

Adaptation to Climate Change: Evidence from US Agriculture[†]

By MARSHALL BURKE AND KYLE EMERICK*

Understanding the potential impacts of climate change on economic outcomes requires knowing how agents might adapt to a changing climate. We exploit large variation in recent temperature and precipitation trends to identify adaptation to climate change in US agriculture, and use this information to generate new estimates of the potential impact of future climate change on agricultural outcomes. Longer run adaptations appear to have mitigated less than half—and more likely none—of the large negative short-run impacts of extreme heat on productivity. Limited recent adaptation implies substantial losses under future climate change in the absence of countervailing investments. (JEL Q11, Q15, Q51, Q54)

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

To understand how agents might adapt to a changing climate, an ideal but impossible experiment would observe two identical Earths, gradually change the climate on one, and observe whether outcomes diverged between the two. Empirical

*Burke: Department of Earth System Science, and Center on Food Security and the Environment, Stanford University, 473 Via Ortega, Stanford, CA 94305 (e-mail: mburke@stanford.edu); Emerick: Department of Economics, Tufts University, 8 Upper Campus Road, Medford, MA 02155 (e-mail: kyle.emerick@tufts.edu). We thank Max Auffhammer, Peter Berck, Melissa Dell, Richard Hornbeck, Sol Hsiang, David Lobell, Wolfram Schlenker, and seminar participants at Berkeley, Cornell, Davis, Minnesota, Stanford, UBC, the World Bank, the 2014 AAEA Annual Meeting, and the 2015 ASSA meeting for helpful comments and discussions. We thank Michael Roberts and Wolfram Schlenker for sharing data. All errors are our own.

[†]Go to <http://dx.doi.org/10.1257/pol.20130025> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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, 2011; Schlenker and Roberts 2009; Dell, Jones, and Olken 2012). Due to omitted variables concerns in the cross-sectional approach, the recent literature has preferred the latter panel approach, noting that while average climate could be correlated with other time-invariant factors unobserved to the econometrician, short-run variation in climate within a given area (typically termed "weather") is plausibly random and thus better identifies the effect of changes in climate variables on economic outcomes.

While using variation in weather helps to solve identification problems, it perhaps more poorly approximates the ideal climate change experiment. In particular, if agents can adjust in the long run in ways that are unavailable to them in the short run,¹ then impact estimates derived from shorter run responses to weather might overstate damages from longer run changes in climate. Alternatively, there could be short-run responses to inclement weather, such as pumping groundwater for irrigation in a drought year, that are not tenable in the 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 United States 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 percent over the same period. We adopt a "long differences" approach and model county-level changes in agricultural outcomes over time as a function of these changes in temperature and precipitation, accounting for time-invariant unobservables at the county level and time-trending unobservables at the state level.

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

¹For example, 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.




FIGURE 1. CHANGE IN TEMPERATURE ($^{\circ}\text{C}$), PRECIPITATION (*percent*), AND LOG CORN YIELDS OVER THE PERIOD 1980–2000

Notes: Values are shown for counties east of the 100th meridian. Temperature and precipitation are measured over the main April–September growing season. Shading of each map corresponds to shading of the bins in the histogram beneath the plot.

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 (86°F) reduces yields at the end of the season by about half a percent, which is a large effect.² The magnitude of this effect is the net of any adaptations made by farmers over the 20-year estimation period, and is robust to using different time periods and differencing lengths.

To quantify the magnitude of any yield-stabilizing adaptations that have occurred, we then compare these long differences estimates to panel estimates of short-run responses to weather. Long-run adaptations appear to have mitigated less than about half of the short-run effects of extreme heat exposure on corn yields, and point estimates across a range of specifications suggest that long-run adaptations have more likely offset *none* of these short-run impacts. We also show limited evidence for adaptation along other margins within agriculture: revenues are similarly harmed by extreme heat exposure, and farmers do not appear to be substantially altering the inputs they use nor the crops they grow in response to a changing climate.

We then examine different explanations for why adjustment to recent climate change has been minimal. For instance, adaptation could be limited because there are few adjustment opportunities to exploit, or alternatively because farmers don't recognize that climate has in fact changed and that adaptation is needed. Which explanation prevails is important for how we interpret our results, and in particular

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

how they extrapolate to future warming scenarios. If farmers failed to adapt in the past because they did not recognize the climate was changing, but in the future they become aware of these changes and quickly adapt, then our findings would be a poor guide to future impacts of warming. On the other hand, if farmers had recognized the need for adaptation but were unable to do so, then their past responses to extreme heat exposure would provide a plausible “business-as-usual” benchmark for the impacts of future warming in the absence of unprecedented adaptation.

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

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

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

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 1930s American Dust Bowl to show that negative environmental shocks can have substantial and lasting effects on productivity. Using data from a more recent period, we examine responsiveness to a slower moving environmental “shock” that is very representative of what future climate change will likely bring. Similar to Hornbeck (2012), we find limited evidence that agricultural productivity has adapted to these environmental changes, with fairly negative implications for the future impacts of climate change on the agricultural sector.

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

I. Model and Empirical Approach

Agriculture is a key sector where future climate change is estimated to have large detrimental effects, and is a primary focus of the empirical literature on climate change impacts. To formalize the ways in which our identification of climate impacts differs from that of past literature, we develop a simple model of farmer adaptation, building on earlier work by Kelly, Kolstad, and Mitchell (2005). The

climate literature generally understands adaptation as any adjustment to a changing environment that exploits beneficial opportunities or moderates negative impacts.⁴ Adaptation thus requires an agent to recognize that something in her environment has changed, to believe that an alternative course of action is now preferable to her current course, and to have the capability to implement that alternative course.

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

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

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

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

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

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

We incorporate climate change as a shift in mean temperature from $\omega \rightarrow \omega'$, with $\omega < \tilde{\omega} < \omega'$. In keeping with evidence from climate science (see Meehl et al. 2007), we assume that this increase in mean is not accompanied by a change in variance, such that after climate change the farmer experiences $z_{it} \sim N(\omega', \sigma^2)$ in each year. A fully informed farmer recognizes this change and immediately adopts the heat-tolerant crop, which we consider “adaptation.” In reality, farmers likely learn about changes in climate over time and only adjust behavior after acquiring strong enough information that climate has changed. Following Kelly, Kolstad, and Mitchell (2005), we assume this learning follows a simple Bayesian process where the farmer has a

⁴ See Zilberman, Zhao, and Heiman (2012) and Burke and Lobell (2010) for an overview.

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




FIGURE 2. PRODUCTIVITY OF TWO DIFFERENT CORN VARIETIES AS A FUNCTION OF TEMPERATURE

prior belief about ω_t but knows that this belief is imperfect. We denote the belief as μ_t and its variance as $1/\tau_t$, such that in period t the farmer believes $\omega \sim N(\mu_t, 1/\tau_t)$. In each period she observes z_{it} and updates her belief about the average temperature to μ_{t+1} using a weighted combination of her prior belief and the new climate realization she experiences. Letting $\rho = 1/\sigma^2$ and \bar{z} the average of the temperature realizations during the previous $T - 1$ years, the farmer’s belief about mean climate after T years is given by DeGroot (1970):

$$(3) \quad \mu_T = \frac{\tau_t \mu_t + T\rho \bar{z}}{\tau_t + T\rho}.$$

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

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

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

Our model can be extended to allow a richer learning environment where a farmer learns about the temperature in her own county from weather realizations in both

her own county and neighboring counties. Extending the model in this way has two implications. First, the farmer will learn about a vector of average temperatures where the vector of beliefs is a (weighted) linear combination of the vector of prior beliefs and the vector of means of realized weather.⁶ Second and more importantly, in settings where mean temperatures are highly correlated across nearby counties but annual weather realizations are less correlated, the ability to learn from other counties can cause the belief about one's own county to converge faster to the true mean temperature—although as we show using simulation in online Appendix A.1, the benefits of this additional information appear relatively modest. In Section III E, we nevertheless extend our empirical approach to test for whether observation of temperature in other counties in the same state accelerated farmer learning and adaptation.

A. Existing Approaches

Returning to Figure 2, the long-term damages imposed by a shift in climate will be $v_0 - v_1$ if adaptation takes place.⁷ Past literature has taken two approaches to estimating this quantity. In pioneering work, Mendelsohn, Nordhaus, and Shaw (1994) use cross-sectional variation in average temperature and precipitation (and their squares) to explain variation in agricultural outcomes across US counties. The cross-sectional specification is

$$(5) \quad y_i = \alpha + \beta_1 w_i + \beta_2 w_i^2 + c_i + \varepsilon_i,$$

where y_i is some outcome of interest in county i , w_i is again the average temperature, and c_i other time invariant factors affecting outcomes (such as soil quality). Mendelsohn, Nordhaus, and Shaw's (1994) preferred dependent variable is land values, which represent the present discounted value of the future stream of profits that could be generated with a given parcel of land, and thus in principle embody any possible long-run adaptation to average climate. Therefore, a county with average temperature of ω will achieve v_0 on average, a county with average temperature of ω' will achieve v_1 , and the estimates of β_1 and β_2 along with a projected rise in average temperatures from ω to ω' would seem to identify the desired quantity of $v_0 - v_1$.

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

⁶Denoting μ_T as the farmer's vector of beliefs after T years, μ as the vector of prior beliefs, \bar{z} as the vector of mean realizations, Φ as the precision matrix of the prior beliefs, and Ω as the precision matrix of the annual observations of weather, the farmer's vector of beliefs after T periods is (DeGroot 1970): $\mu_T = (\Phi + T\Omega)^{-1} (\Phi\mu + T\Omega\bar{z})$.

⁷Kelly, Kolstad, and Mitchell (2005) call this the "equilibrium response," in contrast to the costs incurred when undertaking adaptation (e.g., the purchase of a more expensive heat-tolerant variety), which they term "adjustment costs."

the relationship between agricultural outcomes and variation in temperature and precipitation (Deschênes and Greenstone 2007; Schlenker and Roberts 2009; Welch et al. 2010; Lobell, Schlenker, and Costa-Roberts 2011).⁸ The data generating process in this approach is:

$$(6) \quad y_{it} = \alpha + \beta_1 z_{it} + \beta_2 z_{it}^2 + c_i + \varepsilon_{it}.$$

All time invariant factors are absorbed by the location fixed effects c_i , and impacts of temperature and precipitation on (typically annual) outcomes are thus identified from deviations from location-specific means.⁹ Because this year-to-year variation in temperature and precipitation (typically termed “weather”) is plausibly exogenous, fixed effects regressions overcome omitted variables concerns with cross-sectional models, and the effect of temperature on outcomes such as yield or profits can be interpreted causally.

Many studies then combine the estimated short-run responses from panel regressions with output from global climate models to project potential impacts under future climate change.¹⁰ In making these projections, the implicit assumption is again that short-run responses to variation in weather are representative of how farmers will respond to longer run changes in average climate. It is not obvious this will be the case. Consider a panel covering many years, with a temperature rise from ω to ω' occurring somewhere within these years. The panel model would identify movement along either one of the two curves shown in Figure 2, with the point estimate being a weighted average of the slopes of the two curves, with weights depending on if and when the varietal switch occurred. If the heat-tolerant crop is adopted at the end of the period, then fixed effects estimates will be heavily weighted towards the curve for the conventional crop, overstating equilibrium losses. If adaptation is instantaneous, then fixed effects estimates trace out the curve for the heat-tolerant crop, which could understate impacts if (as drawn) the slope of the response function is positive at ω' . Thus, estimates of short-run responses to weather will not even bound estimates of longer run response to climate. Panel models therefore solve identification problems in the cross-sectional approach, at the cost of more poorly approximating the idealized climate change experiment.

⁸Examples 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).

⁹McIntosh 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 nonlinearly 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.

¹⁰See Burke et al. (2015) for a review of these studies and for the use of global climate models in this context.

B. The Long Differences Approach

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

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

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

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

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

$$(9) \quad \Delta \overline{y}_i = \beta_1 \Delta \overline{z}_i + \beta_2 \Delta(\overline{z}_i)^2 + \Delta \varepsilon_i,$$

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

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

We then use this strategy to quantify the extent of recent adaptation in US agriculture, comparing our long differences estimates to those from an annual panel model. We would interpret more positive long difference estimates as evidence of adaptation: that farmers are better able to adjust to longer run changes in climate than they are to shorter run changes in weather. In Figure 2, if any adaptation takes place, the

long differences approach should identify $v_0 - v_1$. If no adaptation occurs, then long difference regressions will identify $v_0 - v_2$, i.e., the same damages identified by the panel model. We attempt to rule out other explanations for divergence between panel and long-differences estimates—e.g., measurement error, or adaptation outside of agriculture—in Section III.

C. 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.¹¹ 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 one hundredth 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 percent of US corn production and 99 percent of US soy production.

Our climate data are drawn from Schlenker and Roberts (2009) and consist of daily interpolated values of precipitation totals and maximum and minimum temperatures for 4 kilometer (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.¹²

Past literature has demonstrated strong nonlinearities in the relationship between temperature and agricultural outcomes (e.g., Schlenker and Roberts 2009). Such nonlinearities are generally captured using the concept of growing degree days (GDD), which measure the amount of time a crop is exposed to temperatures between a given lower and upper bound. Following Schlenker and Roberts (2009), we use the within-day distribution of temperatures to calculate the percent of each day that cropped area in each county is exposed to temperatures between given lower and upper bounds, and then sum these daily exposures over a fixed growing season (April 1 to September 30th) to get a measure of annual growing degree days for those bounds.

Using this notion of GDD, and using the county agricultural data described above, we model agricultural outcomes as a simple piecewise linear function of temperature and precipitation.¹³ We estimate the long differences model:

¹¹ We thank Michael Roberts for sharing additional census data that are not yet archived online.

¹² We thank Wolfram Schlenker for sharing the weather data and the code to process them.

¹³ 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

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

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 five-year averages to more effectively capture the change in average climate or outcomes over time. That is, for the 1980–2000 period we take averages for each variable over 1978–1982 and over 1998–2002, and difference these two averages.

The lower temperature “piece” in (10) is the sum of GDD between the bounds l_0 and l_1 , and $\Delta GDD_{is;l_0:l_1}$ 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.¹⁴ We also 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.¹⁵ In the estimation we set $l_0 = 0$ and allow the data to determine l_1 and p_0 by looping over all possible thresholds and selecting the model with the lowest sum of squared residuals.

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

There are two potential concerns with our empirical approach. The first is that the inclusion of state fixed effects could absorb most of the variation of interest in our temperature variables. Second, our supposed differential trends in temperature

(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 is no longer limited to location-specific variation over time.

¹⁴ As an example, if $l_0 = 0$ and $l_1 = 30$, then a given set of daily average temperatures of $-1, 0, 1, 10, 29, 31,$ and 35 would result in $GDD_{is;l_0:l_1}$ equal to $0, 0, 1, 10, 29, 30,$ and 30 , and $GDD_{is;l_1:\infty}$ equal to $0, 0, 0, 0, 1,$ and 5 . In practice, we use the within-day distribution of temperatures to allow fractions of days to be spent above or below a given threshold, but the principle is the same.

¹⁵ 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 millimeter (mm) in 1980 to 50 mm in 2000. If the precipitation threshold is 40 mm, then $\Delta Prec_{is;p < p_0} = 5$ and $\Delta Prec_{is;p > p_0} = 10$.

across counties could just be driven by short-run variation in weather around the chosen endpoint years. This has two potential implications. First, there could have been little “true” long-run change in temperature to adapt to. Second, even if temperature was trending differentially across counties, our long differences estimator might mechanically deliver estimates that are similar to those from a county-year panel if interannual variation in temperature around this trend (i.e., weather) was large. In Section A2 of the online Appendix, we explore each of these issues in more detail. We first demonstrate that the residual variation in our temperature changes of interest remains large (even relative to projected future changes) after accounting for state fixed effects, and that this variation very likely represents true long-run increases in temperature rather than large variation in chosen endpoint years. Second, we clarify under what conditions the panel and long differences approaches will differentially identify “true” long-run and short-run responses to changes in temperature, and demonstrate via simulation that our long differences estimates are likely based on agents’ responses to longer run changes in temperature.

Figure 1 displays the variation that is used in our identification strategy. Some US counties have cooled slightly over the past three decades, while others have experienced warming equivalent to over 1.5 times the standard deviation of local temperature—roughly equivalent to the mean warming projected by global climate models to occur over US corn area by 2030. Differential trends in precipitation over the 1980–2000 period have been similarly large, with precipitation decreasing by more than 30 percent in some counties and increasing by 30 percent in others—a range that again almost fully contains the range in climate model projections of future precipitation change over the same area by the mid-21st century. Substantial variation is apparent even within states. For instance, Lee County in the southeastern Iowa experienced an increase in average daily temperature during the main corn growing season of 0.46°C , and Mahaska County—approximately 80 miles to the northwest—experienced a decrease in temperature of 0.3°C over the same period. Corn yields in parts of northern Kentucky declined slightly while rising by 20–30 percent only 100 miles to the south.

Importantly, there remains large variation in our main regressor of interest (exposure to extreme heat) even after conditioning on other climate variables and state fixed effects (see Table A.1 in the online Appendix), and as shown in online Appendix Figure A.4, this variation substantially overlaps projections from global climate models of future changes in extreme heat exposure. This means that applying our estimates to predictions from climate models is not asking our model to extrapolate far out of sample.

While we explore robustness of our results to different time periods and differencing lengths, we focus on the post-1980 period for a number of reasons. First, warming trends since 1980 were much larger than in earlier periods. For instance, over the 1960–1980 period, only half of the counties in our sample experienced average warming, and none experienced warming of more than 1°C (see Figure A.11 in the online Appendix). Second, recognition of climate change was much higher in this later period, which helps alleviate some concerns that a lack of recognition of climate change is what is driving our results. In particular, prior to 1980 there was even significant scientific and popular concern about the risks from “global cooling”

(e.g., Gwynne 1975), and only during the 1980s and 1990s was there growing recognition that the climate was warming and that increasing greenhouse gas emissions meant there would very likely be further warming in the future.

D. Are Recent Climate Trends Exogenous?

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

The effect of aerosols on crops is less well understood (Auffhammer, Ramanathan, and Vincent 2006). While any indirect effect through temperature or precipitation will already be picked up in the data, aerosols become an omitted variables concern if their other influence on crops—namely their effect on solar radiation—have important effects on crop productivity. Because crop productivity is generally thought to be increasing and concave in solar radiation, reductions in solar radiation are likely to be harmful, particularly to C_4 photosynthesis plants like corn that do not become light saturated under typical conditions.¹⁶ However, aerosols also increase the “diffuse” portion of light (think of the relatively even light on a cloudy day), which allows additional light to reach below the canopy, increasing productivity. A recent modeling study finds negative net effects for corn, with aerosol concentrations (circa the year 2000) reducing corn yields over the Midwest by about 10 percent, albeit with relatively large error bars. This would make it likely that, if anything, aerosols will cause us to understate any negative effect of warming on crop yields: aerosols lead to both cooling (which is generally beneficial in our sample) and to a reduction in solar radiation (which on net appears harmful for corn). In any case, the inclusion of state fixed effects means that we would need significant within-state variation in aerosol emissions for this to be a concern.

The second main omitted variables concern is changes in local land use. Evidence from the physical sciences suggests that conversion between types of land (e.g., conversion of forest to pasture, or pasture to cropland), or changes in management practices within pre-existing farmland (e.g., expansion of irrigation) can have significant

¹⁶Crops 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).

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

II. Empirical Results

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

A. 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 (42 centimeter (cm) and 50 cm). In columns 1–4 we run both models under the thresholds selected by the long differences, and in columns 5–8 we fix thresholds at values chosen by the panel.

The panel and long differences models deliver very similar estimates of the responsiveness of corn yields to temperature. Exposure to GDD below 29°C (row 1)

TABLE 1—COMPARISON OF LONG DIFFERENCES AND PANEL ESTIMATES OF THE IMPACTS OF TEMPERATURE AND PRECIPITATION ON US CORN YIELDS

	Diff's (1)	Diff's (2)	Panel (3)	Panel (4)	Diff's (5)	Diff's (6)	Panel (7)	Panel (8)
GDD below threshold	-0.0001 (0.0003)	0.0002 (0.0002)	0.0004 (0.0001)	0.0002 (0.0001)	-0.0001 (0.0003)	0.0003 (0.0002)	0.0005 (0.0001)	0.0003 (0.0001)
GDD above threshold	-0.0053 (0.0010)	-0.0044 (0.0008)	-0.0056 (0.0006)	-0.0062 (0.0007)	-0.0043 (0.0009)	-0.0037 (0.0009)	-0.0048 (0.0005)	-0.0054 (0.0006)
Precip below threshold	0.0515 (0.0194)	0.0297 (0.0125)	0.0118 (0.0027)	0.0095 (0.0048)	0.0253 (0.0123)	0.0115 (0.0046)	0.0068 (0.0015)	0.0057 (0.0019)
Precip above threshold	0.0036 (0.0017)	0.0034 (0.0008)	-0.0008 (0.0005)	0.0001 (0.0004)	0.0024 (0.0015)	0.0029 (0.0007)	-0.0018 (0.0007)	-0.0008 (0.0005)
Constant	0.2655 (0.0319)	0.2397 (0.0124)	3.5721 (0.2491)	4.1872 (0.3013)	0.2674 (0.0307)	0.2400 (0.0115)	3.2423 (0.2647)	3.8577 (0.3349)
Observations	1,531	1,531	48,465	48,465	1,531	1,531	48,465	48,465
R ²	0.258	0.610	0.590	0.863	0.243	0.602	0.593	0.864
Fixed effects	None	State	Cty, Yr	Cty, State-Yr	None	State	Cty, Yr	Cty, State-Yr
T threshold	29°C	29°C	29°C	29°C	28°C	28°C	28°C	28°C
P threshold	42 cm	42 cm	42 cm	42 cm	50 cm	50 cm	50 cm	50 cm

Notes: Data are for US counties east of the 100th meridian, 1980–2000. Specifications 1–2 and 5–6 are estimated with long differences and 3–4 and 7–8 with an annual panel, with different fixed effects shown at bottom; see text for details. Specifications 1–4 use piecewise linear thresholds as chosen by the long differences model, and 5–8 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.




FIGURE 3. RELATIONSHIP BETWEEN TEMPERATURE AND CORN YIELDS

Notes: 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.

have small and generally insignificant effects on yields, but increases in exposure of corn to temperatures above 29°C result in sharp declines in yields, as is seen in the second row of the table and in Figure 3. In our most conservative specification with state fixed effects, exposure to each additional degree-day of heat above 29°C results in a decrease in overall corn yield of 0.44 percent.¹⁷ The panel model delivers a slightly more negative point estimate, a -0.56 percent yield decline for every one degree increase above 29°C, but (as quantified below) we cannot reject that the estimates are the same. We obtain similar results when the two models are run under the temperature and precipitation thresholds chosen by the panel model (columns 5–8), and similar results when the panel model is estimated with state-by-year fixed effects rather than year fixed effects.

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

To test robustness of the corn results, we show in the remainder of this subsection that our results are relatively insensitive to the choice of endpoint years, to the number of years used to calculate endpoints, and to an alternate estimation strategy that further weakens our identification assumptions. In online Appendix A.4, we provide further evidence that our results are insensitive to the exclusion of yield and temperature outliers, and to the inclusion of baseline covariates in the regression.

We first show that our results are largely unchanged when we change the time period under study. In particular, we estimate equation (10) varying T_0 from 1955 to 1995 in five year increments, and for each value of T_0 we estimate 5, 10, 15, 20,

¹⁷ While we prefer the more conservative specifications with state fixed effects in Table 1, one concern with the inclusion of state fixed effects is that farmer responses to increasing temperature might vary meaningfully at the state level, for instance if governments in states that experienced substantial warming helped their farmers invest in adaptation measures. These policies would be absorbed by the state fixed effects, and could obscure meaningful adaptation measures undertaken by farmers. Our results appear inconsistent with a story of state-level adaptation to extreme heat. The effects of extreme heat in the specifications without state fixed effects (columns 1 and 4) are substantially more negative than in the comparable specifications with state fixed effects, which is the opposite of what would be expected if state-level adaptation policies were an omitted variable in these regressions.




FIGURE 4. SENSITIVITY OF RESULTS TO STARTING YEAR AND LENGTH OF DIFFERENCING PERIOD

Notes: Differences between main estimate from 1980–2000 (specification 2 in Table 1) and other estimates under various starting years and differencing lengths. Dots are differences in estimates and whiskers are 95 percent confidence intervals of the differences. Standard errors are calculated by bootstrapping, where 1,000 samples of 31 states were drawn (with replacement) and the difference between estimates was calculated for each bootstrapped sample.

25, and 30-year difference models.¹⁸ Results are shown graphically in Figure 4. We display the difference between the estimate of β_2 for 1980–2000 (our baseline estimate) and the estimate of β_2 for the period determined by the starting year and differencing length. The 95 percent confidence intervals of the differences are calculated by bootstrapping.¹⁹ The average estimate of β_2 across these 39 models is -0.0058 , with only 8 of the estimates of β_2 being statistically different from our main 1980–2000 estimate and none statistically different in the positive direction. This suggests if anything, that our baseline point estimate on the effect of extreme heat is conservative.²⁰ We conduct an analogous exercise for the panel model to make sure that the effect of extreme heat in the panel does not vary with the chosen

¹⁸ Some models of course could not be estimated since our data end at 2005, meaning our five-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 five year average corn area during the starting year. The temperature and precipitation thresholds are fixed at 29°C and 42 cm across models.

¹⁹ We drew 1,000 samples of 31 states with replacement and estimated all regressions for each sample. The differences between the 1980–2000 estimate and all other possible estimates were calculated for each sample. The bootstrapped standard errors are the standard deviations of the differences in estimates.

²⁰ In online Appendix Figure A.10, we display the raw coefficients and their confidence intervals for each period: all estimates are negative, and in only 8 out of 39 cases do we fail to reject a significant negative effect of extreme heat on corn productivity.

time period. Results are plotted in online Appendix Figure A.12, and agree with earlier findings in Schlenker and Roberts (2009) that the effects of interannual deviations in extreme heat have not declined significantly over time.²¹

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

Finally, our estimates in equation (10) would be biased in the presence of within-state time-varying unobservables correlated with both climate and yields. To address this possibility, we use our many decades of data to construct a two period panel of long differences, which further weakens our identification assumption. We estimate the following model:

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

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

In Table 2 we report estimates from both the 1955–1995 period and the 1960–2000 period. In all models the effect of temperature above 29°C remains negative and

²¹ While this unchanging sensitivity of yield to extreme heat over time could be interpreted as additional evidence of a lack of adaptation (as in Schlenker and Roberts 2009), we note that whether responses to short-run variation have changed over time is conceptually distinct from whether farmers have responded to long-run changes in average temperature. As emphasized in our conceptual framework, there is no reason to expect farmers to respond similarly to these two different types of variation. Indeed, farmers could adapt completely to long-run changes in temperature such that average yields do not change—e.g., by adopting a new variety that on average performs just as well in the new expected temperature as the old variety did under the old average temperature—but still face year-to-year variation in yield due to random deviations in temperature about its new long-run average. As such, we view this exercise more as a test of the robustness of the panel model than as evidence of (a lack of) adaptation per se.

TABLE 2—THE EFFECT OF CLIMATE ON YIELDS ESTIMATED WITH A PANEL OF DIFFERENCES

	1955–1995 (1)	1955–1995 (2)	1960–2000 (3)	1960–2000 (4)
GDD below threshold	0.0008 (0.0003)	0.0007 (0.0004)	0.0004 (0.0001)	0.0003 (0.0002)
GDD above threshold	−0.0066 (0.0013)	−0.0058 (0.0020)	−0.0031 (0.0007)	−0.0023 (0.0010)
Precip below threshold	0.0356 (0.0079)	0.0376 (0.0093)	0.0203 (0.0135)	0.0166 (0.0115)
Precip above threshold	0.0017 (0.0015)	0.0033 (0.0017)	0.0008 (0.0015)	0.0014 (0.0020)
Observations	2,060	2,060	2,604	2,604
R^2	0.621	0.565	0.688	0.699
Fixed effects	State Yr	Cty Yr	State Yr	Cty Yr
T threshold	29°C	29°C	29°C	29°C
P threshold	42 cm	42 cm	42 cm	42 cm

Notes: 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.

significant even after the inclusion of county fixed effects. The main coefficients for $GDD > 29$ are also similar to our baseline estimates in Table 1. The main long differences estimates are therefore robust to controlling for a richer set of county-specific time-varying factors.

B. 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 percent 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 1,000 times (sampling states with replacement to account for spatial correlation) and recalculate $1 - \beta_2^{LD}/\beta_2^{FE}$ for each iteration.²² 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 then allow us to test, for each time period of interest, the null hypothesis of “no adaptation” to extreme heat—i.e., that $1 - \beta_2^{LD}/\beta_2^{FE} = 0$.

Results are shown in Figure 5 and suggest that, on the whole, longer run adaptation to extreme heat in corn has been limited. Median estimates from each distribution all

²²That 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 1,000 times for a given time period.




FIGURE 5. PERCENTAGE OF SHORT-RUN IMPACTS OFFSET BY ADAPTATION

Notes: Figure 5 shows the percentage of the short-run impacts of extreme heat on corn productivity that are mitigated in the longer run. Each box plot corresponds to a particular time period as labeled, and represent 1,000 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 fifth to ninety-fifth percentile. The distribution plotted at bottom represents the combination of all the estimates in the above distributions. The values at left give the one-sided p -value on the test that $1 - \beta_2^{LD} / \beta_2^{FE} = 0$.

indicate that adaption has offset less than 25 percent of short run impacts—and point estimates are actually slightly *negative* in two-thirds of the cases. In almost all cases we can conclude that adaptation has offset at most half of the negative shorter run impacts of extreme heat on corn yields. Finally, all confidence intervals span zero, meaning we can never reject that there has been no more adaptation to extreme heat in the long run than has been in the short run (one-sided p -values on the test of the null against the alternative hypothesis that $1 - \beta_2^{LD} / \beta_2^{FE} > 0$ are $p = 0.14$ or greater, as shown on the left of the figure).

C. Soy Productivity

All of our analysis up to this point has focused on corn, the dominant field crop in the United States 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 online Appendix Figure A.14, 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°C 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 percent. 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.

III. Alternate Explanations

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

A. Measurement Error

A key concern with fixed effects estimates of the impact of climate variation is attenuation bias caused by measurement error in climate variables. Fixed effects estimates are particularly susceptible to attenuation since they rely on short-term deviations from average climate to identify coefficients. This makes it more difficult to separate noise from true variation in temperature and precipitation compared to a setting where identification comes from relatively better measured averages over space or time (such as in our long differences results). Therefore one explanation for the limited observed yield adaptation is simply that panel estimates are attenuated relative to long differences estimates, and thus that 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 mismeasurement 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.²³

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. Online Appendix Table A.7 shows that estimates from all three estimators are remarkably similar, providing suggestive evidence that measurement error is not responsible for the similarity between fixed effects and long differences estimates.

B. Selection

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

C. Adaptation along Other Margins

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

²³This result holds so long as the measurement error for temperature and precipitation is uncorrelated with the “true” temperature and precipitation values—i.e., that both exhibit classical measurement error—but does not require the temperature and precipitation errors to be uncorrelated. We have verified this via simulation, with results available upon request.

One way to capture broader economic adjustment to changes in climate is to explore climate impacts on farm revenues or profits, an approach adopted in some of the recent literature (e.g., Deschênes and Greenstone 2007). There are at least two empirical challenges with using profits in particular. The first is that measures of revenues and expenditures are only available every five years when the US Agricultural Census is conducted. Given that our differencing approach seeks to capture change in average farm outcomes over time, if both revenues and costs respond to annual fluctuations in climate, then differencing two “snapshots” from particular years might provide a very noisy measure of the longer term change in profitability. A second concern is that available data on expenses do not measure all relevant costs (e.g., the value of own or family labor on the farm), which might further bias profit estimates if these expenses also respond to changes in climate. As shown in online Appendix Section A.9, long differences regressions with such a measure of “profit” as the dependent variable are indeed very noisy, and we cannot reject that there is no effect on profits, and similarly cannot reject that the effect of extreme heat on profits is a factor of three larger (and more negative) than the effect on corn yields—i.e., that each additional day of exposure to temperatures above 29°C reduced annual profits by 1.4 percent. This does not provide much insight on the relationship between extreme heat exposure and profitability.

We take two alternate approaches to exploring impacts on economic profitability. The first is to construct an annual measure of revenue per acre, which we do by combining annual county-level yield data with annual data on state-level prices.²⁴ 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.²⁵

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.²⁶ Furthermore, we show in online Appendix Table A.8 that trends in

²⁴ Prices are only available at the state level and to our knowledge do not vary much within states within a given year.

²⁵ We attempt to capture changes in average expenditures by averaging two census outcomes near each end-point and then differencing these averaged values. For example, agriculture census data are available in 1978, 1982, 1987, 1992, 1997, and 2002. The change in fertilizer expenditures over the period are constructed as: $\Delta \text{fertilizer expenditure}_{1980-2000} = (\text{fert}_{1997} + \text{fert}_{2002})/2 - (\text{fert}_{1978} + \text{fert}_{1982})/2$.

²⁶ 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

TABLE 3—EFFECTS OF CLIMATE VARIATION ON CROP REVENUES

	Corn		Main spring crops	
	Panel (1)	Diff. (2)	Panel (3)	Diff. (4)
GDD below threshold	0.0005 (0.0001)	0.0003 (0.0002)	0.0002 (0.0001)	0.0003 (0.0003)
GDD above threshold	-0.0046 (0.0005)	-0.0042 (0.0009)	-0.0024 (0.0003)	-0.0023 (0.0011)
Precip. below threshold	0.0068 (0.0016)	0.0107 (0.0048)	0.0058 (0.0014)	0.0116 (0.0058)
Precip. above threshold	-0.0014 (0.0007)	0.0035 (0.0010)	-0.0012 (0.0005)	0.0016 (0.0016)
Constant	3.9556 (0.2539)	-0.0116 (0.0122)	4.7926 (0.3619)	0.0121 (0.0210)
Observations	48,465	1,516	48,465	1,531
Mean of dependent variable	5.55	-0.01	5.36	0.03
R^2	0.568	0.579	0.490	0.454
Fixed effects	Cty, Yr	State	Cty, Yr	State

Notes: 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 six main crops grown during the spring season (corn, soy, cotton, spring wheat, hay, and rice). Revenues are 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°C and precipitation threshold is 50 cm in all regressions. Standard errors are clustered at the state level.

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.²⁷ 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

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 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.

²⁷ As shown in online Appendix Table A.5—and unlike for our yield outcomes—a few outcomes in this table are altered fairly substantially when these five outliers (0.3 percent of the sample) are included. Given that these counties are all geographically distinct (along the Mexico border in southern Texas), and experienced up to 20 times the average increase in exposure to extreme heat than our median county in the sample, it seems reasonable to exclude them from the analysis.

TABLE 4—EFFECTS OF CLIMATE VARIATION ON ALTERNATE ADJUSTMENT MARGINS

	Corn area (1)	Corn share (2)	Farm area (3)	Num. farms (4)	Population (5)
GDD below threshold	0.0010 (0.0012)	0.0003 (0.0001)	-0.0001 (0.0001)	-0.0002 (0.0002)	0.0006 (0.0003)
GDD above threshold	-0.0005 (0.0038)	-0.0009 (0.0004)	0.0000 (0.0004)	-0.0007 (0.0010)	-0.0008 (0.0015)
Precip. below threshold	0.0264 (0.0637)	-0.0004 (0.0034)	0.0037 (0.0035)	0.0021 (0.0029)	-0.0236 (0.0106)
Precip. above threshold	-0.0051 (0.0063)	-0.0016 (0.0010)	0.0007 (0.0007)	-0.0013 (0.0033)	0.0047 (0.0024)
Constant	-0.0130 (0.0687)	-0.0174 (0.0045)	-0.0614 (0.0075)	-0.1836 (0.0157)	0.0144 (0.0160)
Observations	1,511	1,516	1,523	1,526	1,526
Mean of dep. variable	0.075	0.002	-0.068	-0.202	0.045
R^2	0.645	0.418	0.399	0.488	0.392
Fixed effects	State	State	State	State	State
T threshold	29°C	29°C	29°C	29°C	29°C
P threshold	42 cm				

Notes: Dependent variable is difference in log of corn acres (column 1), difference in share of agricultural area planted to corn (column 2), difference in total log farm area (column 3), difference in log number of farms (column 4), and difference in log county population (column 5). All regressions are long differences from 1980–2000, with the sample trimmed of extreme outliers in the temperature data. All regressions are weighted by average agricultural area from 1978–2002. Standard errors are clustered at the state level.

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 percent.²⁸

A final adaptation available to farmers would be to exit agriculture altogether, an option that recent literature has suggested is a possibility. For instance, Hornbeck (2012) shows that population decline was the main margin of adjustment across the Great Plains after the American Dust Bowl. Feng, Oppenheimer, and Schlenker (2015) 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.

²⁸We also explore whether corn planting dates, which are available at the state level for 14 states in our sample for the full study period, responded differentially in areas that warmed—which could be additional evidence of adaptation which we might miss if we’ve fixed the agricultural growing season from April 1 to September 30. Using National Agricultural Statistics Service data, we define the planting date as the week of the year in which 50 percent of that year’s crop has been planted, and find that this date moved an average of eight days earlier between 1980–2000 in our 14-state sample, consistent with overall warming allowing earlier planting. However, we estimate a slight positive but insignificant relationship between change in extreme heat (GDD > 29) and planting date in the sample, providing little conclusive evidence that planting dates shifted in response to changes in extreme heat exposure. Results available upon request.

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 percent decline in the number of farms for the counties experiencing the greatest increase in exposure to extreme heat over our main sample period.²⁹ 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 percent 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.

D. Policy Disincentives to Adapt

A fourth explanation for limited adaptation is that certain governmental agricultural support programs—subsidized crop insurance in particular—could have reduced farmers’ incentives to adapt. In the crop insurance program, the federal government insures farmers against substantial losses while also paying most or all of their insurance premiums, and this plausibly could have reduced farmers’ incentives to undertake costly adaptations.³⁰

As one check on whether observed lack of adaptation is being driven by the existence of subsidized insurance, we utilize the large-scale expansion of the federal crop insurance program in the mid-1990s and compare the impact of long-run changes in temperature before and after the expansion. This expansion, related to a set of revised government policies that were instated beginning in 1994, roughly tripled participation in the crop insurance program relative to the late 1980s, and by the end of our study period over 80 percent of farmers were participating in the program. We find that the effects of temperature in the post-expansion period were the same or even slightly smaller (in absolute value) than the effects in the pre-expansion period, which is the opposite of what would be expected if subsidized insurance had reduced farmers’ incentive to adapt.³¹ While this is not a perfect test—other things could have changed over time that affected farmers’ ability to

²⁹ 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 1980s and 1990s 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.

³⁰ For more details on the program, see <http://www.rma.usda.gov/>. We note that direct income support from the government constitutes a rather small percentage of cash income during our main study period—an average of 7 percent in the Corn Belt during the 1980–2000 period—suggesting that the distortionary effects of these programs on the adaptation decision were likely small. Additional data on farm income over time are available here: <http://www.ers.usda.gov/data-products/farm-income-and-wealth-statistics.aspx>.

³¹ Running the long differences model for 1997–2003 (thus, with five-year average endpoints, utilizing data from 1995–2005) gives a $\beta_{GDD>29} = -0.00438$ (SE = 0.00179), which is almost exactly equal to our baseline estimate for the 1980–2000 period, and less negative than the coefficient for the long differences run over 1980–1993 ($\beta_{GDD>29} = -0.0056$).

adapt—it provides suggestive evidence that our results are not being wholly driven by government programs.

E. Lack of Recognition of Climate Change

Finally, it could be the case that farmers didn't adapt because they didn't realize the climate was changing and that adaptation was needed. Although this doesn't affect the internal validity of our results, it could mean that our results might provide a poor guide to impacts under future climate change if the need for adaptation becomes apparent. Unfortunately, we do not directly observe farmer perceptions of temperature increases, nor their knowledge of the relationship between temperature and crop yields.³² To make progress, and building directly on the model presented in Section I, we first explore whether farmers' responsiveness is a function of characteristics that likely shape their ability to learn about a changing climate. In particular, if adaptation is limited by a difficulty in learning about climate change, then we should observe more adaptation when farmers are given more time to learn about a given change in climate, and more adaptation if they are in an area with a lower temperature variance and thus a clearer "signal" of a given change in climate.

Our data are inconsistent with either of these predictions. First, as shown in Figure 4, point estimates for longer long-difference periods (e.g., the 25-year and 30-year estimates in the bottom right panels) are almost uniformly more negative than estimates for the 1980–2000 period, although we cannot reject that they are the same in most cases. Second, we find little evidence that a lower temperature variance at baseline increased adaptation to a subsequent temperature increase. In the first column of Table 5, we re-estimate our main equation, interacting the 1980–2000 extreme temperature change in a given county with the baseline (1950–1980) standard deviation in extreme heat exposure in that county. The estimate on the interaction term is small and statistically insignificant, providing little evidence that a lower underlying variance helped farmers separate signal from noise. As a third check, and following on recent survey evidence suggesting that past experience informs current beliefs about climate change,³³ we explore whether counties that were rapidly warming prior to our study period were more adaptive during our study period. In particular, we allow the effect of extreme heat over the 1980–2000 period in a given county to depend on the change in extreme heat in that county during the period from 1960–1980, or during 1970–1980 (if farmers weight recent evidence more heavily). As shown in columns 2 and 3 of Table 5, coefficients on either interaction

³² A few existing surveys do ask farmers about their perceptions of different aspects of climate change, but the results are difficult to interpret. For instance, although Iowa is one of the states where temperature has changed the least in recent years, 68 percent of Iowa farmers in a recent survey indicated that they believe that "climate change is occurring" (Iowa State Extension Service 2011), but only 35 percent of them were concerned about the impacts of climate change on their farm operation. Similarly, only 18 percent of North Carolina farmers believed that climate change will decrease average yields by at least 5 percent over the next 25 years (Rejesus 2012), but slightly less than a 5 percent decline by 2030 could be consistent with projected impacts under more conservative warming scenarios, meaning these responses do not necessarily suggest a distorted perception of climate change.

³³ For instance, Myers et al. (2012) and Howe et al. (2012) show that persons residing in areas that have warmed in recent history are more likely to believe in future climate changes.

TABLE 5—HETEROGENOUS EFFECTS OF CLIMATE VARIATION ON CORN YIELDS

	(1)	(2)	(3)	(4)	(5)	(6)
GDD above	-0.427 (0.115)	-0.568 (0.099)	-0.427 (0.104)	-0.424 (0.095)	-0.395 (0.082)	-0.413 (0.085)
GDD above × SD GDD above, 1950–1980	-0.002 (0.005)					
GDD above × GDD above, 1960–1980		0.001 (0.003)				
GDD above × GDD above, 1970–1980			-0.000 (0.004)			
GDD above × State GDD above, 1960–1980				-0.005 (0.006)		
GDD above × Republican vote share, 2000					0.197 (0.424)	
GDD above × High school graduation rate, 1980						-0.002 (0.005)
Observations	1,531	1,531	1,531	1,531	1,530	1,531
Mean of dep. variable	0.24	0.24	0.24	0.24	0.24	0.24
R ²	0.610	0.615	0.610	0.610	0.623	0.614
Fixed effects	State	State	State	State	State	State

Notes: All coefficient estimates and standard errors are multiplied by 100. Dependent variable is the difference (1980–2000) in the log of smoothed corn yields. The sample of counties is limited to the 1980–2000 corn sample from Table 1, and regressions are weighted by 1980 corn acres. All variables in the table other than *GDD above threshold* are demeaned. Only coefficients on GDD above threshold and relevant interactions are reported, but all level effects, GDD below threshold, precipitation above threshold, and precipitation below threshold are also included in the regression. Temperature threshold is 29°C and precipitation threshold is 42 cm in all specifications. Standard errors are clustered at the state level.

are small and insignificant, and only the coefficient on the 1960–1980 interaction has the expected sign.

As a fourth check on the role of beliefs in shaping adaptation, we exploit the fact that beliefs about climate change display well known heterogeneity by political party affiliation, with Republicans consistently less likely than Democrats to believe that climate change is occurring (e.g., Dunlap and McCright 2008). We re-estimate our main equation and include an interaction between our climate variables and George W. Bush's county-level vote share in the 2000 presidential election. Because public debate and awareness about climate change begin in earnest in the late 1980s and early 1990s, this is a reasonable—if highly imperfect—proxy for beliefs about climate change. Results are given in column 5 of Table 5, and again suggest that expectations about climate change, as proxied by political beliefs, had a minimal effect on the responsiveness of farmers to extreme heat exposure: more Republican counties were, if anything, less sensitive to extreme heat exposure over the study period.

A final possibility is that adaptation is limited not by farmers' difficulty in learning about changing climate, but instead by difficulty in learning about the production function with respect to climate—in particular, learning that extreme heat can be damaging to productivity. Although this is a different type of learning, it suggests similar empirical tests as before: farmers should have been more likely to learn about the production function had they been given more time to do so, or had they been exposed to extreme heat in a previous period. As just discussed, we find little evidence that either of these predictions is true. This remains the case when we

expand the latter prediction to include the possibility that counties could learn from other nearby counties' experiences, interacting county-level changes in extreme heat over 1980–2000 with state-level changes in extreme heat over the previous period (column 4 of Table 5). As an alternate check, and building on existing evidence that higher educational achievement accelerates learning about agricultural technologies (Feder, Just, and Zilberman 1985), we allow the effect of extreme heat to vary by county-level educational attainment using data on county-level high school graduation rates from the 1980 US census. As shown in column 6 of Table 5, we find little evidence that county-level educational attainment affected subsequent adaptation.

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

To see whether insurance take-up responded to our observed climate trends, we re-estimate our long-differences model using insurance adoption at the end of our study period (i.e., a five-year average over 1998–2002) as the dependent variable. We explore four measures of take-up: the percent of corn acreage in a county enrolled in any of the multiple government crop insurance programs, the log of acres enrolled in a county, the number of policies sold in each county, and the total premiums paid (including subsidies) in each county. Results from this exercise are shown in online Appendix Table A.11. While the coefficients on the temperature variables are only sometimes significant with state-level clustering, results suggest that participation in the government insurance program by 2000 was higher in counties who saw large increases in exposure to harmful temperatures ($GDD > 29C$) over the previous two decades, and lower in counties that saw increase in exposure to generally helpful temperatures ($GDD0-29C$) over the same period. Moving from the tenth to the ninetieth percentile of the distribution of $GDD > 29C$ changes implies roughly a 5 percentage point increase in the acreage insured, a 23 percent increase in the number of policies sold, and a 20 percent increase in the total premiums paid. Again, however, only one of these estimates is significant at conventional levels with state-level clustering, so we do not wish to oversell these results.

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

Taken as a whole, then, we find little evidence that farmers who were more likely to learn about the effects of extreme heat on yields, or farmers who were more likely to update their expectations about future exposure to extreme heat, were more able

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

IV. 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 United States. 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. 2015). Details of the emissions scenario, the climate models, and their application are provided in the online Appendix.

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

In Figure 6 we present projections of average annual changes in corn yield by 2050 across the 18 climate models. In the top panel we use long differences estimates to generate predictions from precipitation changes, temperature changes, and combined effects of changing both temperature and precipitation. The most substantial negative effects of climate change are driven by increases in temperature, and while the magnitude of the negative effects of temperature vary across climate models, all predict fairly substantial negative effects of future warming on corn productivity. For instance, under climate change projections from the commonly used Hadley CM3 climate model, our long differences estimates deliver a predicted decrease in yields of approximately 27.3 percent 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).




FIGURE 6. PROJECTED IMPACTS OF CLIMATE CHANGE ON CORN YIELDS BY 2050

Notes: Panel A: 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 percent CI, and colors represent projections using only precipitation changes (gray), temperature changes (black), or both combined (white circles). Projections are separate for each climate model, as labeled. Panel B: projected impacts of combined temperature and precipitation changes across the same climate models, based on long differences (white) or panel estimates (black) of historical sensitivities to climate. The median projection is shown as a dashed line.

Panel B of the figure compares projections from long difference and panel models for each of the 18 different climate models. The similarity of regression estimates in the historical data results in projections that are comparable for both long differences and fixed effects, although the long differences estimates are somewhat noisier. We note that this noise is almost entirely due to the coefficient and standard error on GDD below 29°C, which is much less precisely estimated in the long differences than in the panel. Since a given temperature rise increases exposure to both harmful and beneficial GDD for almost all counties in our sample, the noise in the GDD below 29°C 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 percent by mid-century, with some models projecting yield losses as low as 7 percent and other losses as high as 64 percent. To put these projected losses in perspective, the 2012 drought and heat

wave that was considered one of the worst on record and that received extensive attention in the press decreased average corn yields for the year by 15–25 percent relative to the prior few years.³⁴ Our median projection suggests that by 2050, every year will experience losses roughly this large. Valued at production quantities and prices averaged over 2006–2010 for our sample counties, 15 percent yield losses would generate annual dollar losses of \$6.7 billion by 2050.

V. Conclusions

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

We exploit large variation in multi-decade changes in temperature and precipitation across US counties to estimate how farmers have responded to longer run changes in climate. We argue that these changes are plausibly exogenous and show that their magnitude is on par with future changes in climate projected by global climate models, making them an ideal source of variation to identify historical responses to longer run changes in climate and in 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 (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. We also provide evidence that this lack of adaptation was not driven by a lack of recognition that climate was changing, perhaps suggesting that farmers either lacked adaptation options or found them too expensive to exploit.

Using climate change projections from 18 global climate models, we project potential impacts on corn productivity by mid-century. If future adaptations are as effective as past adaptations in mitigating the effects of exposure to extreme heat, our median estimate is that future climate change will reduce annual corn productivity in 2050 by roughly 15 percent, which is on par with the effect of the

³⁴ For instance, see <http://www.ers.usda.gov/topics/in-the-news/us-drought-2012-farm-and-food-impacts.aspx>. The range in estimated losses depend on whether the comparison is against previous season's yield or the yield projected at planting in 2012.

highly-publicized “extreme” drought and heat wave experienced across the US Corn Belt in the summer of 2012. Given that these projections account for farmers’ present adaptive abilities, our results imply substantial losses under future climate change in the absence of unprecedented adaptation.

REFERENCES

- Auffhammer, Maximilian, and Anin Aroonruengsawat.** 2011. “Simulating the impacts of climate change, prices and population on California’s residential electricity consumption.” *Climatic Change* 109 (1S): 191–210.
- Auffhammer, Maximilian, V. Ramanathan, and Jeffrey R. Vincent.** 2006. “Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in India.” *Proceedings of the National Academy of Sciences* 103 (52): 19668–72.
- Bound, John, Charles Brown, and Nancy Mathiowetz.** 2001. “Measurement error in survey data.” In *Handbook of Econometrics*, Vol. 5, edited by James J. Heckman and Edward Leamer, 3705–3843. Amsterdam: North-Holland.
- Burke, Marshall, John Dykema, David B. Lobell, Edward Miguel, and Shanker Satyanath.** 2015. “Incorporating Climate Uncertainty into Estimates of Climate Change Impacts.” *Review of Economics and Statistics* 97 (2): 461–71.
- Burke, Marshall, and Kyle Emerick.** 2016. “Adaptation to Climate Change: Evidence from US Agriculture: Dataset.” *American Economic Journal: Economic Policy*. <http://dx.doi.org/10.1257/pol.20130025>.
- Burke, Marshall, and David Lobell.** 2010. “Food Security and Adaptation to Climate Change: What Do We Know?” In *Climate Change and Food Security*, edited by David Lobell and Marshall Burke, 133–53. New York: Springer.
- Burke, Marshall B., Edward Miguel, Shanker Satyanath, John A. Dykema, and David B. Lobell.** 2009. “Warming increases the risk of civil war in Africa.” *Proceedings of the National Academy of Sciences* 106 (49): 20670–74.
- Cutler, David M., Grant Miller, and Douglas M. Norton.** 2007. “Evidence on early-life income and late-life health from America’s Dust Bowl era.” *Proceedings of the National Academy of Sciences* 104 (33): 13244–49.
- Davis, Donald R., and David E. Weinstein.** 2002. “Bones, Bombs, and Break Points: The Geography of Economic Activity.” *American Economic Review* 92 (5): 1269–89.
- DeGroot, Morris H.** 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken.** 2012. “Temperature Shocks and Economic Growth: Evidence from the Last Half Century.” *American Economic Journal: Macroeconomics* 4 (3): 66–95.
- Deschênes, Olivier, and Michael Greenstone.** 2007. “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather.” *American Economic Review* 97 (1): 354–85.
- Deschênes, Olivier, and Michael Greenstone.** 2011. “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US.” *American Economic Journal: Applied Economics* 3 (4): 152–85.
- Dunlap, Riley E., and Aaron M. McCright.** 2008. “A widening gap: Republican and Democratic views on climate change.” *Environment* 50 (5): 26–35.
- Feder, Gershon, Richard E. Just, and David Zilberman.** 1985. “Adoption of agricultural innovations in developing countries: A survey.” *Economic Development and Cultural Change* 33 (2): 255–98.
- Feng, Shuaizhang, Michael Oppenheimer, and Wolfram Schlenker.** 2015. “Weather Anomalies, Crop Yields, and Migration in the US Corn Belt.” <http://www.columbia.edu/~ws2162/articles/FengOppenheimerSchlenker.pdf>.
- Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker.** 2012. “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment.” *American Economic Review* 102 (7): 3749–60.
- Greenwald, Roby, M. H. Bergin, Jin Xu, Daniel Cohan, Gerrit Hoogenboom, and W. L. Chameides.** 2006. “The influence of aerosols on crop production: A study using the CERES crop model.” *Agricultural Systems* 89 (2): 390–413.
- Griliches, Zvi, and Jerry A. Hausman.** 1986. “Errors in variables in panel data.” *Journal of Econometrics* 31 (1): 93–118.

- Gwynne, Peter. 1975. "The cooling world." *Newsweek*, April 28. http://denisdutton.com/newsweek_coolingworld.pdf.
- Hornbeck, Richard. 2012. "The Enduring Impact of the American Dust Bowl: Short and Long-Run Adjustments to Environmental Catastrophe." *American Economic Review* 102 (4): 1477–1507.
- Howe, Peter D., Ezra M. Markowitz, Tien Ming Lee, Chia-Ying Ko, and Anthony Leiserowitz. 2012. "Global perceptions of local temperature change." *Nature Climate Change* 3: 352–56.
- Iowa State Extension Service. 2011. *Iowa Farm and Rural Live Poll: 2011 Summary Report*. Iowa State University Extension and Outreach. Ames, IA, December.
- Kelly, David L., Charles D. Kolstad, and Glenn T. Mitchell. 2005. "Adjustment costs from environmental change." *Journal of Environmental Economics and Management* 50 (3): 468–95.
- Leibensperger, E. M., L. J. Mickley, D. J. Jacob, W. T. Chen, J. H. Seinfeld, A. Nenes, P. J. Adams, et al. 2011. "Climatic effects of 1950–2050 changes in US anthropogenic aerosols, Part 2: Climate response." *Atmospheric Chemistry and Physics Discussions* 11: 24127–64.
- Lobell, David B. 2013. "Errors in Climate Datasets and Their Effects on Statistical Crop Models." *Agricultural and Forest Meteorology* 170: 58–66.
- Lobell, David B., and Gregory P. Asner. 2003. "Climate and Management Contributions to Recent Trends in U.S. Agricultural Yields." *Science* 299 (5609): 1032.
- Lobell, D. B., G. Bala, and P. B. Duffy. 2006. "Biogeophysical impacts of cropland management changes on climate." *Geophysical Research Letters* 33 (6): L06708–12.
- Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate Trends and Global Crop Production Since 1980." *Science* 333 (6042): 616–20.
- McIntosh, Craig T., and Wolfram Schlenker. 2006. "Identifying Non-Linearities in Fixed Effects Models." https://gps.ucsd.edu/_files/faculty/mcintosh/mcintosh_research_identifying.pdf.
- Meehl, Gerald A., Julie M. Arblaster, and Grant Branstator. 2012. "Mechanisms Contributing to the Warming Hole and the Consequent U.S. East–West Differential of Heat Extremes." *Journal of Climate* 25 (18): 6394–6408.
- Meehl, Gerald A., Thomas F. Stocker, William D. Collins, Pierre Friedlingstein, Amadou T. Gaye, Jonathon M. Gregory, Akio Kitoh, et al. 2007. "Global Climate Projections." In *IPCC Fourth Assessment Report, Working Group I*, edited by Susan Solomon, Dahe Qin, Martin Manning, Melinda Marquis, Kirsten Averyt, Melinda M. B. Tignor, Henry LeRoy Miller, Jr., et al., 747–845. Cambridge: Cambridge University Press.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw. 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *American Economic Review* 89 (4): 753–71.
- Miguel, Edward, and Gérard Roland. 2011. "The long-run impact of bombing Vietnam." *Journal of Development Economics* 96 (1): 1–15.
- Myers, Teresa A., Edward W. Maibach, Connie Roser-Renouf, Karen Akerlof, and Anthony A. Leiserowitz. 2012. "The relationship between personal experience and belief in the reality of global warming." *Nature Climate Change* 3: 343–47.
- Olmstead, Alan L., and Paul W. Rhode. 2011. "Adapting North American wheat production to climatic challenges, 1839–2009." *Proceedings of the National Academy of Sciences* 108 (2): 480–85.
- Rejesus, Roderick M. 2012. *Farmer Perceptions and Beliefs about Climate Change: A North Carolina Perspective*. North Carolina Cooperative Extension Service. Raleigh, March/April.
- Samuelson, Paul A. 1947. *Foundations of Economic Analysis*. Cambridge: Harvard University Press.
- Schlenker, Wolfram. 2006. "Interannual Weather Variation and Crop Yields." Unpublished.
- Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. 2005. "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach." *American Economic Review* 95 (1): 395–406.
- Schlenker, Wolfram, and David B. Lobell. 2010. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5 (1).
- Schlenker, Wolfram, and Michael J. Roberts. 2009. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of Sciences* 106 (37): 15594–98.
- Viner, Jacob. 1958. *The Long View and the Short: Studies in Economic Theory and Policy*. New York: The Free Press.
- Welch, Jarrod R., Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. 2010. "Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures." *Proceedings of the National Academy of Sciences* 107 (33): 14562–67.
- Zilberman, David, Jinhua Zhao, and Amir Heiman. 2012. "Adoption Versus Adaptation, with Emphasis on Climate Change." *Annual Review of Resource Economics* 4: 27–53.