estimates (3) Panel	s of the impacts (4) Panel	ot temperati (5) Diffs	ure and prec (6) Diffs	1pitation on (7) Panel	US corn yields (8) Panel
GDD below threshold	-0.0001	0.0002	0.0004***	0.0002^{**}	-0.0001	0.0003*	0.0005***	0.0003***
	(0.0003)	(0.0002)	(0.0001)	(0.0001)	(0.0003)	(0.0002)	(0.0001)	(0.0001)
GDD above threshold	-0.0053^{***}	-0.0044^{***}	-0.0056^{***}	-0.0062^{***}	-0.0043^{***}	-0.0037^{***}	-0.0048***	-0.0054^{***}
	(0.0010)	(0.0008)	(0.0006)	(0.0007)	(0.0009)	(0.0009)	(0.0005)	(0.0006)
Precip below threshold	0.0515^{**}	0.0297^{**}	0.0118^{***}	0.0095^{*}	0.0253^{**}	0.0115^{**}	0.0068^{***}	0.0057^{***}
	(0.0194)	(0.0125)	(0.0027)	(0.0048)	(0.0123)	(0.0046)	(0.0015)	(0.0019)
Precip above threshold	0.0036^{**}	0.0034^{***}	-0.0008	0.0001	0.0024	0.0029^{***}	-0.0018^{**}	-0.0008
	(0.0017)	(0.0008)	(0.0005)	(0.0004)	(0.0015)	(0.0007)	(0.0007)	(0.0005)
Constant	0.2655^{***}	0.2397^{***}	3.5721^{***}	4.1872^{***}	0.2674^{***}	0.2400^{***}	3.2423^{***}	3.8577^{***}
	(0.0319)	(0.0124)	(0.2491)	(0.3013)	(0.0307)	(0.0115)	(0.2647)	(0.3349)
Observations	1531	1531	48465	48465	1531	1531	48465	48465
R squared	0.258	0.610	0.590	0.863	0.243	0.602	0.593	0.864
Fixed Effects	None	\mathbf{State}	Cty, Yr	Cty, State-Yr	None	\mathbf{State}	Cty, Yr	Cty, State-Yr
T threshold	29C	29C	29C	29C	28C	28C	28C	28C
P threshold	$42 \mathrm{cm}$	$42 \mathrm{cm}$	$42 \mathrm{cm}$	$42 \mathrm{cm}$	$50 \mathrm{cm}$	$50 \mathrm{cm}$	$50 \mathrm{cm}$	$50 \mathrm{cm}$
Data are for US counties earlier	ast of the 100t	<u>h meridian, 19</u>	<u>30-2000. Speci</u>	fications 1-2 and 5	-6 are estimate	<u>ed with long di</u>	ifferences and 3	-4 and 6-8 with
an annual panel, with differ	rent fixed effec	cts shown at bo	ttom; see text	for details. Specif	ications 1-4 us	e piecewise lin	ear thresholds :	as chosen by the
long differences model, and by 1980-2000 average corn **, and 10% * levels.	area (panel).	nolds as cnosen Standard errors	by the paner i s are clustered	model. Regressions at the state level.	s are weighted Asterisks indi	by 1980 count cate statistical	y corn area (lor significance at	ig differences) or the 1% ***,5%

	(1)	(2)	(3)	(4)
	1955 - 1995	1955 - 1995	1960-2000	1960-2000
GDD below threshold	0.0008***	0.0007^{*}	0.0004^{***}	0.0003*
	(0.0003)	(0.0004)	(0.0001)	(0.0002)
GDD above threshold	-0.0066***	-0.0058***	-0.0031***	-0.0023**
	(0.0013)	(0.0020)	(0.0007)	(0.0010)
Precip below threshold	0.0356^{***}	0.0376^{***}	0.0203	0.0166
	(0.0079)	(0.0093)	(0.0135)	(0.0115)
Precip above threshold	0.0017	0.0033^{*}	0.0008	0.0014
	(0.0015)	(0.0017)	(0.0015)	(0.0020)
Observations	2060	2060	2604	2604
R squared	0.621	0.565	0.688	0.699
Fixed Effects	State Yr	Cty Yr	State Yr	Cty Yr
T threshold	29	29	29	29
P threshold	42	42	42	42

Table 2: The effect of climate on yields estimated with a panel of differences.

Dependent variable in all regressions is the difference in the log of smoothed corn yields. Data are a two period panel with 20 year differences. Periods are 1955-1975 and 1975-1995 in Columns 1-2. Periods are 1960-1980 and 1980-2000 in Columns 3-4. The sample of counties is limited to the 1980-2000 corn sample from Table 1. Regressions in Columns 1-2 are weighted by 1955 smoothed corn acres. Regressions in Columns 3-4 are weighted by 1960 smoothed corn acres. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

	Co	orn	Main Spri	ing Crops
	(1)	(2)	(3)	(4)
	Panel	Diffs	Panel	Diffs
GDD below threshold	0.0005^{***}	0.0003	0.0002	0.0003
	(0.0001)	(0.0002)	(0.0001)	(0.0003)
GDD above threshold	-0.0046***	-0.0042***	-0.0024***	-0.0023**
	(0.0005)	(0.0009)	(0.0003)	(0.0011)
Precip below threshold	0.0068***	0.0107**	0.0058***	0.0116^{*}
	(0.0016)	(0.0048)	(0.0014)	(0.0058)
Precip above threshold	-0.0014**	0.0035***	-0.0012**	0.0016
-	(0.0007)	(0.0010)	(0.0005)	(0.0016)
Constant	3.9556***	-0.0116	4.7926***	0.0121
	(0.2539)	(0.0122)	(0.3619)	(0.0210)
Observations	48465	1516	48465	1531
Mean of Dep Variable	5.55	-0.01	5.36	0.03
R squared	0.568	0.579	0.490	0.454
Fixed Effects	Cty, Yr	State	Cty, Yr	State

Table 3: Effects of climate variation on crop revenues

In Columns 1 and 2 the dependent variable is log of agricultural revenue per acre from corn. Dependent variable in Columns 3 and 4 is log of agricultural revenue per acre from 6 main crops grown during the spring season (corn, soy, cotton, spring wheat, hay, and rice). Revenues calculated as yield per acre multiplied by state-level annual prices. Panel regressions are weighted by average area cultivated to corn (Column 1) and main crops (Column 3) from 1978-2002. Long differences regressions are weighted by smoothed corn area in 1980 (Column 2) and smoothed area cultivated to main crops (Column 4). Temperature threshold is 28 and precipitation threshold is 50 in all regressions. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

	(1)	(2)	(3)	(4)	(5)
	Corn area	Corn share	Farm area	Num. farms	Population
GDD below threshold	0.0010	0.0003***	-0.0001	-0.0002	0.0006**
	(0.0012)	(0.0001)	(0.0001)	(0.0002)	(0.0003)
GDD above threshold	-0.0005	-0.0009**	0.0000	-0.0007	-0.0008
	(0.0038)	(0.0004)	(0.0004)	(0.0010)	(0.0015)
Precip below threshold	0.0264	-0.0004	0.0037	0.0021	-0.0236**
	(0.0637)	(0.0034)	(0.0035)	(0.0029)	(0.0106)
	0.0051	0.0016	0.0007	0.0010	0.0047*
Precip above threshold	-0.0051	-0.0016	0.0007	-0.0013	0.0047*
	(0.0063)	(0.0010)	(0.0007)	(0.0033)	(0.0024)
Constant	0.0190	0 0174***	0 001 4***	0 1090***	0.0144
Constant	-0.0130	-0.0174	-0.0014	-0.1830	0.0144
	(0.0687)	(0.0045)	(0.0075)	(0.0157)	(0.0160)
Observations	1511	1516	1523	1526	1526
Mean of Dep Variable	0.075	0.002	-0.068	-0.202	0.045
R squared	0.645	0.418	0.399	0.488	0.392
Fixed Effects	State	State	State	State	State
T threshold	29	29	29	29	29
P threshold	42	42	42	42	42

Table 4: Effects of climate variation on alternate adjustment margins

Dependent variable is difference in log of corn acres (Column 1), difference in share of agricultural area planted to corn (Column 2), difference in total log farm area (Column 3), difference in log number of farms (Column 4), and difference in log county population (Column 5). All regressions are long differences from 1980-2000, with the sample trimmed of extreme outliers in the temperature data. 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.

Table 5:	Heterogenous	effects of clima	te variation or	corn yields		
	(1)	(2)	(3)	(4)	(5)	(9)
GDD above	-0.00427^{***} (0.00115)	-0.00568^{***} (0.00099)	-0.00427^{***} (0.00104)	-0.00424^{***} (0.00095)	-0.00395^{***} (0.00082)	-0.00413^{***} (0.0085)
GDD above*SD GDD above, 1950-1980	-0.00002 (0.0005)					
GDD above*GDD above, 1960-1980		0.00001 (0.00003)				
GDD above*GDD above, 1970-1980			-0.00000 (0.00004)			
GDD above*State GDD above, 1960-1980				-0.00005 (0.0006)		
GDD above*Republican vote share, 2000					0.00197 (0.00424)	
GDD above*High school grad. rate, 1980						-0.00002 (0.00005)
Observations	1531	1531	1531	1531	1530	1531
Mean of Dep Variable	0.24	0.24	0.24	0.24	0.24	0.24
R squared	0.610	0.615	0.610	0.610	0.623	0.614
Fixed Effects	State	State	State	State	State	State
Dependent variable is the difference (1980-200 from Table 1, and regressions are weighted by coefficients on GDD above threshold and relev threshold, and Precipitation below threshold a in all specifications. Standard errors are cluste	(0) in the log of sn 1980 corn acres. Ant interactions a are also included i ered at the state l	and the corn yield All variables in the reported, but and the regression. Evel. Asterisks inc	ds. The sample of the table other that all level effects, G Temperature thre dicate statistical i	f counties is limit in <i>GDD above thr</i> DD below thresh sshold is 29C and significance at the	ed to the $1980-20$ eshold are demea old, Precipitation precipitation thr $1\%^{***}, 5\%^{**}$, al	00 corn sample ned. Only above eshold is 42cm nd 10% * levels.

A Appendix - For Online Publication

A.1 Understanding changes in climate and agriculture over time

Figure A.1 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 ~ 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.2 maps these changes in GDD, showing that extreme-heat outliers are clustered among a few counties in southern Texas.

Figure A.1: 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.

Figure A.2: 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.1.



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 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.3 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 Figure A.3: 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.



sample by running the regression

$$ln(clim_t) = \alpha + \beta t + \varepsilon_t, \tag{12}$$

where t is the sample year. Results for our main GDD>29C variable are shown in Figure A.4 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 β s estimated in Equation (12)), 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 $\Delta GDD>29$ 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.5, with the observed distributions of Δ GDD>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.4: 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	g yield		DV	7: GDD>29)
	\hat{eta}	\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.5: 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.1.1 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:

$$y_{it} = \alpha_i + \delta_1 \bar{T}_i + \delta_2 \nu_{it} + \epsilon_{it} \tag{13}$$

Change in the outcome in county *i* in year *t* responds to both average climate \overline{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:

$$T_{it} = T_{i0} + f_i(t) + \nu_{it} \tag{14}$$

where T_{i0} is some baseline average temperature before any warming starts, $f_i(t)$ is a countyspecific 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.:

$$y_{it} = \alpha + \beta T_{it} + c_i + \theta_i * t + \varepsilon_{it} \tag{15}$$

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:

$$\Delta \bar{T}_i = \bar{T}_{ib} - \bar{T}_{ib} \tag{16}$$

$$= (T_{i0} - T_{i0}) + f_i(b) - f_i(a) + (\bar{\nu_{ib}} - \bar{\nu_{ia}})$$
(17)

$$= f_i(b) - f_i(a) + \Delta \bar{\nu}_i \tag{18}$$

So some part of the $\Delta \overline{T}_i$ 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 $\Delta \overline{T}_i = \theta_i * (b - a) + \Delta \overline{\nu}_i$). So in the long differences regression:

$$\Delta \bar{y}_i = \alpha + \beta_{LD} \Delta \bar{T}_i + \epsilon_i \tag{19}$$

 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-toyear weather ν_{it} . Results are shown in Figure A.6. 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 $\beta_{LD} \rightarrow \delta_2$ (See Figure A.6).

This gives us a useful prediction we can take to the data. Since $\bar{\nu_{it}} \to 0$ 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 $\Delta \bar{\nu_{it}} \to 0$. 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.7, 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 1970-2005 and the average in years 1950-1976, and the estimate using the least data differences the average in years 1991-1993 and 1962-1964.

The results are shown in Figure A.8. 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.6: 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.7: 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.8: 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.2 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.

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

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.3 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.

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

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.

		-		
	(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**
Treep selen uneshera	(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***
Constant	(0.0124)	(0.0125)	(0.0115)	(0.0110)
	(0.0124)	(0.0123)	(0.0113)	(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

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

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.1. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

	Table A.5: 1	Robustness of	results on a	lternate adag	otations to rel	moval of outli	iers	
	(1)	(2)	(3)	(4)	(5)	(9)	(2)	(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	\mathbf{State}	State	\mathbf{State}	\mathbf{State}	\mathbf{State}	\mathbf{State}	\mathbf{State}	\mathbf{State}
T threshold	29	29	29	29	29	29	29	29
P threshold	42	42	42	42	42	42	42	42
Dependent variable is diffe	prence in log of	corn acres (Col	umns $1-2$), diffe	erence in share	of agricultural	area planted to	corn (Columns	3-4), difference
in total log farm area (Col	umns 5-6), and	difference in lo	g number of far	mers (Column	s 7-8). All regre	ssions are long o	differences from	t 1980-2000, and
even numbered regressions	s have the samp	le trimmed of c	utliers indicate	d in Figure A.	1. All regression	is are weighted	by average agri	cultural area
from 1978-2002. Standard	errors are clust	ered at the stat	te level. Asteris	sks indicate sta	tistical significa	nce at the 1% *	***,5% **, and 1	-0% * levels

A.4 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.10, 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 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.9 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.11 (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.9: 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.11 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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1980-2000	1980-2000	1950-2005	1950-2005	1950-2005	1950-2005	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)
D · 1 1	0.0004***	0.0045***	0.0400***	0.0400***	0.0457***	0.0450***	0.0170***
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***
Ĩ	(0.0019)	(0.0019)	(0.0033)	(0.0035)	(0.0036)	(0.0039)	(0.0003)
0	0.0005***	0.000***	0 5000***	0 5000***		0 F11 F444	0 - 11 0 * * *
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

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

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.10: 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.5 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 effects and 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.11: 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.6 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

	0	<u> </u>	1
	(1)	(2)	(3)
	Fixed Effects	Random Effects	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

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

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.12 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.7 Effects on soy productivity

Estimates of the impact of extreme heat on (log) soy yields are shown in Figure A.13. 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.12: Relationship between corn yields and temperature. Estimates represent the change in log corn yield under an additional day of exposure to a given $^{\circ}C$ 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.8 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. profits₂₀₀₀ = revenues₂₀₀₀ - costs₂₀₀₀ 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).³⁴ The second concern is that both costs and revenues will likely respond to annual variation

³⁴In 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.13: 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:

$$\Delta log profits_{1980-2000} = ln(profit_{1997} + profit_{2002})/2 - ln(profit_{1978} + profit_{1982})/2 \quad (20)$$

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.³⁵ 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.

	1 1			
	(1)	(2)	(3)	(4)
	Fertilizer	Seed	Chemicals	Petroleum
GDD below threshold	0.0005	0.0008^{**}	0.0011^{*}	0.0002
	(0.0004)	(0.0004)	(0.0006)	(0.0004)
	0.000 -	0.0000	0.0001	0.0000
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
I recip below unresitora	(0.0220)	(0.0105)	(0.0115)	(0.0010)
	(0.0229)	(0.0125)	(0.0113)	(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
Constant	(0.0210)	(0.0200)	(0.0220)	(0.0201)
	(0.0270)	(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

Table A.8: Effects of Climate Variation on Input Expenditu
--

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.

³⁵As with the profit measure described above, the change in fertilizer expenditures over the period are constructed as: Δ fertilizer expenditure₁₉₈₀₋₂₀₀₀ = (fert₁₉₉₇ + fert₂₀₀₂)/2 - (fert₁₉₇₈ + fert₁₉₈₂)/2

A.9 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,

$$ln(farms)_{ist} - ln(farms)_{is1978} = \beta_t * Extreme heat_{is} + \alpha_{st} + \varepsilon_{ist}, \tag{21}$$

where $Extreme heat_{is}$ 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 statespecific 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.10 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

	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***
G	(0.0177)	(0.0107)	(0.0231)	(0.0057)
1987*Extreme Heating	-0.0579^{***} (0.0190)	-0.0352^{**} (0.0166)	-0.0727^{***} (0.0205)	-0.0455^{**} (0.0216)
1992*Extreme Heating	-0.0351 (0.0223)	-0.0396 (0.0240)	-0.0460^{**} (0.0191)	-0.0430^{**} (0.0179)
1997*Extreme Heating	0.0051 (0.0296)	-0.0155 (0.0318)	-0.0016 (0.0216)	-0.0221 (0.0171)
2002*Extreme Heating	$0.0174 \\ (0.0351)$	-0.0169 (0.0299)	0.0045 (0.0352)	-0.0617^{*} (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

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

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.

characteristics that are correlated with productivity changing differentially between places that heated and those that did not. In Table A.10 we regress the percentage of farms owing 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.11 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

	(1)	(2)
	Diffs, 1978-1997	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)
Procip below threshold	0.2114	0.0608
I Tecip below tilleshold	(0.1470)	0.0008
	(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

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

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.

test for mean reversion in temperature we compare our $\Delta GDD > 29$ over the main 1980-2000 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:

$$\Delta GDD_{is}^{1980-2000} = \alpha + \beta GDD_{is}^{1955-1975} + \eta_s + \varepsilon_{is}$$
(22)

(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.11.1 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.14 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.

	(1)	(2)	(3)	(4)
	% Acreage Enrolled	log Acres Enrolled	Policies Sold	Total Premiums
GDD below threshold	-0.0006*	-0.0005	-1.6411**	-3.7636***
	(0.0003)	(0.0006)	(0.7021)	(1.2821)
			0.0000t	
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***
Compositio	(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

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

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.14: 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.12 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.15. 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.15: 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.

