Adaptation to Climate Change: Evidence from US Agriculture

Marshall Burke and Kyle Emerick*

October 14, 2012

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 60% – 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

*Department of Agricultural and Resource Economics, University of California at Berkeley. Burke: marshall.burke@berkeley.edu. Emerick: kemerick@berkeley.edu. We thank Max Auffhammer, Peter Berck, Richard Hornbeck, Sol Hsiang, David Lobell, Wolfram Schlenker, and seminar participants at UC Berkeley 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, and such adjustment limits the resulting economic damages associated with climate change, then the role for public policy in addressing climate change would appear limited. But if agents appear slow or unable to adjust on their own, and economic damages under climate change appear likely to otherwise be large, then this would suggest a much more substantial role for public policy in addressing future climate threats.

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 shorter-run fluctuations in climate helps to solve identification problems, it perhaps more poorly approximates the ideal climate change experiment. In particular, if agents can adjust in the longer-run in ways that are unavailable to them in the short-run[1], then impact estimates derived from short-run responses to weather might overstate damages

[1] e.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.

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 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 shorter-run 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: impacts of climate on agricultural productivity are identified from variation in *trends* in climate over time, and thus take into account how farmer behavior responds to these longer-run changes. Evidence from the natural sciences suggests that these trends are plausibly exogenous, and differencing eliminates concerns about time-invariant omitted variables that plague cross-sectional attempts to quantify adaptation. Second, observed variation in climate change 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

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

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 compare these long differences estimates to panel estimates of short-run responses to climate. Long run adaptations appear to have mitigated less than 60% of the short-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.

Results are less conclusive as to *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. While survey-based evidence seems to support the latter explanation – existing surveys of US farmers suggest limited recognition of climate change – responses to recent climate trends do not appear to vary as a function of education or political attitudes, factors that might predict recognition of or belief in climate change. Additionally, we cannot rule out that land values have decreased substantially in areas that experienced substantial recent warming, which would suggest warming-related trends are both recognized by farmers and expected to be persistently harmful.

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 future adjustment capabilities are constrained to what farmers were capable of in the recent past. Nevertheless, our projections improve on past work in two ways: as mentioned above they are less dependent on large out-of-sample extrapolation, and relative to panel-based projections that use shorter-run responses in the past to inform estimates of longer-run responses in the future, our estimates account for farmers' recent ability to adapt to longer-run changes in climate. 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 expected yield losses resulting from the well-publicized "extreme" drought and heat wave that struck the US midwest in the summer of 2012. 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 global damages 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 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, 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 outline our estimation approach and discuss how it differs from previous literature. Section 3 describes our main results on the extent of past adaptation. In Section 4 we try to rule out alternative explanations for our results. 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 Approach

Agriculture is a key sector where future climate change is estimated to have large detrimental impacts, and is a primary focus of the empirical literature on climate change impacts. How might key agricultural outcomes respond to a changing climate? Suppose that at a given location there is a "best response" function of any agricultural variable of interest (yield, profits, land values, etc) to longer-run changes in climate at that location - i.e. the maximum value of the outcome that can be obtained given different long-run climate realizations. The dark line in Figure 2 traces this hypothetical climate response function with respect to temperature. Imagine also that farmers' responses to short-run fluctuations in temperature do not necessarily coincide with these longer-run responses. These short run response curves (dotted lines in Figure 2) could lie fully below the longer-run response curve, as in the panel A of Figure 2, indicating a situation where the short-run options available to farmers in the face of climate fluctuations are less effective than the longer-run options. Short-run response functions could also lie more flat than the longer run response curve (panel B in Figure 2), indicating that farmers have short-run options – pumping irrigation water, drawing down stocks – that are not available in the long run.

Because a primary anticipated effect of future climate change will be a long-run increase in average temperatures, understanding how agricultural outcomes will be affected by climate change thus requires knowledge of the longer-run response curve in Figure 2. In pioneering work, Mendelsohn, Nordhaus, and Shaw (1994) use cross-sectional variation in average temperature and precipitation (and their squares) to explain variation in land values across US counties. Land values represent the present discounted value of the future stream of profits that could be generated with a given parcel of land, and so in principle should embody any possible adaptation to average climate. Using cross-sectional variation with land values on the left-hand-side would then appear to trace out the desired long-run relationship between average climate and the value of productivity, and a projected rise in average temperatures from T_0 to T_1 would be interpreted as a movement from point (a) to point (c) in Figure 2.

Cross sectional models in this setting make two oft-criticized assumptions: that average climate is not correlated with other unobserved factors (soil quality, labor productivity, etc) that also affect outcomes of interest (Schlenker, Hanemann, and Fisher, 2005; Deschênes and Greenstone, 2007), and that farmers recognize that they have experienced a shift in climate and have the knowledge and resources to re-optimize appropriately (Hanemann, 2000). For instance, suppose that achieving the point (c) requires a different crop variety and different crop management than is used at (a). The farmer must know what variety to switch to, have access to it, and know how to adjust his management. It's not obvious whether - or how quickly - these things will happen.

Given these concerns, more recent work has used panel data to explore the relationship between climate variables and agricultural outcomes (Deschênes and Greenstone (2007); Schlenker and Roberts (2009); Welch et al. (2010); Lobell, Schlenker, and Costa-Roberts (2011)).[3] The impact of climate variables on annual outcomes of interest are identified from inter-annual deviations from location-specific means, and this year-to-year variation in climate (typically termed "weather") is plausibly exogenous.[4] Many studies then combine these estimated short-run responses with output from global climate models to project potential impacts under future climate change.

In making these projections, the implicit assumption is again that short-run responses to inter-annual variation in weather are representative of how farmers will respond to longer-run changes in average climate. As shown conceptually in Figure 2, it is not clear that this will be the case. Given a predicted rise in temperature from T_0 to T_1, panel-based estimates would predict that outcomes shift to point (b), and the estimated impact of that shift, $(b) - (a)$, might either over- or under-estimate the actual impact of the increase in mean temperature once farmers have had time to adapt.

2.1 The long differences approach

We attempt to simultaneously overcome the limitations of both the cross-sectional and panel approaches by long differencing. Suppose that data on longer-run mean climate and mean outcomes are available for a set of locations at two different points in time - for instance, decadal averages for two different decades. Let $\overline{y_{i1}}$ represent average outcomes in location i in time period 1, $\overline{y_{i2}}$ average outcomes in the second period, and similarly for the climate variables ($\overline{T_{i1}}$ is the average temperature in i in period 1, etc). Suppose the within-period

[3] Examples 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).

[4] McIntosh 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.

data generating process is given by:

$$\overline{y_{i1}} = \beta_1 \overline{T_{i1}} + \beta_2 \overline{P_{i1}} + \alpha_i + \varepsilon_{i1} \tag{1}$$

where α_i is an unobserved time invariant factor correlated with the climate variables that will bias cross-sectional estimates of β_1 and β_2. Differencing over the two periods gives the following:

$$\overline{y_{i2}} - \overline{y_{i1}} = \beta_1(\overline{T_{i2}} - \overline{T_{i1}}) + \beta_2(\overline{P_{i2}} - \overline{P_{i1}}) + (\varepsilon_{i2} - \varepsilon_{i1}) \tag{2}$$

with the unobserved α_i's dropping out. Rewritten this gives:

$$\Delta\overline{y_i} = \beta_1 \Delta\overline{T_i} + \beta_2 \Delta\overline{P_i} + \Delta\varepsilon_i, \tag{3}$$

Generating unbiased estimates of β_1 and β_2 requires that changes in temperature and precipitation 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.

Our primary analysis focuses on the effect of these longer-run changes in climate on productivity of the primary US field crops, corn and soy. These two crops account for about 65% of the cropped area in the US and 30% of the value of total US agricultural production, and high quality corn and soy yield data – the most basic measure of productivity[5] – are available on the county level across the US for at least the last half century.

Estimating the impact of climate on productivity of these crops with the long differences approach in (3) 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 trends that are randomly assigned to different regions on Earth. 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 decadal trends in climate 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.

[5]Yield is output per unit area, and is defined as bushels per acre for corn and soy in the US.

2.2 Adapting to a changing climate

The long differences approach also allows us to quantify the extent of recent climate adaptation in agriculture. The climate literature generally understands adaptation to mean any adjustment to a changing environment that exploits beneficial opportunities or moderates negative impacts.[6] 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 take two approaches to quantifying recent adaptation to longer-run changes in climate. In the first approach, we compare coefficient estimates in (3) to coefficient estimates from an annual panel model. The panel approach will estimate the impact of fluctuations in climate net of any short-run responses that farmers can undertake, and the long differences approach will capture the impact of climate net of any longer-run adaptations farmers can undertake. Comparing the two then provides an estimate of whether farmers have been able to adjust to longer-run changes in climate better than they have been able to respond to shorter-run changes, which we interpret as evidence of adaptation. As suggested above, to the extent that warmer temperatures are harmful to agricultural productivity, economists would generally expect a longer-run increase in temperature to be less damaging than a shorter-run increase of the same magnitude, although panel B in Figure 2 and the discussion above suggests this will not necessarily be the case. We attempt to rule out other explanations for divergence between panel and long-differences estimates - e.g. measurement error - in Section 4.

Conditional on recognizing that the climate has changed and believing that adjustment is necessary, a farmer could in principle adjust along many possible margins: a farmer growing a particular crop could alter her planting date, switch crop varieties, change her input mix, and/or utilize (or invest in new) irrigation technology, or she might choose to grow different crops or to leave agriculture altogether. Our second approach to quantifying adaptation is thus to use our long differences approach and available data to directly examine whether farmers have adjusted along these margins in response to a changing climate.

Finally, given that we find limited evidence of adaptation along any of these margins, we attempt to shed light on *why* adaptation has been limited: is it because farmers did not see the need to adapt, or because adaptation options were limited or appeared costly? As direct evidence we review data from existing interviews on farmer perceptions of climate change. As indirect evidence we examine whether the impacts of climate trends vary as a function of education or political attitudes, factors that might predict recognition of or belief

[6] See Zilberman, Zhao, and Heiman (2012) and Burke and Lobell (2010) for an overview.

in climate change, and examine whether climate trends have had demonstrable impacts on the evolution of land values across counties. Land values represent the present discounted value of the future stream of profits that could be generated with a given piece of land, and if land values and agricultural productivity both respond strongly to climate trends, then this suggests that markets believe that exploitable adaptation opportunities are somewhat limited or expensive.

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.[7] 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 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.[8]

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). GDD measure the amount of time a crop is exposed to temperatures between a given lower and upper bound, with daily exposures summed over the growing season to get a measure of annual growing degree days. Denoting the lower bound as t_l and the upper bound as t_h, if t_d is the average temperature on a given day d, then degree days for that day are calculated as:

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

[8] We thank Wolfram Schlenker for sharing the weather data and the code to process them.

$$GDD_{d;t_l:t_h} = \begin{cases} 0 & \text{if } t_d \leq t_l \\ t_d - t_l & \text{if } t_l < t_d \leq t_h \\ t_h - t_l & \text{if } t_h < t_d \end{cases}$$

Daily degree days are then summed over all the days in the growing season (typical April 1 to September 30th for corn in the United States) to get an annual measure of GDD.

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.[9] We estimate the long differences model:

$$\Delta y_{is} = \beta_1 \Delta GDD_{is;l_0:l_1} + \beta_2 \Delta GDD_{is;l_1:\infty} + \beta_3 \Delta Prec_{is;p<p_0} + \beta_4 \Delta Prec_{is;p>p_0} + \alpha_s + \Delta\varepsilon_{is}, \quad (4)$$

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 (4) is the sum of GDD between the bounds l_0 and l_1, and $\Delta GDD_{is;l_0:l_1}$ 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 $\Delta GDD_{is;l_1:\infty}$ term measures the change in these GDD between the two periods.[10] In contrast to Schlenker and Roberts (2009), we measure precipitation in a county as a piecewise linear function with a kink at p_0. The variable $Prec_{is;p<p_0}$ is therefore the difference between precipitation and p_0 interacted with an indicator variable for precipitation being below the threshold p_0. $Prec_{is;p>p_0}$ is similarly defined for precipitation above the threshold.[11] In the estimation we set $l_0 = 0$ and allow the data to determine l_1 and p_0 by looping over all possible

[9]We 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.

[10]As an example, if $l_0 = 0$ and $l_0 = 30$, then a given set of observed temperatures -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.

[11]A 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 $\Delta Prec_{is;p>p_0} = 10$.

thresholds and selecting the model with the lowest sum of squared residuals.

Importantly, we also include in (4) 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 (4), 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.

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. Differential trends in precipitation have been similarly large, with precipitation decreasing by more than 30% in some counties and increasing by 30% in others. By way of comparison, the upper end of the range in these recent temperature trends is roughly equivalent to the mean warming projected by global climate models to occur over US corn area by 2030, and the range in precipitation trends almost fully contains the range in climate model projections of future precipitation change over the same area by the mid-21st century. More importantly, substantial variation is apparent 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 yields rose by 20-30% only 100 miles to the south.

One concern is that these differences over time might simply reflect large short-run variation around the endpoint years rather than substantial longer-run trends in climate: e.g. a very hot or cold year within one of the five-year endpoints from which our differences are constructed could conceivably shift the average for that endpoint substantially, causing us to conflate short-run variation with longer-run trends. To further demonstrate that our differences measure is actually picking up substantial differences in longer-run climate trends across counties, we estimate the annual growth rate in temperature or precipitation for each county over our main 1978-2002 study period. The within-state distribution of annual growth in extreme heat exposure is plotted in Figure A.3, demonstrating cross-county variation in annual growth rates of up to 4 percentage points in many key corn-growing states, which are substantial when aggregated over our 20+ year study period. This indicates that our differences measure is indeed picking up longer-run trends in climate.

2.4 Are recent climate trends exogenous?

There are a few potential violations to the identifying assumption in (4). 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^{o}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.[12] However, aerosols also increase the "diffuse" portion of light (picture 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

[12] Crops 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)).

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 pre-existing 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.

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 (e.g. Meehl, Arblaster, and Branstator (forthcoming)). As such, these trends appear to represent a true "natural experiment", and are likely exogenous with respect to the outcomes we wish to measure.

3 Empirical Results

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 1-3 we run both models under the thresholds selected by the long differences, and in Columns 4-6 we fix thresholds at values chosen by the panel. Given that the inclusion of state fixed effects might eliminate a significant amount of variation if climate trends vary only little within states, we show results for the long differences model with and without state fixed effects. We adopt the more conservative model with state fixed effects as our preferred specification.

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° results in a decrease in overall corn yield of 0.44%. The panel model delivers a slightly more negative point estimate – a -0.56% yield decline for every one degree increase above 29° – but 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 4-6), and we attempt to quantify what these estimates imply about adaptation below. Interpreting the total effect of an increase in temperature requires considering both GDD coefficients. As an example, consider a county that experienced an increase in daily temperatures of 1° for 20 days of the growing season. If 10 of those days were originally below 29° and the remaining 10 days were above 29°, then the total predicted effect on log corn yields from our model is 10*0.0002 - 10*0.0044 = -0.042.

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.

This is potentially of interest given the surprisingly muted role that precipitation has played in past estimates of the impacts of climate change on agriculture (e.g. Schlenker and Roberts (2009); Schlenker and Lobell (2010)). We offer two potential explanations. First, there could be adaptations available in the short run that may not be available over longer time horizons – e.g. pumping irrigation water in response to a dry year, but not being able to do so over the longer run if the resource is depletable. A second explanation is measurement error. If precipitation is measured with error, and fixed effects models amplify attenuation bias concerns associated with errors-in-variables, then such error is a candidate explanation for the smaller fixed effects estimates (measurement error is also an important concern for temperature, which we return to below). While we cannot definitively rule out either explanation, we note that even the long differences precipitation estimates remain small: on a sample growing season precipitation 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%. As we show below, 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.

3.2 Robustness of corn results

One potential concern is that differencing between two endpoints introduces a large amount of noise in the measurements of longer term changes in climate and yields. While our estimates in Table 1 take 5 year averages of each climate variable and only consider differences in these "smoothed variables" in the regressions, our choice of endpoints remains relatively arbitrary. As a robustness check, we estimated the main long differences specification using a combination of different time periods and differencing lengths. More specifically, we estimated (4) varying T_0 from 1955 to 1995 in 5 year increments. For each value of T_0 we estimated 5, 10, 15, 20, 25, and 30 year difference models.[13] Results are shown graphically in Figure 4. For each starting point and differencing length we display the point estimate of β_2 (the responsiveness of corn yield to extreme heat) and its 95% confidence interval. In each panel the horizontal line is at the value -0.0044, which is our estimate of β_2 in Column 2 of Table 1. The average estimate of β_2 across these 39 models is -0.0058, with only 9 of the estimates of β_2 being statistically different from -0.0044 and only one statistically different in the positive direction. Furthermore, if more adaptation is possible over a longer time frame, we would expect estimates on the the impact of extreme heat to be less negative for longer differencing lengths (e.g. 25 or 30 years). Point estimates if anything suggest the opposite: the effect of exposure to extreme heat is at least as negative for these longer periods.

We verify that yield results are not being driven by outliers. As shown in Figures A.1 and A.2, there are a handful of counties that experienced yield trends that were substantially more negative than the rest of the distribution, and a handful of counties along the border of south Texas that saw increases in exposure to extreme heat that were substantially beyond what other counties experienced. Table A.1 shows that our corn yield results are robust to dropping these outliers.

Our estimates would be biased in the presence of time-varying unobservables correlated with both climate and yields. To deal with the possibility of time-trending unobservables within states, we use our many decades of data to construct a two period panel of long

[13] Some 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.

differences, which further weakens our identification assumption. We estimate the following model:

$$\Delta y_{it} = \beta_1 \Delta GDD_{it;l_0:l_1} + \beta_2 \Delta GDD_{it;l_1:\infty} + \beta_3 \Delta Prec_{it;p<p_0} + \beta_4 \Delta Prec_{it;p>p_0} + \alpha_i + \delta_t + \varepsilon_{it}, \quad (5)$$

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 trends 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 trends across the two periods.

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 above 29 are also similar to our baseline panel estimates in Table 1. The main long differences estimates are therefore robust to controlling for a richer set of county specific time-varying factors.

3.3 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 short-run 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.[14] 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 bootstrapped adaptation estimates for each time period are shown in Figure 5. Results 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*

[14]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. The distribution of accumulated estimates for each time period is shown in Figure 5.

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.

3.4 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.4 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 provide little evidence that any longer-run adaptation measures undertaken by corn and soy farmers over the last few decades have mitigated the shorter-run effects of extreme heat on corn and soy productivity. We now explore the extent to which limited apparent adaptation in crop yields is an artifact of measurement error or selection, and/or to what extent it might obscure adaptation along other margins.

4.1 Measurement error and alternate adaptation measures

A key concern with fixed effect 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 if panel estimates are attenuated relative to long differences estimates, one explanation for the limited observed yield adaptation is 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.

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 rely entirely on within-county variation. Table A.3 shows that results from all three estimators are remarkably similar, providing suggestive evidence that measurement error is not responsible for the lack of difference between estimates from fixed effects and long differences.

As a final check that our comparison of long differences and panel estimates is not generating an artificially low estimate of adaptation for corn and soy yields, we adopt a slightly different approach to estimating adaptation that does not depend on making this comparison. In general, adaptation to a changing climate would imply that exposure to a shock in a given year should reduce the impact of a similarly sized shock in the future. We implement this idea in our panel setting by including a set of interactions between lagged climate variables and current climate: if adaptation is occurring then the effects of current unfavorable temperature should be smaller when past shocks are large. The full specification is provided in the Appendix, and the results are given in Table A.4. Estimating a model with up to 10 lags and their interactions with current climate, we find that past exposure to extreme heat reduces the effect of current exposure by a maximum of about 12%, again suggesting limited recent adaptation to extreme heat.

4.2 Adaptation along other margins

Our analysis up to this point can only identify adaptations that directly affect the magnitude of temperature and precipitation effects on yields. While this will capture many of the oft-mentioned modes of adaptations, such as switching to different seed varieties or applying more irrigation water to a particular crop, it clearly does not capture all possible margins of adjustment, and as such might underestimate the extent of recent adaptation to climate change.

One way to capture broader economic adjustment to variation 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 this approach. The first is that such measures 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.

We take two alternate approaches to exploring impacts on economic productivity. 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.[15] 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, where we attempt to capture changes in *average* expenditures by averaging two census outcomes near each endpoint and then differencing these averaged values.[16]

[15] Prices are only available at the state level and to our knowledge do not vary much within states within a given year.

[16] For example, ag census data are available in 1978, 1982, 1987, 1992, 1997, and 2002. The change in

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.[17] Furthermore, we show in Table A.5 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 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.[18] 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%.

A final adaptation available to farmers would be to exit agriculture altogether, an option that recent literature has suggested is not uncommon. For instance, Hornbeck (2012) shows that population decline was the main margin of adjustment across the Great Plains after

---

fertilizer expenditures over the period are constructed as: Δfertilizer expenditure$_{1980-2000}$ = (fert$_{1997}$ + fert$_{2002}$)/2 - (fert$_{1978}$ + fert$_{1982}$)/2

[17]Coefficient 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, adaptation-related 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 in *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.

[18]As shown in Table A.2 - and unlike for our yield outcomes - a few outcomes in this table are altered fairly substantially when these five outliers (0.003% 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, we feel comfortable dropping them.

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 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.[19] 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.

4.3 Why is crop yield adaptation limited?

The possibility that farmers might exit agriculture in response to climate change offers one explanation for the limited adaptation in corn and soy yields: selection. In particular, if exit is selective, and if those that exit are somehow better performing (e.g. they are also the first to realize climate is changing for the worse) leaving a remaining population with lower yields, then we could mistake a change in the composition of the farming population for a lack of adaptation. 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.7 provides suggestive evidence that this is not the case. The percentage of farms owning more than $20,000 equipment, which is positively correlated

[19]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.

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.

Why else might recent adaptation have been limited? We offer two explanations: there were few adaptation opportunities to exploit (or existing options were prohibitively expensive), or there was little recognition that the climate was in fact changing and that adaptation was needed. Which explanation it is could have important implications for projections of the impact of future climate change on agricultural outcomes. If adaptation opportunities are available but are unexploited because farmers do not yet recognize that they are needed – e.g. it might not be obvious that small increases in exposure to extreme heat cause large yield declines – then future adaptation could be larger than past adaptation if the need for it becomes apparent. On the other hand, if available adaptations are already exploited, then past impacts could provide a useful guide for what future impacts will look like in the absence of concerned investment in adaptive capacity.

One way to shed light on which it might be is to ask farmers about their perceptions of recent climate change. While the ideal survey for our purposes is unavailable[20], a few surveys do ask farmers about their perceptions of different aspects of climate change. For instance, 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 yields by at least 5% over the next 25 years (Rejesus, 2012). While these results are a little difficult to interpret – Iowa is one of the states where temperature has changed the *least* in recent years – they do not suggest overwhelming recognition of the role of recent temperature changes in farm outcomes.

An alternate approach is to estimate whether farmers' responsiveness to longer-run climate trends varies as a function of county characteristics that are plausibly correlated with their beliefs in or recognition of climate change. For instance, 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)). Similarly, it could be that better-educated farmers are more likely to recognize that the climate is changing. Evidence for heterogeneity on either of these fronts would then suggest that an observed lack of adaptation could be driven by perceptions about recent climate change.

[20]Such a survey would elicit farmer beliefs about how their crop responds to extreme heat and how their exposure to extreme heat has changed in recent decades, and allow the researcher to map these beliefs to observed data on temperature changes.

To test for heterogeneity along these lines, 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, or an interaction with the percent of residents in each county 25 and older who have at least a high school diploma.[21] Point estimates on the interaction between our extreme heat measure and either the vote share or high school graduation rate are small and statistically insignificant (results available upon request), providing some evidence that beliefs about or perception of climate change are not driving our observed lack of adaptation.

As a final approach, we explore how land values respond to increased exposure to extreme heat. Land values represent the present discounted value of the future stream of profits that can be realized from a given piece of land, and if markets believe that productivity is affected by exposure to extreme heat and that these effects will persist into the future net of any adaptation efforts, then land values should decline in places that have differentially warmed. Long differences estimates of the effect of climate on land values are shown in the last column of Table 4. The point estimate on the response of land values to extreme heat is indeed negative but imprecisely estimated. The 95% confidence interval on the predicted change in land values for counties in our sample that warmed the most is [-12%, 6%], meaning we can't rule out both no effect and fairly large negative effects.[22]

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., 2011). 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, our long differences approach offers two ad-

[21] Data on high school graduation rates are from the 2000 Census, and data on 2000 presidential vote shares are from http://spatialnews.geocomm.com/features/election2000/

[22] Results for the untrimmed sample, which includes 5 counties in southern Texas that suffered inordinate warming, are substantially more negative and significant at the 10% level (see Table A.2). We do not emphasize this result since it appears to be substantially driven by these 5 counties.

vantages 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 much 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 increase in temperature 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 is projected at the time of writing to decrease average corn yields for the year by 15-20% relative to the prior few years.[23] Our median projection suggests that by 2050, every year will experience losses this large. Valued at production quantities and prices averaged over 2006-2010 for our sample counties, 15%

[23]As forecast by the Oct 11th 2012 version of the seasonally-updated USDA forecast, available at http://www.usda.gov/oce/commodity/wasde/latest.pdf.

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 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 climate to estimate climate responses, 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 recent variation in temperature and precipitation trends across US counties to estimate how farmers have responded to longer-run trends in climate. We argue that these trends 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 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 longer-run responses are indistinguishable from estimates of how the same crops responded to short-run (i.e. 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 time-trending unobservables, and farmers do not appear to be adapting along other margins within agriculture.

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 impact estimate suggests that future changes in climate 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 abilities to adapt to climate change, our results imply substantial losses under future climate change in the absence of efforts to help farmers better adapt to extreme heat.

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. 2011. "Global climate models and climate data: A user guide for economists." *Working paper* .

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. 2011. "Incorporating climate uncertainty into estimates of climate change impacts, with applications to US and African agriculture." *NBER Working paper 17092* .

Burke, M. and D. Lobell. 2010. "Food Security and Adaptation to Climate Change: What Do We Know?" In *Climate Change and Food Security: Adapting Agriculture to a Warmer World*. 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.

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." *The American Economic Review* 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.

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* , forthcoming.

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." *Journal of Econometrics* 31 (1):93–118.

Hanemann, W.M. 2000. "Adaptation and its measurement." *Climatic Change* 45 (3):571–581.

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

Iowa State Extension Service. 2011. "Iowa Farm and Rural Live Poll." Tech. rep., Iowa State University, Extension and Outreach.

Kilian, L. 2008. "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature* :871–909.

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, 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., G. Arblaster, and G. Branstator. forthcoming. "Mechanisms contributing to the warming hole and the consequent US east-west differential of heat extremes." *J. Climate* .

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.

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.

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." *The American Economic Review* 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 (^{o}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.

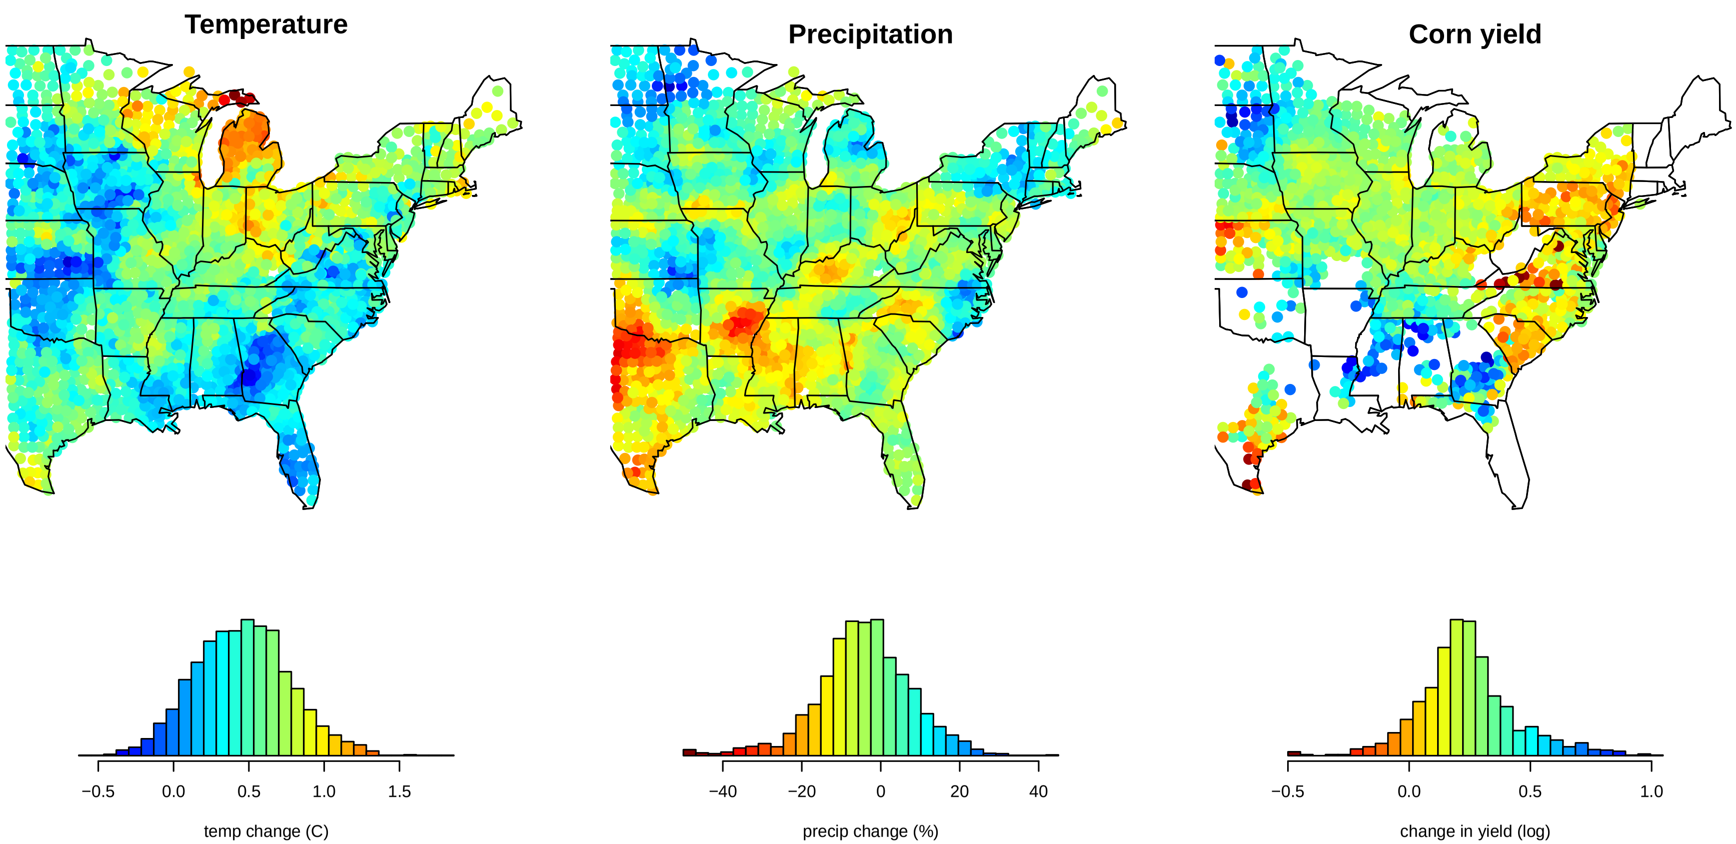

Figure 2: Hypothetical response of agricultural outcomes to temperature. In panel A, short run response options are less effective than longer-run response options. In panel B, there are adaptation options available in the short run that are not available in the longer-run.

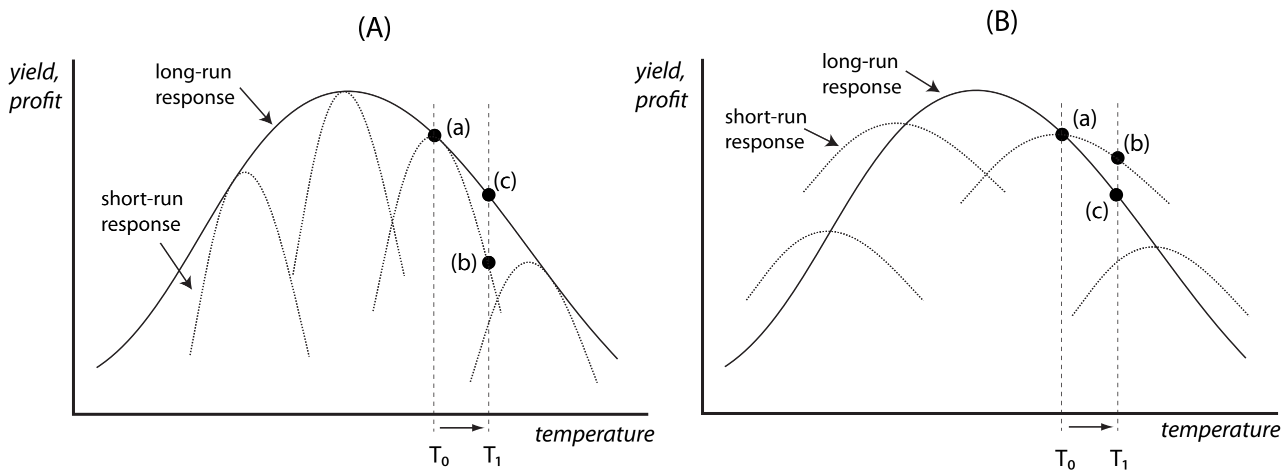

Figure 3: Relationship between temperature and corn yields. Estimates represent the change in log corn yield under an additional day of exposure to a given °C temperature, relative to a day spent at 0°C, as estimated by long differences (dark black line) and panel models (dashed grey line). The shaded area gives the confidence interval around the long differences estimates.

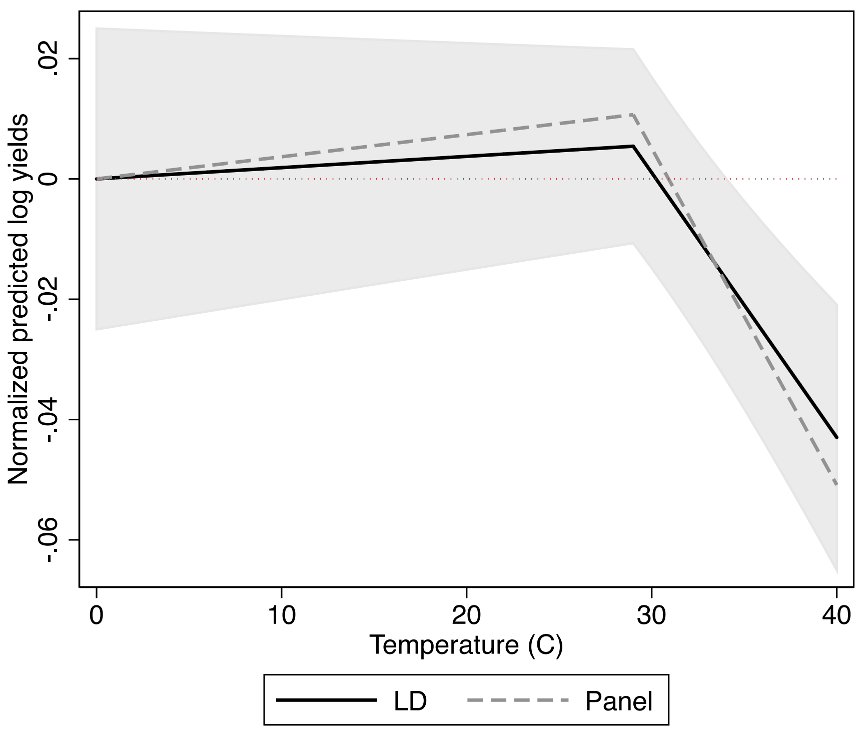

Figure 4: Effects of extreme heat on corn yields under various starting years and differencing lengths. The horizontal line is the point estimate from Specification 2 in Table 1 estimated over 1980-2000.

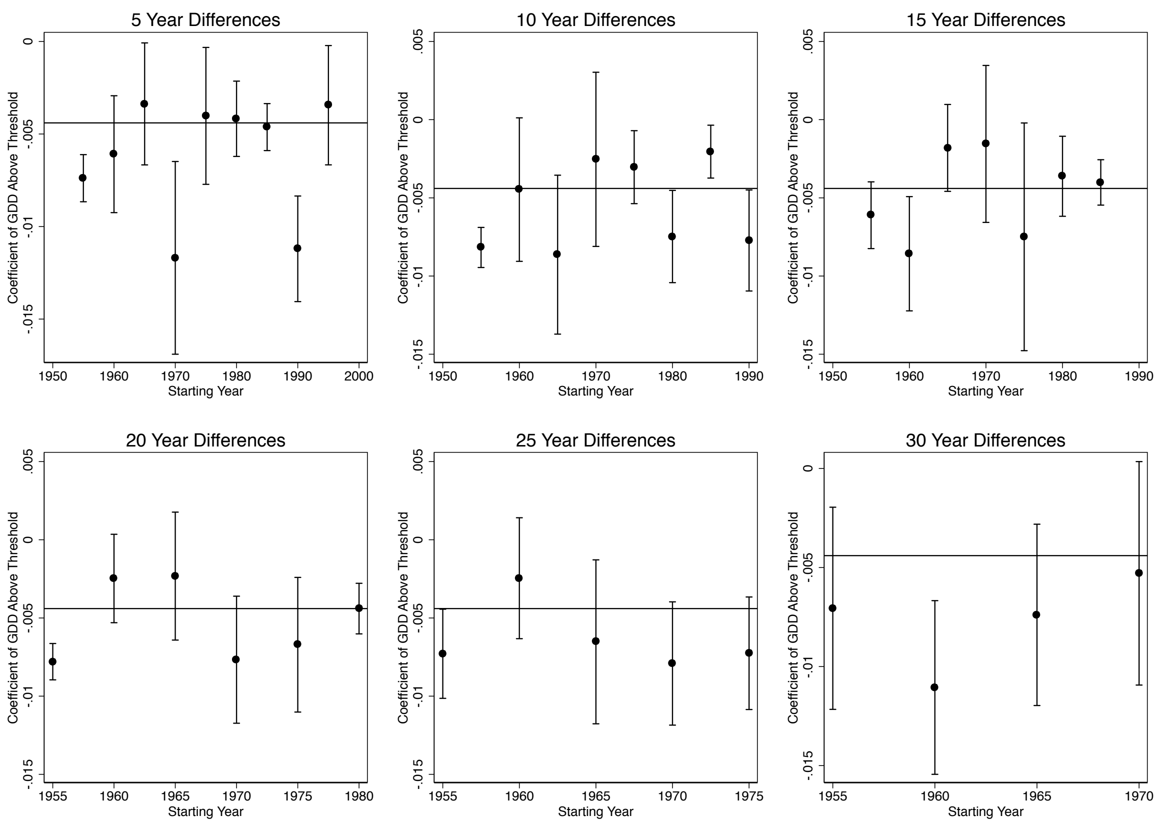

Figure 5: Percentage of the short run impacts of extreme heat on corn productivity that are mitigated in the longer run. Each boxplot corresponds to a particular time period as labeled, and represent 1000 bootstrap estimates of $1-\beta_2^{LD}/\beta_2^{FE}$ for that time period. See text for details. 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.

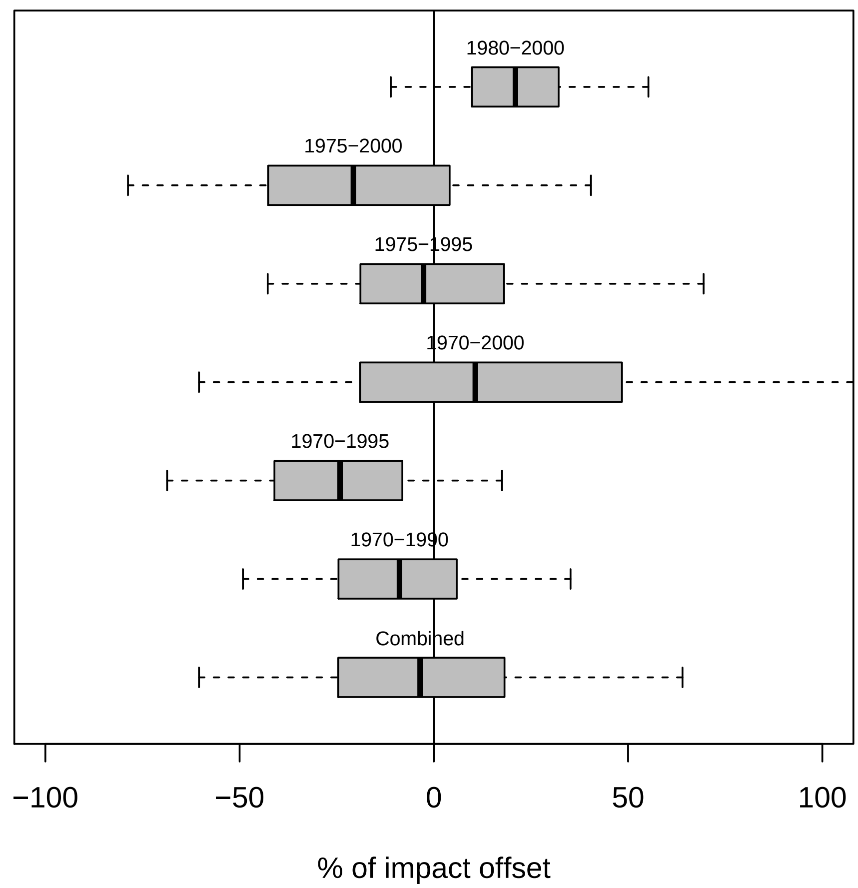

Figure 6: Projected impacts of climate change on corn yields by 2050. Top panel: impacts as projected by the long differences model, for each of the 18 climate models reporting the A1B ("business as usual") climate scenario. Circles represent projection point estimates, whiskers the 95% CI, and colors 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.

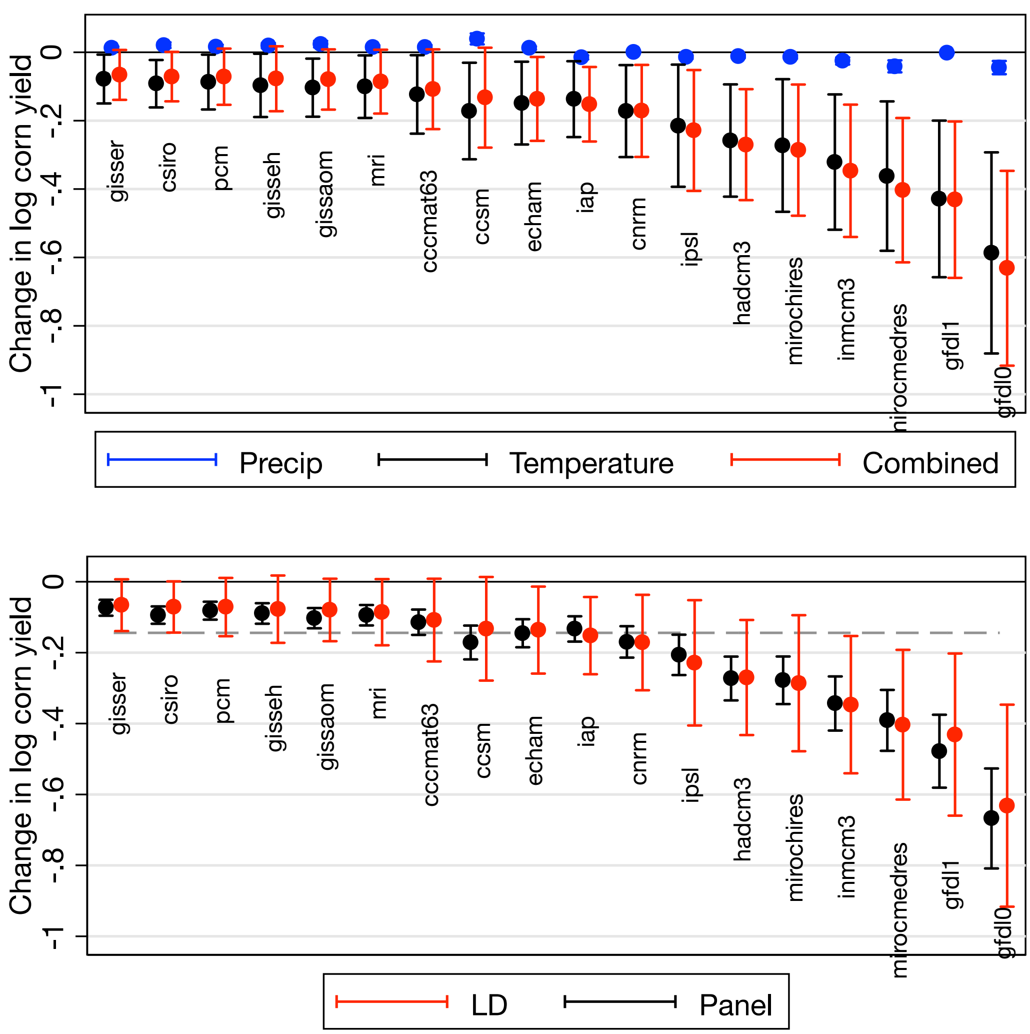

Tables

Table 1: Comparison of long differences and panel estimates of the impacts of temperature and precipitation on US corn yields

	(1) Diffs	(2) Diffs	(3) Panel	(4) Diffs	(5) Diffs	(6) Panel
GDD below threshold	-0.0001	0.0002	0.0004***	-0.0001	0.0003*	0.0005***
	(0.0003)	(0.0002)	(0.0001)	(0.0003)	(0.0002)	(0.0001)
GDD above threshold	-0.0053***	-0.0044***	-0.0056***	-0.0043***	-0.0037***	-0.0048***
	(0.0010)	(0.0008)	(0.0006)	(0.0009)	(0.0009)	(0.0005)
Precip below threshold	0.0515**	0.0297**	0.0118***	0.0253**	0.0115**	0.0068***
	(0.0194)	(0.0125)	(0.0027)	(0.0123)	(0.0046)	(0.0015)
Precip above threshold	0.0036**	0.0034***	-0.0008	0.0024	0.0029***	-0.0018**
	(0.0017)	(0.0008)	(0.0005)	(0.0015)	(0.0007)	(0.0007)
Constant	0.2655***	0.2397***	3.5721***	0.2674***	0.2400***	3.2423***
	(0.0319)	(0.0124)	(0.2491)	(0.0307)	(0.0115)	(0.2647)
Observations	1531	1531	48465	1531	1531	48465
R squared	0.258	0.610	0.590	0.243	0.602	0.593
Fixed Effects	None	State	Cty, Yr	None	State	Cty, Yr
T threshold	29C	29C	29C	28C	28C	28C
P threshold	42cm	42cm	42cm	50cm	50cm	50cm

Data are for US counties east of the 100th meridian, 1980-2000. Specifications 1-2 and 4-5 are estimated with long differences and 3 and 6 with an annual panel; see text for details. Specifications 1-3 use piecewise linear thresholds as chosen by the long differences model, and 4-6 use thresholds as chosen by the panel model. Regressions are weighted by 1980 county corn area (long differences) or by 1980-2000 average corn area (panel). Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

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

	(1) 1955-1995	(2) 1955-1995	(3) 1960-2000	(4) 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

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.

Table 3: Effects of climate variation on crop revenues

	Corn		Main Spring 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

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.

Table 4: Effects of climate variation on alternate adjustment margins

	(1) Corn area	(2) Corn share	(3) Farm area	(4) Num. farms	(5) Population	(6) Land Values
GDD below threshold	0.0010	0.0003***	-0.0001	-0.0002	0.0006**	0.0004
	(0.0012)	(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0005)
GDD above threshold	-0.0005	-0.0009**	0.0000	-0.0007	-0.0008	-0.0007
	(0.0038)	(0.0004)	(0.0004)	(0.0010)	(0.0015)	(0.0013)
Precip below threshold	0.0264	-0.0004	0.0037	0.0021	-0.0236**	-0.0148
	(0.0637)	(0.0034)	(0.0035)	(0.0029)	(0.0106)	(0.0127)
Precip above threshold	-0.0051	-0.0016	0.0007	-0.0013	0.0047*	0.0008
	(0.0063)	(0.0010)	(0.0007)	(0.0033)	(0.0024)	(0.0036)
Constant	-0.0130	-0.0174***	-0.0614***	-0.1836***	0.0144	0.3555***
	(0.0687)	(0.0045)	(0.0075)	(0.0157)	(0.0160)	(0.0296)
Observations	1511	1516	1523	1526	1526	1525
Mean of Dep Variable	0.075	0.002	-0.068	-0.202	0.045	0.375
R squared	0.645	0.418	0.399	0.488	0.392	0.609
Fixed Effects	State	State	State	State	State	State
T threshold	29	29	29	29	29	29
P threshold	42	42	42	42	42	42

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), difference in log county population (Column 5) and difference in log of value of agricultural land and buildings (Column 6). 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.

A Appendix - For Online Publication

A.1 Understanding changes in climate and outcomes over time

Figure A.1 plots changes in GDD 0-29C and GDD above 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 substantial outliers that might distort results (the ~ 10 points plotted as white circles in the figure), we run regressions with and without these outliers. 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 over 1980-2000, for corn-growing counties east of the 100th meridian. Left plot: changes in GDD 0-29C and GDD above 29C. Right plot: change in log corn yields and GDD above 29C. To check robustness we run the long differences regressions with and without the points shown as white circles.*

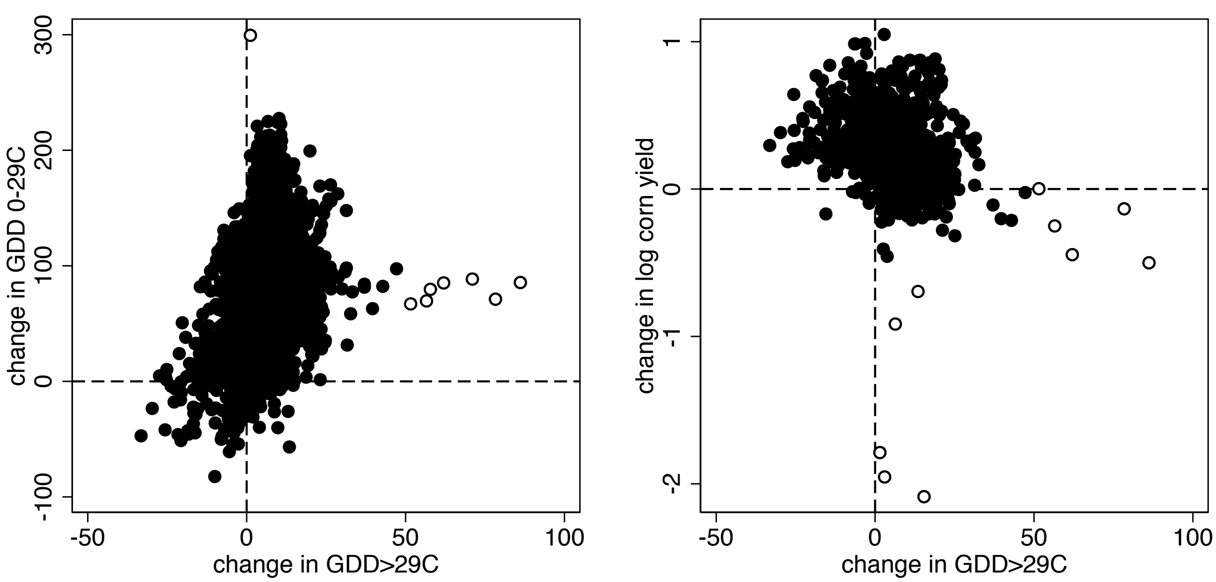

Robustness of our corn yield results to dropping outliers is explored in Table A.1. Point estimates decline slightly when outliers are dropped – not surprising given that nearly all of the outliers experienced both yield declines and large increases in exposure to extreme heat – but coefficients are statistically indistinguishable from estimates on the full sample.

Robustness of the results on alternate adaptation margins to dropping outliers is shown in Table A.2. 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.

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.*

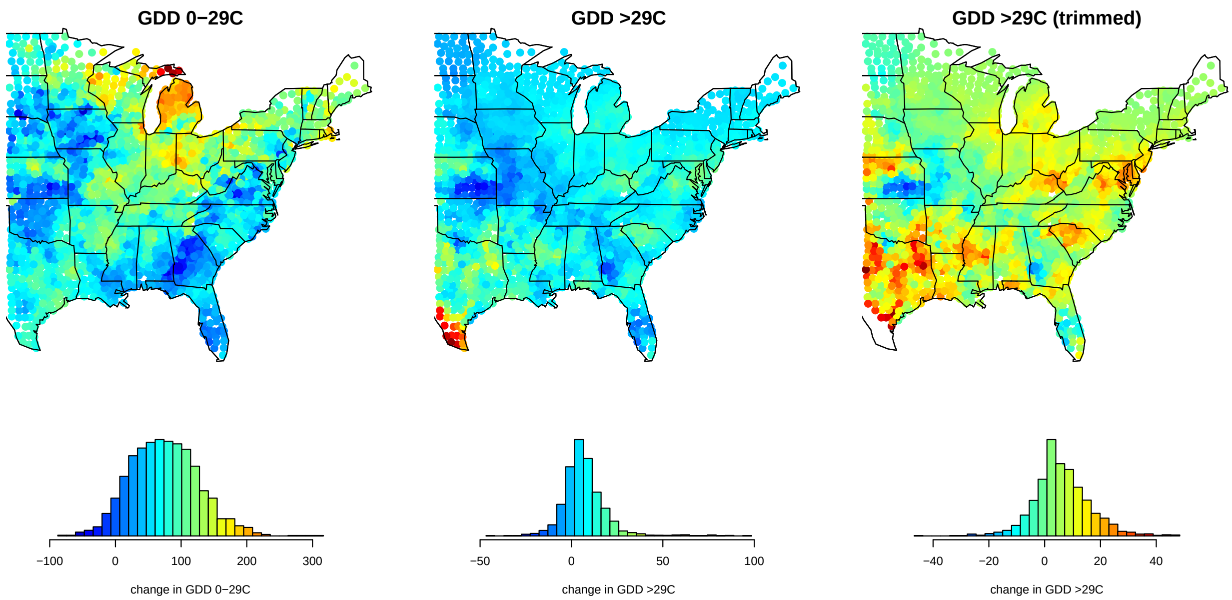

To further demonstrate the large variation in climate trends within states, and to verify that these longer-run changes do not just reflect short-run variation (e.g. single hot or cold years that are driving the difference in end-points), we estimate the trend in temperature and precipitation from 1978-2002 for each county in our main sample by running the regression

$$ln(clim_t) = \alpha + \beta t + \varepsilon_t, \tag{6}$$

where t is the sample year. Results for our main GDD>29C variable are shown in Figure A.3 for the main corn belt states. Plots represent the distribution of annual percentage changes in GDD>29C across counties within a given state (i.e. the kernel density of βs estimates in (6)), 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.

A.2 Measurement error

If fixed effects estimators are more likely than long differences estimates to suffer attenuation bias if climate variables are measured with error, then any comparison of the two estima-

Figure A.3: *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.*

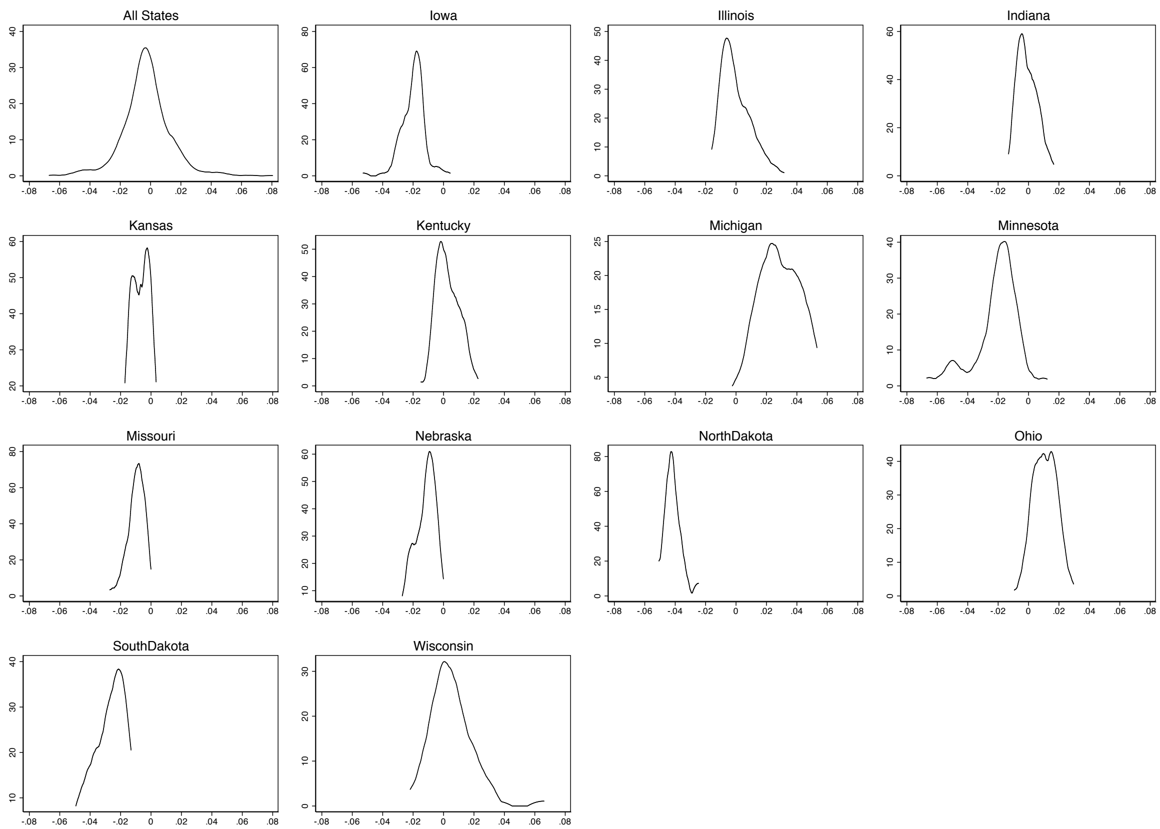

tors might understate past adaptation – panel estimates of the impact of extreme heat are less negative than they otherwise would be, causing the difference between panel and first differences to be smaller than if there were no measurement error. Following Griliches and Hausman (1986) we explore differences among panel estimators to provide insight into the extent to which fixed effects estimates are attenuated by measurement error. If measurement error in our climate variables is a problem, then Griliches and Hausman (1986) show that 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.

Table A.3 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.

Table A.1: Robustness of corn yield results to dropping outliers

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

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

A.3 Alternate measure of adaptation for corn

We verify that our main conclusions on corn yields are robust to an alternate method of estimating adaptation that does not depend on comparing long differences and panel estimates. Adaptation to a changing climate would imply that the damage caused by a current shock in climate be less severe when farmers have been exposed to more shocks in the recent past. To implement this idea, we estimated a "distributed lag" version of the panel model, where we include a set of interactions between lagged climate variables and current climate. Our panel model is now the following:

$$\begin{aligned} ln(yield)_{it} = &\sum_{j=0}^{K} \alpha_j GDDlow_{it-j} + \sum_{j=1}^{K} \alpha_j^h GDDlow_{it-j} * GDDlow_{it} + \sum_{j=0}^{K} \beta_j GDDhigh_{it-j} \\ &+ \sum_{j=1}^{K} \beta_j^h GDDhigh_{it-j} * GDDhigh_{it} + \sum_{j=0}^{K} \delta_j Preciplow_{it-j} \\ &+ \sum_{j=1}^{K} \delta_j^h Preciplow_{it-j} * Preciplow_{it} + \sum_{j=0}^{K} \gamma_j Preciphigh_{it-j} \\ &+ \sum_{j=1}^{K} \gamma_j^h Preciphigh_{it-j} * Preciphigh_{it} + \theta_i + \varepsilon_{it}. \end{aligned} \tag{7}$$

Table A.2: Robustness of results on alternate adaptations to removal of outliers

	(1) Corn area	(2) Corn area	(3) Corn %	(4) Corn %	(5) Farm area	(6) Farm area	(7) # farm	(8) # farm	(9) Land Val.	(10) Land Val.
GDD below threshold	0.0013	0.0010	0.0003***	0.0003***	-0.0000	-0.0001	-0.0001	-0.0002	0.0005	0.0004
	(0.0012)	(0.0012)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0002)	(0.0005)	(0.0005)
GDD above threshold	-0.0038	-0.0005	-0.0007**	-0.0009**	-0.0010**	0.0000	-0.0023*	-0.0007	-0.0017	-0.0007
	(0.0048)	(0.0038)	(0.0003)	(0.0004)	(0.0004)	(0.0004)	(0.0011)	(0.0010)	(0.0010)	(0.0013)
Precip below threshold	0.0404	0.0264	-0.0006	-0.0004	0.0067	0.0037	0.0077**	0.0021	-0.0103	-0.0148
	(0.0485)	(0.0637)	(0.0030)	(0.0034)	(0.0043)	(0.0035)	(0.0036)	(0.0029)	(0.0109)	(0.0127)
Precip above threshold	-0.0079	-0.0051	-0.0015	-0.0016	-0.0001	0.0007	-0.0023	-0.0013	0.0001	0.0008
	(0.0067)	(0.0063)	(0.0010)	(0.0010)	(0.0008)	(0.0007)	(0.0032)	(0.0033)	(0.0035)	(0.0036)
Constant	-0.0273	-0.0130	-0.0164***	-0.0174***	-0.0649***	-0.0614***	-0.1881***	-0.1836***	0.3500***	0.3555***
	(0.0661)	(0.0687)	(0.0045)	(0.0045)	(0.0079)	(0.0075)	(0.0156)	(0.0157)	(0.0298)	(0.0296)
Observations	1516	1511	1521	1516	1528	1523	1531	1526	1530	1525
Mean of Dep Variable	0.073	0.075	0.002	0.002	-0.069	-0.068	-0.202	-0.202	0.374	0.375
R squared	0.642	0.645	0.418	0.418	0.387	0.399	0.478	0.488	0.607	0.609
Fixed Effects	State	State	State	State	State	State	State	State	State	State
T threshold	29	29	29	29	29	29	29	29	29	29
P threshold	42	42	42	42	42	42	42	42	42	42

Dependent variable is difference in log of corn acres (Columns 1-2), difference in share of agricultural area planted to corn (Columns 3-4), difference in total log farm area (Columns 5-6), difference in log number of farmers (Columns 7-8), and and difference in log of value of agricultural land and buildings (Columns 9-10). All regressions are long differences from 1980-2000, and even numbered regressions have the sample trimmed of outliers indicated in Figure A.1. 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 A.3: Understanding measurement error through the comparison of panel estimators

	(1) Fixed Effects	(2) Random Effects	(3) First Difference
GDD below threshold	0.0004***	0.0003***	0.0003***
	(0.0001)	(0.0001)	(0.0001)
GDD above threshold	-0.0045***	-0.0042***	-0.0045***
	(0.0005)	(0.0005)	(0.0004)
Precip below threshold	0.0045**	0.0040**	0.0054***
	(0.0018)	(0.0017)	(0.0019)
Precip above threshold	-0.0011*	-0.0011	-0.0009
	(0.0006)	(0.0007)	(0.0006)
Constant	3.2154***	3.6483***	0.0703**
	(0.2881)	(0.1648)	(0.0342)
Observations	48465	48465	45405
R squared	0.463		0.494
Fixed Effects	Cty, Yr	Yr	Yr
T threshold	28	28	28
P threshold	50	50	50

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

In this specification the degree of adaptation is measured by the α_j^h, β_j^h, δ_j^h and γ_j^h terms. If farmers are adapting to higher temperatures then the effects of current unfavorable temperature should be smaller when past shocks are large, that is $\sum_{j=1}^{K} \beta_j^h$ should be large and positive. We note that this measure of adaptation will only account for "reactive" adaptations that farmers will undertake after exposure to climate change, and will not account for "proactive" adaptations that farmers could undertake to prepare for climate change before being exposed (Zilberman, Zhao, and Heiman, 2012).

In Table A.4 we report results for $K = 3$, $K = 5$ and $K = 10$. All lagged variables are standardized by dividing by the within-county standard deviation in order to ease interpretation of interaction terms. The sums of the different interaction terms and their standard errors are reported in the bottom of the table. A first noteworthy point is that the sum of interactions is only statistically significant for GDD above the threshold of 28°. This suggests that increased frequency of extreme temperatures is the only event causing reactive adaptation amongst corn farmers. The value of $\sum_{j=1}^{K} \hat{\beta}_j^h$ is of interest because it measures the *decrease* in the negative impacts of current heating when counties have been exposed to significant heating in the recent past. The interpretation of $\sum_{j=1}^{K} \hat{\beta}_j^h = 0.00036$ in Column 2 is that the effect of GDD above 28° is lower in absolute value by 0.00036 when GDD above 28° have been a standard deviation larger during each of the last 3 years. The magnitude of the sum suggests that the effect of current extreme heating is only reduced by 7.5% through adaptation to past shocks. In Column 3 when the number of lags is $K = 5$ we estimate an approximate 10.2% reduction in the effect of current heat. Finally in Column 4 with $K = 10$ our estimate is 12.5%. The small magnitude of the combined effects is consistent with the previous estimates suggesting that agricultural adaptation to climate change has been fairly minimal.

Table A.4: Panel estimates of adaptation from a distributed lag model

	(1) 0 lags	(2) 3 lags	(3) 5 lags	(4) 10 lags
GDD below threshold	0.0005***	0.0009**	0.0004	0.0004
	(0.0001)	(0.0004)	(0.0004)	(0.0006)
GDD above threshold	-0.0048***	-0.0061***	-0.0064***	-0.0068***
	(0.0005)	(0.0007)	(0.0009)	(0.0010)
Precip below threshold	0.0068***	0.0060***	0.0060***	0.0067***
	(0.0015)	(0.0018)	(0.0020)	(0.0019)
Precip above threshold	-0.0018**	-0.0016***	-0.0022***	-0.0021***
	(0.0007)	(0.0004)	(0.0005)	(0.0007)
Constant	3.2423***	1.5479	4.6367***	4.3314*
	(0.2647)	(1.3004)	(1.4339)	(2.1404)
Observations	48465	48465	48465	48465
R squared	0.593	0.619	0.635	0.651
Fixed Effects	Cty, Yr	Cty, Yr	Cty, Yr	Cty, Yr
T threshold	28C	28C	28C	28C
P threshold	50cm	50cm	50cm	50cm
Sum of lag interactions, GDD Below		-0.00001	0.00000	0.00000
SE on sum, GDD Below		0.00001	0.00001	0.00002
Sum of lag interactions, GDD Above		0.00036	0.00048	0.00059
SE on sum, GDD Above		0.00012	0.00017	0.00023
Sum of lag interactions, Precip Below		-0.00057	-0.00102	-0.00001
SE on sum, Precip Below		0.00137	0.00141	0.00147
Sum of lag interactions, Precip Above		-0.00032	0.00038	0.00030
SE on sum, Precip Above		0.00061	0.00061	0.00069

Data are from 1978-2002. Dependent variable is the log of corn yields. Individual coefficients for interaction terms not reported. All regressions are weighted by 1978-2002 average corn area. Standard errors are clustered at the state level. Asterisks indicate statistical significance at the 1% ***,5% **, and 10% * levels.

A.4 Effects on soy productivity

Estimates of the impact of extreme heat on (log) soy yields are shown in Figure A.4. 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.4: 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.

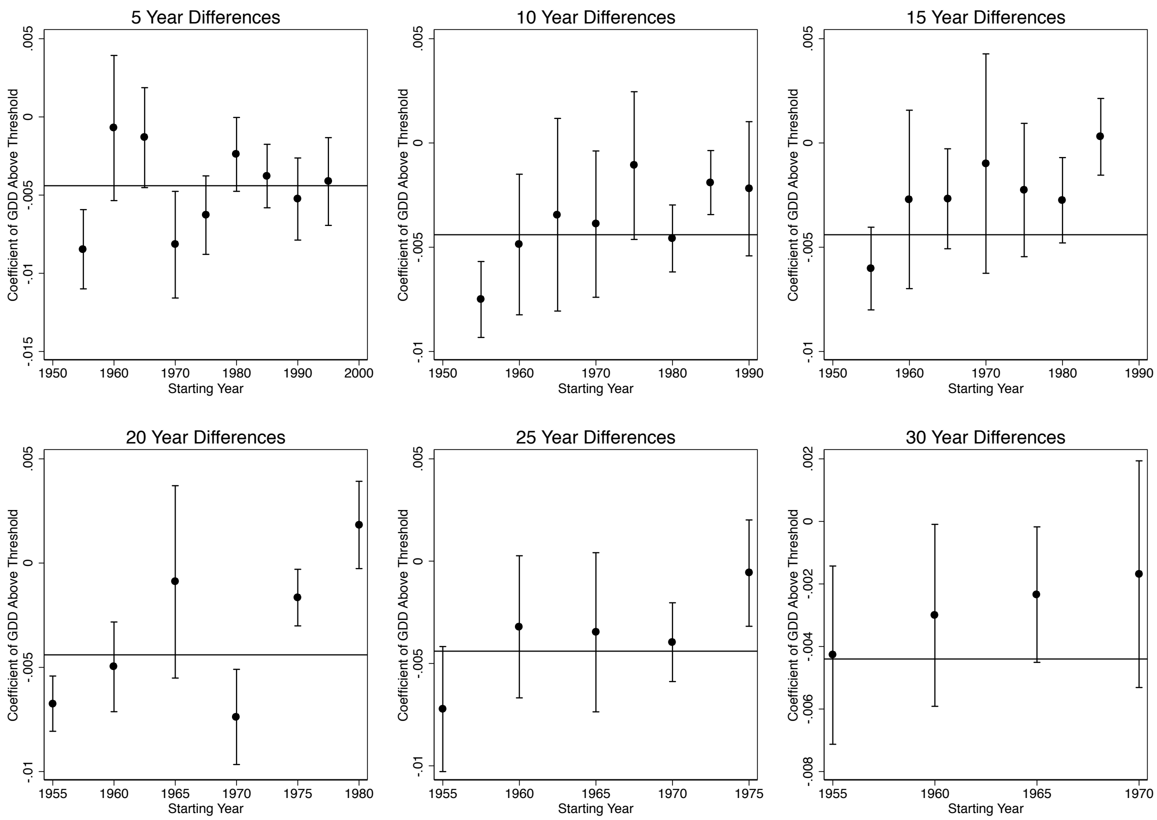

A.5 Revenues and profits

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

We choose not to construct 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).[24]

[24] In recent years, the value of own labor appears to represent about 10% of operating costs for corn, based

The second concern is that both costs and revenues will likely respond to annual variation 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.

We take two alternate approaches to exploring profitability. 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.[25] As shown in Table A.5, 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.

A.6 Exit from agriculture

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

$$ln(farms)_{ist} - ln(farms)_{is1978} = \beta_t * Extremeheat_{is} + \alpha_{st} + \varepsilon_{ist}, \tag{8}$$

where $Extremeheat_{is}$ is an indicator variable for a large change in GDD above 29. An important note is that the census defines a farm to be any place where at least $1000 in agricultural products was sold during that year. Table A.6 reports estimates with and without state-specific time fixed effects. The state specific time-effect eliminates all state-specific factors varying over time. For instance, if heating was more heavily concentrated in some states and those states had different policies over time, the state-specific time effects

on cost estimates available at http://www.ers.usda.gov/data-products/commodity-costs-and-returns.aspx. Hired labor expenditures are minimal for corn.

[25]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: Δfertilizer expenditure$_{1980-2000}$ = (fert$_{1997}$ + fert$_{2002}$)/2 - (fert$_{1978}$ + fert$_{2002}$)/2

Table A.5: Effects of Climate Variation on Input Expenditures

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

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

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.7 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

Table A.6: Estimated Differences in Log Number of Farms by Amount of Warming

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

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

data could be due to the change in the ability of the farming population that results from climate change. This possibility would become especially problematic if farmers that were more productive and had access to better quality land also had a larger opportunity cost of being in farming. If selection of this type is driving our estimates then we should see characteristics that are correlated with productivity changing differentially between places that heated and those that did not. In Table A.7 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.

Table A.7: Effects of climate variation on equipment ownership.

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

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

A.8 Non-monotonic climate trends within periods

While our long differences results on the impact of extreme heat on corn and soy yields generally appear robust to varying the starting period and period length over which differences are taken (see Figures 4 and A.4), point estimates do depend somewhat on the particular period under consideration. One explanation for this is that longer-run climate trends are not monotonic, meaning that particular periods might experience both an increase and decrease in temperature with the net result being little or no change.

In Figure A.5 we plot the annual GDD >29C averaged across our study sample. As shown, average exposure to extreme heat declined between 1980-1995, and rose quickly after that. This might help to explain why we see a positive but insignificant effect of extreme temperature exposure on soy yields between 1980-2000, but negative and significant coefficients for the 1980-1995 period when run on its own, and for the 1995-2000 period when run on its own. We speculate that this does not appear problematic in our corn results given the somewhat larger distribution of corn area across US counties and thus more overall variation in temperature trends.

Figure A.5: *GDD>29C over time, averaged across corn-growing counties in our sample. The black dotted line is the annual average, and the red line the 5-year running mean. The dashed blue line plots the trend between 1980-1994, and the solid blue line the trend 1995-2000.*

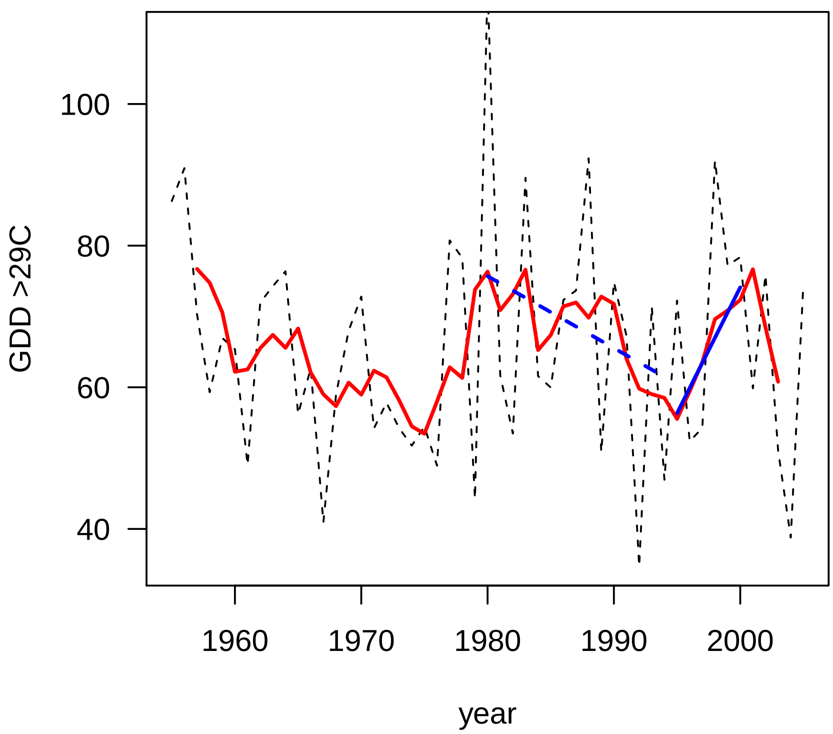

A.9 Climate change projections

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. (2011) and Burke et al. (2011). 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., 2011).

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.6. 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.6: *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.*