

The Impact of Climate Change on U.S. Agriculture: New Evidence on the Role of Heterogeneity and Adaptation

By MICHAEL KEANE AND TIMOTHY NEAL*

This paper investigates the impact of climate change on the productivity of crop production using U.S. county-level yield and weather data between 1950 and 2015. It finds that the pooled estimators used in previous studies underestimate the sensitivity of crops to high temperatures by ignoring slope heterogeneity, and underestimate the damage of future climate change on yield. Furthermore, explicitly modelling this heterogeneity provides a natural approach to measuring the degree of adaptation to climate change in the data. It concludes with evidence that further adaptation may mitigate up to half of the substantial losses to crop productivity forecast by 2050.

JEL: C23, Q15, Q51, Q54

Many scientific and environmental agencies are warning of the extreme impacts that future climate change may have on the world, including the productivity of agriculture and global food supply. In contrast, studies that have attempted to estimate the sensitivity of agriculture to the climate in hopes of obtaining insight into the effects of future climate change have obtained very mixed results. The forecast impact of climate change on U.S. agriculture in particular ranges from extreme damage to crop produc-

* Keane: Nuffield College, Oxford University, Manor Road Building, Oxford OX1 3UQ (email: michael.keane@economics.ox.ac.uk); Neal: Department of Economics, University of New South Wales, Anzac Parade, NSW 2000 (email: timothy.neal@unsw.edu.au).

tivity (e.g. Schlenker and Roberts 2009) and crop quality (Kawasaki and Shinsuke 2016) to very minor damage or even net benefits to agriculture (e.g. Deschenes and Greenstone 2007 and Mendelsohn, Nordhaus and Shaw 1994). Resolving this disagreement and uncertainty concerning the relationship between crop yield and the climate has been viewed as one of the top priorities for the improvement of future climate change impact assessments (Lobell and Burke 2008).

Critical to the discussion is the degree of adaptation to climate change that has occurred in the past and can be expected from agricultural producers in the future. Adaptation to climate change may involve switching crop cultivars to ones that are more tolerant of heat, increasing water retention in fields, irrigation, fertilizers, altering the planting and harvesting dates, shifting the spatial distribution of agricultural production to cooler areas, or switching to more heat-resistant crops entirely. The very existence of adaptation between regions and over time implies that there is slope heterogeneity in the relationship between crop yield and high temperatures. This heterogeneity could be modelled as occurring between counties or over time, or indeed in both dimensions.

In this article we investigate these issues using temperature and crop yield data for U.S. counties between 1950 and 2015 (thereby including the extreme drought experienced by the Midwest in 2012-13). We focus on estimating the climate sensitivity of corn (maize) and soybeans. These crops are by far the two largest in the United States in terms of tonnage, and accordingly are important for national food security and global food supply. Understanding their sensitivity to the climate is important as temperatures in excess of a certain threshold can significantly decrease crop yield by directly damaging

the plant tissue or enzymes through heat stress. Several recent studies investigate this topic using panel data regressions of crop yields on temperature (see, e.g., Schlenker and Roberts 2009, Butler and Huybers 2013, and Burke and Emerick 2016).

Our first key finding relevant to this literature is that ignoring the slope heterogeneity in the relationship between yield and the climate leads to significant underestimation of crop sensitivity to temperature, and hence to forecasts of future damage from climate change. This is because the heterogeneity, representing adaptation or lack thereof, is positively correlated with the amount of time a crop is exposed to harmful temperatures over a growing season, which is the regressor of interest in crop yield regressions. This correlation arises because counties or time periods that experience consistently high levels of heat (meaning high values of the regressor) have a strong financial incentive to adopt adaptation techniques resulting in the crop being less sensitive to heat (meaning a slope coefficient less negative and closer to zero). After comparing the results from a series of regressions using a pooled panel estimator versus an estimator that accounts for slope heterogeneity between counties and over time, the empirical results suggest that pooled panel regressions underestimate the sensitivity of crop yields by around 35 percent.

Analysing the distribution of heterogeneity in the temperature coefficient also provides a natural and direct approach to measuring the degree of historical adaptation to climate in the data. Results from an estimator that accounts for heterogeneity provides strong evidence for significant amounts of adaptation occurring between counties, yet no evidence that average levels of adaptation have increased over time since the 1980s. They also provide

useful insights into the extent to which adaptation can be expected to mitigate future damage from climate change.

This approach to measuring adaptation is one of the main contributions of our work. There is no established method to measure adaptation in the literature, and previous attempts have used a variety of techniques. Butler and Huybers (2013), Schlenker and Roberts (2009), and Burke and Emerick (2016) all use different approaches to measuring adaptation (see Section I.A for details). The first paper argues that adaptation could cut future yield losses from climate change in half, while the latter two conclude that adaptation has been rather limited. The diversity of results highlights the need for further work in this area.

Lastly, we combine our model estimates with forecasts from a range of climate models that project temperature and precipitation forward to 2050 under the A1B emissions scenario. Using the simple fixed effects estimator, the mean forecast (across climate models) for crop yield damage in this emissions scenario is 29 percent for corn and 35 percent for soybeans. The mean forecast damage increases to 44 percent for corn and 55 percent for soybeans when using an estimator that accounts for slope parameter heterogeneity between counties and over time. These increases of 50 to 60 percent are significant and reflect the bias of the econometric estimates that ignore heterogeneity.

We also consider more optimistic scenarios that contain further adaptation and are obtained by allowing poorly adapted counties to become one half or one standard deviation more adapted by the forecast year of 2050. This exercise suggests that the predicted damage to corn yield could be reduced to somewhere in the range of 21-34 percent, depending on the degree of

further adaptation that is assumed to be reasonable over the next three decades. The results are much less optimistic for soybean yield, however, as there is significantly less heterogeneity in the estimates of climate sensitivity for soybeans. Of course, it is important to note that forecasting decades into the future using these techniques involves a number of strong assumptions relating to future technological improvement and also excludes more drastic adaptation techniques (such as switching crops or land use).

The rest of this paper is structured as follows. The first section outlines the previous approaches to this topic found in the literature, why pooled estimates of climate sensitivity are biased, the econometric methodology that we adopt, and finally the sources for the data used in this article. The second section details the results of regressions for both pooled estimators as well as estimators that incorporate two-dimensional slope heterogeneity. The third section conducts a forecasting exercise to 2050 using the results of the second section, while the fourth section concludes.

I. Empirical Approach

A. Previous Approaches

Several recent papers attempt to estimate the impact of climate change on agriculture using what we will call the ‘GDD’ approach. It relates climate to the production or yield (yield meaning production per acre of land) of specific crops. A GDD, or growing degree day, refers to the amount of time that a crop in its growing season is exposed to temperatures within a specific band of temperature. Temperature can be either beneficial or harmful to a crop depending on its value, motivating a separation of temperature into bands based on a threshold value.

Examples of the GDD approach include Schlenker and Roberts (2009), Burke and Emerick (2016), Kawasaki and Shinsuke (2016), and Lobell et al. (2011). These studies typically estimate an equation similar to the following:

$$(1) \quad y_{it} = \alpha_i + \alpha_t + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

where y_{it} is the log of a particular crop yield for county i at year t , GDD_{it} is the number of days that the crop was exposed to beneficial temperatures, KDD_{it} or ‘killing degree days’ refers to the number of days that the crop was exposed to harmful temperatures, and $PREC_{it}$ is the sum measure of precipitation that the crop experienced during the growing season. α_i absorbs all unobserved time-invariant intercept heterogeneity that is related to county i (such as soil quality), while α_t absorbs all unobserved intercept heterogeneity that is constant across U.S. counties but varies between years t (such as farming technology). The threshold temperature used to distinguish harmful from beneficial temperature is typically 29 degrees Celsius for both corn and soybeans. The key parameter of interest is β_2 which determines the extent to which high temperatures reduce crop yield.¹

¹An alternative approach to estimating the impact of climate change on agricultural productivity is the hedonic (or ‘Ricardian’) approach (e.g. Mendelsohn, Nordhaus and Shaw 1994 and Deschenes and Greenstone 2007). It relies on estimating the relationship between climatic variables and agricultural land values. By linking climate to land values, as opposed to crop yield, these studies are able to examine the entire agricultural sector. Furthermore, they are able to account for a broad range of adaptations and shifts in farmer behaviour, such as switching crops or changes in land use.

The limitation of the hedonic approach is that its validity relies on some strong assumptions. For instance, it assumes output and input prices remain constant, which is problematic if there are climate-induced price changes (see Darwin 1999*a*). Cross-sectional estimates also suffer from omitted variable bias, due to variables that are related to both climate and land values, such as irrigation and soil quality. Indeed, pooling irrigated and non-irrigated farms has been shown to lead to underestimation of damage from future climate change (Schlenker, Hanemann and Fisher 1994 and Darwin 1999*b*).

Several studies also use the GDD approach to determine the extent to which adaptation to higher temperatures has occurred. Adaptation can include a wide range of activities, including switching corn cultivars to one that has a higher heat tolerance (usually by producing more heat-resistant proteins, as well as several other properties that provide drought resistance), improving water retention in the field, the use of fertilisers, switching crops entirely, and even relocating agricultural production to cooler areas. The main motivation for understanding the extent of historical adaptation is that it may substantially alter the forecast damage to crop productivity from climate change.

Schlenker and Roberts (2009) tested for evidence of historical adaptation to extreme heat by running regressions after splitting the sample into northern and southern U.S. states and also by 1950-1977 and 1978-2005 periods. They argued that since the coefficients did not change significantly after splitting the sample either by region or time, there is evidence that historical adaptation has been very limited. In contrast, Butler and Huybers (2013) ran regressions on each county separately using data from 1981 to 2008 and concluded from the collection of coefficients that there was adaptation occurring between counties. When forecasting the future damage of climate change on crop yields they argued that yield losses could be halved with further adaptation.

Finally, Burke and Emerick (2016) adopt what they call a ‘long differences’ approach to measuring the extent of historical adaptation. They estimate (1) using a U.S. county panel dataset from 1980 to 2000, and then calculate the 1978-1982 and 1998-2002 averages of climate and yield and estimate the

following long difference regression:

$$(2) \quad \begin{aligned} \Delta y_{is} = & \alpha_s + \beta_1 \Delta GDD_{is} + \beta_2 \Delta KDD_{is} + \beta_3 \Delta PREC_{is} \\ & + \beta_4 \Delta PREC_{is}^2 + \Delta \epsilon_{is} \end{aligned}$$

where ΔGDD_{is} is the difference between the 1998-2002 average and the 1978-1982 average for GDD in county i and state s , and similarly for the other variables. The extent of adaptation is then calculated as $1 - \frac{\hat{\beta}_2^{LD}}{\hat{\beta}_2^{FE}}$, where $\hat{\beta}_2^{LD}$ is the point estimate from the long difference model and $\hat{\beta}_2^{FE}$ is the point estimate from the panel fixed effects model. They conclude from this exercise that adaptation has been fairly minor in magnitude.

A key point is that both the GDD and long difference methods rely on the panel fixed effects model in (1) providing a consistent and unbiased estimate of the sensitivity of crop yield to yearly variations in temperature. The next subsection will demonstrate why this is unlikely to be true.

B. Bias in Previous Approaches

The very existence of adaptation techniques implies the presence of heterogeneity in the sensitivity parameter to high temperatures, β_2 . This heterogeneity may occur between counties, because some counties invest in adaptation techniques more than others, or over time, as better cultivars become available or the availability of irrigated water increases or decreases. This multidimensional heterogeneity can be represented in the following equation:

$$(3) \quad y_{it} = \alpha_i + \alpha_t + \beta_{1it} GDD_{it} + \beta_{2it} KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

where $\beta_{1it} = \lambda_1 + \lambda_{1i} + \lambda_{1t}$ and $\beta_{2it} = \lambda_2 + \lambda_{2i} + \lambda_{2t}$. This specification

allows for additive random or fixed effects in the sensitivity of crop yield to temperature.

To see the bias of pooled estimators in a model containing adaptation through slope heterogeneity, first simplify (3) by excluding precipitation and stacking the variables:

$$(4) \quad y_{it} = \alpha_i + \alpha_t + \mathbf{z}'_{it}\boldsymbol{\theta}_{it} + \epsilon_{it}$$

where $\mathbf{z}'_{it} = (GDD_{it}, KDD_{it})$ and $\boldsymbol{\theta}_{it} = (\beta_{1it}, \beta_{2it})' = (\lambda_1 + \lambda_{1i} + \lambda_{1t}, \lambda_2 + \lambda_{2i} + \lambda_{2t})'$. Consider a two-way fixed effects pooled OLS regression of (4) to remove the fixed effects and obtain an estimate of the average slope coefficients:

$$(5) \quad \tilde{y}_{it} = \tilde{\mathbf{z}}'_{it}\boldsymbol{\theta} + v_{it}$$

where $\tilde{y}_{it} = y_{it} - N^{-1} \sum_{i=1}^N y_{it} - T^{-1} \sum_{t=1}^T y_{it} + NT^{-1} \sum_{i=1}^N \sum_{t=1}^T y_{it}$ and similarly for $\tilde{\mathbf{z}}_{it}$, $\boldsymbol{\theta} = (\lambda_1, \lambda_2)'$, and v_{it} is defined as:

$$(6) \quad v_{it} = \tilde{\mathbf{z}}'_{it}\boldsymbol{\theta}_i + \tilde{\mathbf{z}}'_{it}\boldsymbol{\theta}_t + \epsilon_{it}$$

where $\boldsymbol{\theta}_i = (\lambda_{1i}, \lambda_{2i})'$ and $\boldsymbol{\theta}_t = (\lambda_{1t}, \lambda_{2t})'$.

The OLS estimate of $\boldsymbol{\theta}$ will be:

$$(7) \quad \hat{\boldsymbol{\theta}} = \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it} \tilde{\mathbf{z}}'_{it} \right)^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it} \tilde{y}_{it} \right)$$

Expanding on \tilde{y}_{it} and simplifying yields:

$$(8) \quad \hat{\theta} = \theta + \mathbf{Q}_{zz,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{z}_{it} \tilde{z}'_{it} \theta_i \right) \\ + \mathbf{Q}_{zz,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{z}_{it} \tilde{z}'_{it} \theta_t \right) + \mathbf{Q}_{zz,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{z}_{it} \epsilon_{it} \right)$$

where $\mathbf{Q}_{zz,NT}^{-1} = \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{z}_{it} \tilde{z}'_{it} \right)^{-1}$. If \tilde{z}_{it} , θ_i , and θ_t are independent this will simplify to:

$$(9) \quad \hat{\theta} = \theta + \frac{1}{N} \sum_{i=1}^N \theta_i + \frac{1}{T} \sum_{t=1}^T \theta_t + \mathbf{Q}_{zz,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{z}_{it} \epsilon_{it} \right)$$

and the pooled estimate of the average slope will be unbiased and asymptotically consistent as $(N, T) \xrightarrow{j} \infty$ (provided the usual assumptions are valid). However, if \tilde{z}_{it} , θ_i , and θ_t are not independent then this will not hold and there will be bias in the pooled estimates.

Moreover, there is every reason to believe that λ_{2i} and λ_{2t} , the heterogeneity components in the slope coefficient for harmful temperatures, are positively related to the amount of harmful temperatures KDD_{it} . In counties and years that feature hotter temperatures than average there is a financial incentive to use adaptation techniques that alleviate the damage of heat to the crop yield. Accordingly there will be a positive relation between KDD_{it} and the heterogeneity components of its slope coefficient.

Indeed, Butler and Huybers (2013) provided evidence that the slope coefficient for a county is significantly and positively correlated with its relative experience of heat. Since the correlation is positive, pooled regressions will likely overestimate this parameter (i.e. underestimate the sensitivity of crop

yield to high temperatures), and therefore underestimate the forecast damage due to climate change.

C. Econometric Methodology

We will begin by estimating a suite of pooled models in order to verify that they are biased and as a point of comparison in the forecasting exercise. In this preliminary exercise, the following equation will be estimated using a variety of assumptions on the fixed effects term α_{it} :

$$(10) \quad y_{it} = \alpha_{it} + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

Following the literature, states west of the 100th Meridian that are more reliant on widespread irrigation are excluded from the sample (although we found it did not meaningfully affect parameter estimates in this dataset).² Regression weights are based on the average area in the county used to produce the specific crop, and the standard errors are also clustered at the state level. The temperature threshold used to split GDD and KDD is 29 degrees Celsius for both corn and soybeans, again following the literature. Butler and Huybers (2013) noted that although 29 degrees may appear low as a threshold for damaging temperatures, the temperature experienced by the corn plant itself is higher than the air temperature above the crop canopy.

We consider models where the intercept α_{it} includes year fixed effects, county and year fixed effects, and county fixed effects with state quadratic time trends. If the model includes both county and year fixed effects, pa-

²Note: The 100th Meridian separates the Great Plains to the east from the semi-arid lands to the west.

rameters are identified from deviations in county weather from the county average after removing any annual weather shocks that are common to all counties.

The next step will be to introduce slope heterogeneity into the model. This allows us to determine whether the average coefficient changes once the heterogeneity bias is removed. It also allows us to determine if there is evidence of adaptation by looking at the distribution of the coefficients. First consider a model containing heterogeneity only between counties:

$$(11) \quad y_{it} = \alpha_i + \beta_{1i}GDD_{it} + \beta_{2i}KDD_{it} + \beta_{3i}PREC_{it} + \beta_{4i}PREC_{it}^2 + \epsilon_{it}$$

This can be estimated using mean-group OLS (or ‘MG-OLS’) which was proposed in Pesaran and Smith (1995). It involves simply running regressions for each county across the entire sample period of 1950-2015, collecting all of the beta coefficients, and then averaging them. The strong limitation of this approach in our context is its neglect of time effects in the intercept or the slope coefficients, features which are of central importance for the questions we address.

Accordingly, we estimate a model that allows the intercept and slope coefficients to vary between counties and over time:

$$(12) \quad \begin{aligned} y_{it} = & \alpha_i + \alpha_t + \beta_{1it}GDD_{it} + \beta_{2it}KDD_{it} + \beta_{3it}PREC_{it} \\ & + \beta_{4it}PREC_{it}^2 + \epsilon_{it} \end{aligned}$$

This can be estimated using mean-observation OLS (or ‘MO-OLS’) which was proposed in Neal (2016). Crucially, the estimator allows the regressors to be correlated with the heterogeneity, and also allows county and year fixed

effects to be included in the model. It assumes that the multidimensional heterogeneity is additive as in (3). Very briefly, the rationale behind the estimator is to obtain a pooled estimate, say generically β , then run a series of regressions for each individual to collect β_i , then a series of regressions for each time period to collect β_t , and then construct $\beta_{it} = \beta_i + \beta_t - \beta$.³ The MO-OLS estimate is then a simple or weighted average of these bias-corrected β_{it} estimates. The distribution of the coefficients provides an account of the degree of adaptation that can be found between counties and over time simultaneously.

It is important to be clear about the types of adaptation that will be captured by these approaches and those that will not. Heterogeneity between counties will pick up all forms of adaptations related to the individual crop, whether short or long term in nature. This includes irrigation, cultivar adoption, fertilisers, improving water retention, and other techniques. Time-based heterogeneity will capture adaptation techniques that are geographically widespread but were not available throughout the sample. This will include any technological advancement or adaptation in response to a changing climate. But the methods we discuss here will not capture certain other forms of adaptation such as changing the planting and harvest date for the crop, closing down farms no longer viable due to warming temperatures, crop switching, or land use changes. These forms of adaptation may prove to be important, especially when climate change becomes more severe. Future work can attempt to investigate whether there is evidence of

³While this procedure will remove most of the heterogeneity bias, there is a residual bias which the estimator corrects by repeatedly using sample approximations of the residual bias until it becomes sufficiently small (in practice usually zero to several decimal places). For more details on the bias correction procedure and proofs of asymptotic consistency and normality, please see Neal (2016).

planting and harvest date changes, and also whether the spatial distribution of agricultural production is also changing in line with the climate.

D. Data Sources

The temperature and precipitation data used in this study were obtained from Schlenker and Roberts (2009), they contain daily observations on minimum temperature, maximum temperature, and precipitation in a grid across the continental United States from 1950 to 2015.⁴ They also provided the code to parse this grid-based weather data into counties, weighting the grid locations according to where the agricultural production is located in each county. From the daily maximum and minimum observations it is possible to approximate the amount of time each day that a crop is exposed to one-degree Celsius temperature intervals, and that is easily extended to the amount of time in the entire growing season by summing across all days. The growing seasons for both corn and soybeans are assumed to be between the 1st of May and the 30th of September, in line with the literature. Using these single-degree temperature intervals across the growing season, the variables GDD_{it} and KDD_{it} can be aggregated by summing across the relevant intervals. The former is summed across 0 to 29 degrees Celsius, while the latter is summed across the degree days of all temperatures above 29 degrees Celsius.

Data on crop yield was obtained from the United States Department of Agriculture's National Agricultural Statistics Service. Information is provided at the county-level and covers the same annual sample period of 1950

⁴The raw data comes from the PRISM Climate Group. Temperature observations prior to 1950 are available, but the further back in time the observations go the harder it is to convert them into a grid as fewer weather stations are operational.

to 2015, although not all counties contain crop yield data across all years (creating an unbalanced panel). While data is reported across a range of agricultural produce, we focus on corn and soybeans, the two largest crops. These data also provide the average crop area of each county, which we use for regression weights.

II. Results

A. Pooled Regressions

Table 1 presents the results for corn yield from a number of pooled regressions that provide a basis of comparison for our subsequent models containing slope heterogeneity. The table lists four specifications where the first features no fixed effects, the second contains year fixed effects, the third has county and year fixed effects, and finally the fourth has county fixed effects and state-specific quadratic time trends. Each model includes precipitation in levels as well as in squares, although the results do not change meaningfully when these are removed from the model.

The estimated sensitivity of corn to high temperatures when the model contains no fixed effects is -0.0073 , indicating that exposure to an additional degree-day of heat above 29°C will lead to a decrease in overall corn yield of 0.73 percent. Including county and time fixed effects, or county fixed effects with a state-specific quadratic time trend, decreases this sensitivity to be a 0.63 percent reduction in corn yield for each additional degree-day over 29°C . These results are similar to those reported in Burke and Emerick (2016). Although unreported, we also found that this sensitivity does not change significantly when states west of the 100th Meridian line are included or whether the regression results are weighted by harvest area or not.

TABLE 1—POOLED PANEL ESTIMATES OF THE IMPACTS OF TEMPERATURE ON CORN YIELDS

Specification	(1)	(2)	(3)	(4)
GDD	0.0004 (0.0001)	0.0002 (0.0001)	0.0003 (0.0001)	0.0003 (0.0001)
KDD	-0.0073 (0.0016)	-0.0053 (0.0012)	-0.0063 (0.0007)	-0.0063 (0.0005)
Precipitation	0.0009 (0.0009)	0.0013 (0.0009)	0.0010 (0.0002)	0.0013 (0.0003)
Precipitation ²	-6.9e-07 (6.4e-07)	-1.2e-06 (6.2e-07)	-9.2e-07 (2.3e-07)	-1.11e-06 (2.5e-07)
Constant	3.1497 (0.6528)	2.7941 (0.6460)	2.7188 (0.2661)	2.4767 (0.2123)
Fixed Effects	No	Yr	Cty, Yr	Cty, State-Yr
Obs.	126,043			
R^2	0.1576	0.6993	0.8245	0.8354

Notes: Results exclude states west of the 100th Meridian line. Sample range is 1950-2015, with specifications (1) - (4) differing by type of fixed effects as outlined in the table. Regressions are weighted by average county harvest area for each crop. Standard errors are reported in parentheses, and are clustered at the state level.

The results for soybean yield are presented in Table 2 and are similar in that including fixed effects for counties and time decreases the estimated sensitivity parameter relative to the baseline specification. The sensitivity for soybeans is lower than it is for corn, except in the model that only contains year fixed effects. Overall, the pooled results are broadly consistent with previous studies and are not particularly sensitive to changes in modelling assumptions.

B. Heterogeneity across Counties and Time

Estimating a model with both dimensions of heterogeneity will provide a complete overview of the extent of adaptation both between counties and

TABLE 2—POOLED PANEL ESTIMATES OF THE IMPACTS OF TEMPERATURE ON SOY-BEAN YIELDS

Specification	(1)	(2)	(3)	(4)
GDD	0.0003 (0.0001)	0.0003 (0.0001)	0.0005 (0.0001)	0.0004 (0.0000)
KDD	-0.0058 (0.0008)	-0.0057 (0.0008)	-0.0048 (0.0003)	-0.0043 (0.0003)
Precipitation	0.0016 (0.0004)	0.0017 (0.0004)	0.0012 (0.0001)	0.0012 (0.0002)
Precipitation ²	-1.2e-06 (2.8e-07)	-1.3e-06 (3.1e-07)	-9.1e-07 (1.1e-07)	-9.2e-07 (1.4e-07)
Constant	1.9542 (0.3711)	1.6434 (0.4169)	0.9127 (0.2702)	1.2853 (0.1461)
Fixed Effects	No	Yr	Cty, Yr	Cty, State-Yr
Obs.	87,767			
R^2	0.2213	0.6237	0.7523	0.7509

Notes: Results exclude states west of the 100th Meridian line. Sample range is 1950-2015, with specifications (1) - (4) differing by type of fixed effects as outlined in the table. Regressions are weighted by average county harvest area for each crop. Standard errors are reported in parentheses, and are clustered at the state level.

over the sample period. The Mean-Observation OLS estimator is used to model both dimensions of heterogeneity, and the results for corn are presented in Table 3. The table outlines the unweighted and weighted mean of the distribution of slope coefficients, as well as other statistics concerning the distribution. The mean coefficient for KDD is significantly more negative than the pooled model estimates in Table 1, reflecting the heterogeneity bias in the pooled estimates that was outlined in section I.B. The unweighted mean is -0.0096, meaning that one extra degree day of temperatures over 29° will lead to a 0.96 percent reduction in crop yield in that county. This coefficient represents a 35 percent increase in sensitivity compared to the

pooled model that also featured county and time fixed effects. Weighting the mean by the harvested area of corn in each county reduces it slightly to -0.0089.

Meanwhile, there is a significant degree of heterogeneity in the distribution of the coefficients. The standard deviation of the slope coefficients for KDD is 0.0068, with the 10th to 90th percentile range being -0.0161 to -0.0034. Using the Spearman correlation coefficient the correlation between $\hat{\beta}_{2it}$ and KDD_{it} is estimated to be 0.43, supporting the argument made in Section I.B that the heterogeneity will be correlated with this regressor (leading to bias which we also observe), and that the correlation will be positive leading to overestimation in particular.

TABLE 3—MEAN-OBSERVATION ESTIMATES OF THE IMPACTS OF TEMPERATURE ON U.S. CORN YIELDS

	Mean	Weighted Mean	Median	Standard Deviation	10th Percentile	90th Percentile
GDD	0.0005 (0.0000)	0.0005	0.0005	0.0005	-0.0001	0.0011
KDD	-0.0096 (0.0003)	-0.0089	-0.0085	0.0068	-0.0161	-0.0034
Precipitation	0.0011 (0.0002)	0.0015	0.0010	0.0028	-0.0020	0.0044
Precipitation ²	-1.1e-06 (1.4e-07)	-1.4e-06	-8.5e-07	2.5e-06	-3.8e-06	1.4e-06
Constant	2.7080 (0.1208)	2.8947	2.8225	2.1363	0.1713	4.8419
Obs.	126,043					

Notes: Results exclude states west of the 100th Meridian line. The sample range is 1950-2015, and standard errors are reported in parentheses.

Figure 1 plots the distribution of the KDD slope coefficients between counties and over time. We observe that there is significant heterogeneity be-

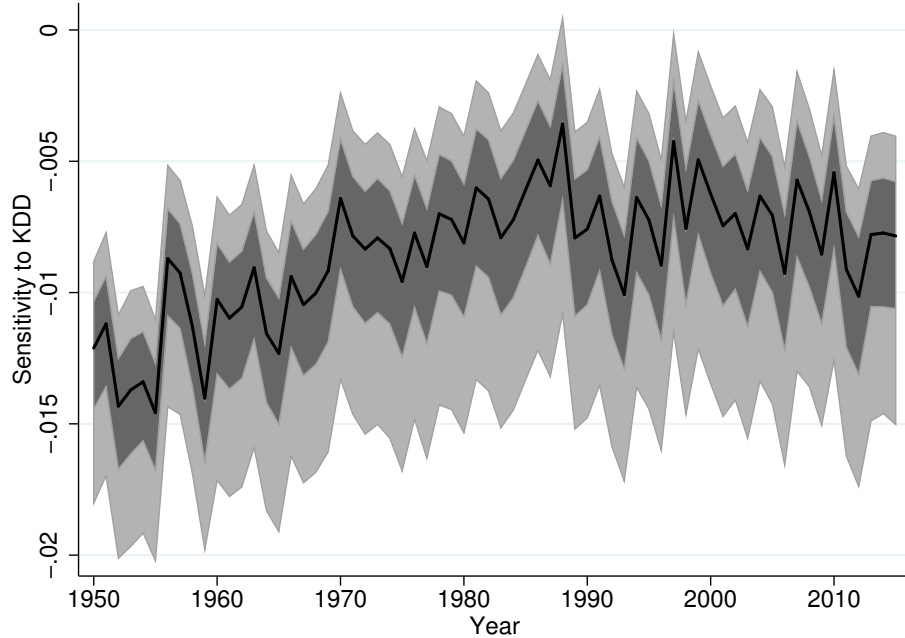
tween counties in the estimates. The 10th to 90th percentiles of coefficients at each year span slightly over -0.01 of space, which is a significant amount. Due to the fact that the lower light grey area, representing the 10th to 25th percentile of coefficients, is significantly wider than the 75th to 90th percentile it is clear that the left tail of the distribution is fat. It is also reassuring that almost all estimated coefficients below the 90th percentile are well below zero.

The trend of the median coefficient also offers an interesting account of aggregate adaptation over time. Table 4 presents a regression of the median coefficient for each year on a linear time trend. While the estimated trend is positive over the sample, an inspection of the graph suggests that it flattens following the late 1980s. A test for a structural break in the parameters of this model is also presented in Table 4, following the methodology of Andrews (1993). The results show that a significant structural break occurred in 1989, which coincides with an extreme drought that occurred in the Midwest of the United States during 1988-89.

Table 4 also presents a regression that accounts for this structural break by incorporating a post-1988 dummy as well as an interaction term between this dummy and the time trend. The fitted line incorporating the structural break is illustrated in Figure 2. Following the structural break the trend of aggregate adaptation is negative. Indeed, the median coefficient in 2015 is similar to that found in the 1970s, suggesting that no progress has been made in adaptation on aggregate over the last four decades.

The results for soy yields are presented in Table 5. This time, the mean coefficient for KDD, whether unweighted or weighted, is very similar to the pooled model estimates in Table 2. Moreover, there is significantly

FIGURE 1. DISTRIBUTION OF KDD SLOPE COEFFICIENTS ACROSS TIME AND COUNTIES FOR U.S. CORN



Note: The black line plots the median coefficient of the KDD variable that is reported in Table 3, while the dark grey area represents the 25th to 75th percentile of coefficients, and the light grey area represents the 10th to 90th percentile of coefficients.

less heterogeneity in the distribution. This can be seen in Figure 3, where the 10th to 90th percentiles of the distribution of KDD slope coefficients comprise a much tighter range than they did with corn.⁵ This is evidence for comparatively little historical adaptation between counties. Using the Spearman correlation coefficient the correlation between $\hat{\beta}_{2it}$ and KDD_{it} is estimated to be 0.25, which is weaker than we observed for corn.

The time trend of aggregate adaptation is somewhat difficult to interpret due to the significant degree of random fluctuation in the median coefficient

⁵The relatively small amount of heterogeneity for soybeans presumably explains why the MO-OLS and pooled results are not very different. This is in sharp contrast to the situation we observed for corn.

TABLE 4—ANALYSING THE MEDIAN COEFFICIENTS TO KDD FOR CORN AND SOY

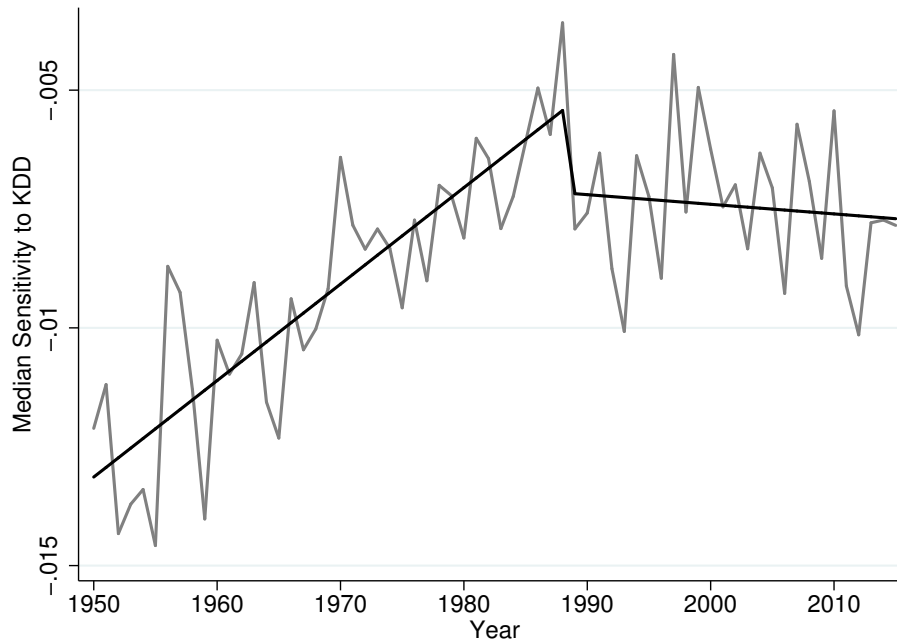
Regression Results		Corn		Soybeans	
t	8.1e-05 (1.3e-05)	0.0002 (2.0e-05)	4.6e-05 (1.3e-05)	1.0e-05 (6.9e-05)	
Constant	-0.0112 (0.0005)	-0.0133 (0.0006)	-0.0066 (0.0011)	-0.0066 (0.0006)	
$d_{t>break}$		0.0070 (0.0020)		0.0023 (0.0014)	
$t^*d_{t>break}$		-0.0002 (4.1e-05)		-9.6e-06 (7.2e-05)	
Structural Break Test	Statistic	p-value	Statistic	p-value	
Supremum Wald	49.32	0.00	7.52	0.24	
Average Wald	31.58	0.00	2.96	0.18	
Supremum LR	38.63	0.00	7.55	0.24	
Average LR	26.83	0.00	3.07	0.17	

Notes: HC3 standard errors are reported in parenthesis. The estimated structural break date in the trend and constant is 1989 for corn and 1972 for soybeans.

between years. Table 4 presents a regression of the median KDD coefficient on a linear trend term. The estimated trend is only very slightly positive, suggesting there has been very little improvement over the last four decades in the aggregate adaptation of soy yield to high temperatures. A structural break test could not reject the null hypothesis that there is no structural break in the parameters of this regression.

In summary, this section has examined the extent of historical adaptation to high temperatures found in the data both between counties and over time. Accounting for heterogeneity was found to be very important in obtaining unbiased estimates of the average sensitivity in the case of corn, especially heterogeneity between counties, as that is where the vast majority of adap-

FIGURE 2. MEDIAN SENSITIVITY TO KDD OVER TIME FOR CORN



Note: The grey line plots the median coefficient of the KDD variable that is reported in Table 3, while the black fitted line is from the regression reported in Table 4.

tation could be found. Furthermore, after accounting for heterogeneity over time it is possible to conclude that there is little evidence to support the hypothesis that the extent of adaptation has notably improved over the last four decades, potentially reflecting a lack of improvements in adaptation technology or a lack of financial incentives to implement them. These results will be used in the next section to better inform forecasts of future damage to crop productivity from severe climate change.

III. Projecting Future Crop Yields

The chief motivation for correctly estimating the sensitivity of crop yield to high temperatures, as well as for measuring the degree of adaptation that

TABLE 5—MEAN-OBSERVATION ESTIMATES OF THE IMPACTS OF TEMPERATURE ON U.S. SOY YIELDS

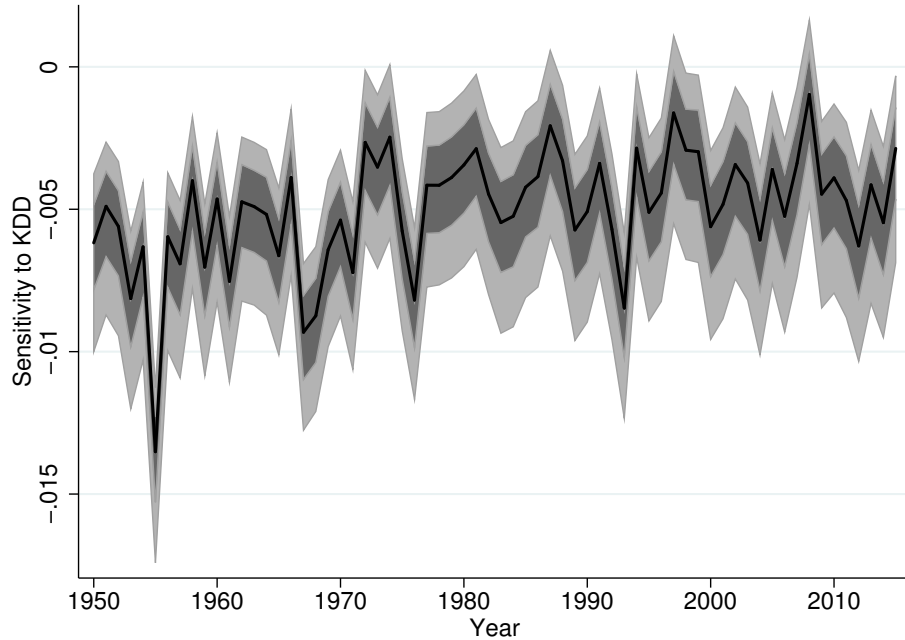
	Mean	Weighted Mean	Median	Standard Deviation	10th Percentile	90th Percentile
GDD	0.0004 (0.0000)	0.0004	0.0004	0.0005	-0.0001	0.0010
KDD	-0.0053 (0.0002)	-0.0048	-0.0049	0.0037	-0.0096	-0.0014
Precipitation	0.0018 (0.0001)	0.0017	0.0015	0.0025	-0.0009	0.0048
Precipitation ²	-1.5e-06 (1.5e-07)	-1.4e-06	1.2e-06	2.1e-06	-3.8e-06	6.2e-07
Constant	1.5328 (0.1440)	1.6671	1.6375	1.9314	-0.8738	3.6623
Obs.	92,373					

Notes: Results exclude states west of the 100th Meridian line. The sample range is 1950-2015, and standard errors are reported in parentheses.

has occurred in the past between regions and over time, is to understand the effect that future climate change will have on crop productivity. Even if national governments embrace the emission abatement targets of the recent Paris Agreement, the world will still undergo up to a further one degree of warming and climate change. Crucial in any forecast is an acknowledgment of the role that future adaptation to high temperatures may have in mitigating the damage of climatic change.

This article presents projections of the impact of further climate change on crop yield in 2050. To do this, temperature and precipitation projections for 2050 are collected for each U.S. county under a range of climate models working under the A1B emissions scenario. The A1B scenario projects a future of rapid economic and technological growth, population that peaks in 2050, and an approach to energy generation that is balanced between

FIGURE 3. DISTRIBUTION OF KDD SLOPE COEFFICIENTS ACROSS TIME AND COUNTIES FOR U.S. SOY



Note: The black line plots the median coefficient of the KDD variable that is reported in Table 3, while the dark grey area represents the 25th to 75th percentile of coefficients, and the light grey area represents the 10th to 90th percentile of coefficients.

fossil intensive and non-fossil energy sources. Among the range of scenarios used by the IPCC, the A1B scenario falls somewhere in the middle in terms of global emission levels at 2050. Using this information, each climate model (or global circulation model) forecasts temperatures on a global grid according to their assumptions on the sensitivity of the climate to these emissions. Importantly, the distribution of temperature changes do vary by climate model, highlighting the need to account for between-county heterogeneity in the temperature sensitivity parameters. For more details on these climate models see Burke and Emerick (2016), who also provided the

climate projections used to conduct this exercise.⁶ Reporting results from a range of climate models enables us to account for model uncertainty in the forecasts.

Projected temperature (using the same GDD/KDD form with a 29°C threshold) and precipitation are plugged into one of the models presented above, and the estimate of future crop yield is compared with the same model’s estimated crop yield in 2005.⁷ The base year was selected to balance the two principles of a year that is as recent as possible, but also a year that provided maximum coverage among the U.S. counties.

More formally, we first calculate a 2050 forecast using the pooled panel estimator:

$$(13) \quad \hat{y}_{mit'} = \hat{\alpha} + \hat{\alpha}_t + \hat{\beta}_1 GDD_{mit'} + \hat{\beta}_2 KDD_{mit'} + \hat{\beta}_3 PREC_{mit'} + \hat{\beta}_4 PREC_{mit'}^2$$

where m is the climate model, i is the county, $t = 2005$, and $t' = 2050$. α_t is the estimated fixed effect for the year 2005. All parameter estimates are obtained from specification (3) in Table 1, for corn, or Table 2 for soybeans. These forecasts are then compared with the estimated 2005 yield values:

$$(14) \quad \hat{y}_{it} = \hat{\alpha} + \hat{\alpha}_t + \hat{\beta}_1 GDD_{it} + \hat{\beta}_2 KDD_{it} + \hat{\beta}_3 PREC_{it} + \hat{\beta}_4 PREC_{it}^2$$

where $t = 2005$. The results of this exercise can be seen under the ‘Pooled’ column in Table 6 for corn and Table 7 for soybeans. The rows of the tables

⁶To avoid potential aggregation bias by comparing the forecasts of the climate models with actual weather station observations (which are interpolated at different resolutions), Burke and Emerick (2016) first calculated the change in temperatures for each county according to the climate model and then added that difference to the actual observations.

⁷The estimate of 2050 crop yield could also be compared to the actual crop yield in 2005, but the results are very similar since the models considered in this section very closely estimate the actual 2005 county average crop yield.

give the names of different climate models, while the forecasted damage to crop productivity is listed in the column labelled ‘Pooled’.

A 2050 forecast of crop yield using MO-OLS is then calculated:

$$(15) \quad \hat{y}_{mit'} = \hat{\alpha}_{it} + \hat{\beta}_{1it}GDD_{mit'} + \hat{\beta}_{2it}KDD_{mit'} + \hat{\beta}_{3it}PREC_{mit'} + \hat{\beta}_{4it}PREC_{mit'}^2$$

and compared to the estimated 2005 yield value:

$$(16) \quad \hat{y}_{it} = \hat{\alpha}_{it} + \hat{\beta}_{1it}GDD_{it} + \hat{\beta}_{2it}KDD_{it} + \hat{\beta}_{3it}PREC_{it} + \hat{\beta}_{4it}PREC_{it}^2$$

where $t = 2005$ and $t' = 2050$. Note that the 2005 coefficients are used for each county when forecasting forward to 2050. This model allows for heterogeneity yet assumes that no adaptation beyond current levels occur in the future, consistent with our findings in Figures 1 and 3. Since adaptation at the aggregate level was not estimated to have changed significantly between 1988 and 2015, changing base years has very little impact on the results. The results of this exercise can be found in Table 6 for corn and Table 7 for soybeans under the ‘MO-OLS’ column.

Comparing the ‘MO-OLS’ column with the ‘Pooled’ column reveals the degree to which accounting for heterogeneity and adaptation in the model affects the forecast of damage from climate change. It is important to note that the percentage changes listed in the table are not the predicted change in crop yield in the future relative to today’s crop yield, as further improvements in farming technology may increase yield in spite of a changing climate, but rather the percentage change of yield with climate change relative to a world without climate change.

The results in Table 6 and Table 7 contain some important insights. First of all, there is a significant amount of variation in the forecast damage to crop productivity between climate models, which is to be expected since the climate models contain significant variation in projected average temperature increase as well as the distribution of that increase across counties. Thus, to better summarise the information, the mean, maximum, and minimum forecast for each column are also provided at the bottom of the table.

Second, removing the heterogeneity bias by using MO-OLS significantly increases the average amount of forecast damage to productivity. The forecast damage to crop yield increases from 29 percent to 44 percent for corn and increases from 35 percent to 55 percent for soybeans relative to the pooled estimates. This increase of more than 50 percent in the damage forecast of both crops is due to the bias found in the pooled estimates, as well as the impact of properly accounting for heterogeneity and adaptation across counties in the forecast.

Finally, we consider two more optimistic scenarios that allow for further adaptation to high temperatures in 2050 relative to 2005. We call these the low adaptation and high adaptation scenarios. These experiments are motivated by the idea that counties with a current high degree of sensitivity to high temperatures could potentially pay a cost to possess a lower sensitivity through adaptation. Specifically, taking the standard deviation of the distribution of slope coefficients on the KDD variable across counties in 2005, we postulate that counties could plausibly move half of one (low adaptation) or one (high adaptation) standard deviation towards zero by 2050. In order to reflect exponentially increasing costs and biological limits to adaptation for counties already highly adapted, we weight the movement

TABLE 6—THE EFFECTS OF CLIMATE CHANGE IN 2050 ON CORN PRODUCTIVITY

Climate Model	Pooled	MO-OLS	Low Adaptation	High Adaptation
CCCMAT63	-27	-48	-42	-36
CNRM	-26	-42	-34	-24
CSIRO	-45	-70	-68	-65
GFDL0	-28	-25	08	57
GFDL1	-29	-35	-15	13
GISSAOM	-40	-62	-59	-55
GISSEH	-34	-57	-53	-49
GISSER	-42	-66	-63	-61
IAP	-32	-52	-47	-41
INMCM	-20	-30	-12	10
IPSL	-04	-08	09	29
UNIF1C	-46	-71	-69	-68
MIROCHIRES	-09	-09	11	36
MIROC MEDRES	-17	-17	06	36
ECHAM	-29	-49	-42	-35
MRI	-36	-60	-57	-53
CCSM	-21	-36	-26	-16
PCM	-39	-63	-61	-58
HADCM3	-27	-40	-27	-13
Mean	-29	-44	-34	-21
Min	-46	-71	-69	-68
Max	-04	-08	11	57

Notes: Results are presented in percentage terms. 2050 temperature forecasts based on the A1B emissions scenario are compared to a base year of 2005.

of the slope coefficient by its proximity to zero.

More formally, consider a forecast of 2050 allowing for further adaptation:

$$(17) \quad \hat{y}_{mit'} = \hat{\alpha}_{it} + \hat{\beta}_{1it}GDD_{mit'} + \left(1 + \frac{s_t}{\gamma(-0.01)}\right) \hat{\beta}_{2it}KDD_{mit'} \\ + \hat{\beta}_{3it}PREC_{mit'} + \hat{\beta}_{4it}PREC_{mit'}^2$$

where $t = 2005$, $t' = 2050$, s_t is the standard deviation of $\hat{\beta}_{2it}$, and $\gamma = 2$ in

the low adaptation scenario or $\gamma = 1$ in the high adaptation scenario.

The results for the low and high adaptation scenarios are presented in the last two columns of Tables 6 and 7. The results for corn imply that a noteworthy portion of the damage due to climate change could potentially be mitigated by further adaptation. The mean damage to productivity decreases from 44 percent to 34 percent under the low adaptation scenario, and decreases further still to 21 percent under the high adaptation scenario. The same cannot be said for soybeans, however, and this is to be expected given the small amount of historical adaptation found in the data. Further adaptation may only decrease the amount of damage to productivity very modestly, from 55 percent under current amounts of adaptation to 52 percent under high adaptation.

It is important to list the limitations of this forecasting exercise before summarising the results. The model estimates do not consider future improvements to farming technology, and accordingly the projected damages to productivity hold constant the productivity of farming that is not related to the climate. Also, the results do not account for changes in planting and harvest dates that may further mitigate damage from climate change. Nor does it account for changes in geographical location of agriculture away from counties that will become too hot to be viable and towards areas that may previously have been too cold. And our forecasts do not consider any major technological innovation in corn seeds and soybeans between now and 2050, as the forecast is relying on the existing trend in time varying heterogeneity to base the projected damage to yield.

It is possible to calculate the damage to yield under a whole suite of assumptions about these sources of adaptation, but until further work is

TABLE 7—THE EFFECTS OF CLIMATE CHANGE IN 2050 ON SOYBEAN PRODUCTIVITY

Climate Model	Pooled	MO-OLS	Low Adaptation	High Adaptation
CCCMAT63	-45	-54	-53	-51
CNRM	-37	-54	-52	-50
CSIRO	-70	-58	-57	-56
GFDL0	17	-60	-55	-51
GFDL1	-09	-58	-54	-50
GISSAOM	-63	-56	-55	-54
GISSEH	-58	-55	-54	-53
GISSER	-68	-57	-56	-56
IAP	-52	-55	-54	-53
INMCM	-08	-55	-53	-50
IPSL	08	-51	-49	-46
UNIF1C	-73	-58	-58	-57
MIROCHIRES	12	-51	-49	-46
MIROCMEDRES	09	-54	-50	-47
ECHAM	-45	-55	-53	-52
MRI	-59	-56	-55	-54
CCSM	-31	-53	-51	-50
PCM	-64	-56	-55	-54
HADCM3	-29	-55	-53	-50
Mean	-35	-55	-53	-52
Min	-73	-60	-58	-57
Max	17	-51	-49	-46

Notes: Results are presented in percentage terms. 2050 temperature forecasts based on the A1B emissions scenario are compared to a base year of 2005.

done these would be merely speculative. The adaptation scenarios that are presented in our scenarios at least have a foundation in the amount of adaptation that already exists between counties. Furthermore, it is important to note that these adaptations, or any others beyond the scope of this paper, may not be financially viable on the part of the farmer or the government. Some forms of adaptation can be very expensive and environmentally costly, such as widespread use of irrigation, and as such it would be unrealistic to

assume that they are possible by 2050, particularly as this would depend on the degree to which prices of agricultural goods increase in the future. Accordingly, there is a large amount of uncertainty attached to any estimate of damage to productivity decades into the future.

In summary, this section has shown the significant degree to which accounting for heterogeneity in the econometric model heightens the expected damage to crop yield from worsening climate change. This occurs for both corn and soybean crops. Positing scenarios of further adaptation indicate that with corn there is potential for further adaptation to mitigate some damage to crop yield at least in the United States. However, the same exercise showed that for soybeans there is much less potential for further adaptation to mitigate damage, unless major improvements in technology were to occur. Regardless, the expected damage to corn and soybeans yields from climate change is substantial.

IV. Conclusion

This paper has attempted to make two contributions to the literature concerned with correctly forecasting the damage to crop yield from climate change. The first contribution was to demonstrate that previous papers that rely on pooled panel estimators are underestimating the true average climate sensitivity of crops due to the presence of heterogeneity bias. Indeed, properly accounting for heterogeneity between counties increases the estimated average sensitivity to high temperatures by around 35 percent, and also meaningfully worsens the expected impact of climate change on yields.

The second contribution was to directly measure the extent of historical

adaptation to climate changes found in the data, both between counties and over time. We argued that a natural way to do this is to represent adaptation as slope coefficient heterogeneity in the econometric model. Estimating these two dimensions of heterogeneity reveals strong evidence for the existence of significant adaptation between counties, but no evidence to suggest that the average amount of adaptation has increased since the 1980s. This may be due to a lack of awareness or the lack of technological progress in the area of adaptation over the last four decades. The extent of adaptation found in the data does suggest that it may be possible to mitigate part of the forecast damage to corn productivity through further adaptation, although it is not clear from the analysis in this paper how costly or feasible such widespread use of adaptation methods would be. The results for soybeans, in contrast, suggest that further adaptation may only marginally mitigate the damage to yield from climate change.

Even with further adaptation, climate change is expected to affect crop yield significantly by 2050. Developing countries with less resources to invest in adaptation are likely to be even more exposed to climatic changes. It is also clear that if runaway climate change were to occur the effects on crop yield by the end of this century would be even more severe. Accordingly, it remains important for governments to also pursue other forms of damage mitigation, particularly emissions reductions, and more drastic adaptation methods such as changing the spatial distribution of agricultural production and further development in cultivars to increase a crop's tolerance to high temperatures.

REFERENCES

- Andrews, D.** 1993. “Tests for Parameter Instability and Structural Change with Unknown Change Point.” *Econometrica*, 61(4): 821–856.
- Burke, M., and K. Emerick.** 2016. “Adaptation to Climate Change: Evidence from US Agriculture.” *American Economic Journal: Economic Policy*, 8(3): 106–140.
- Butler, E., and P. Huybers.** 2013. “Adaptation of US maize to temperature variations.” *Nature Climate Change*, 3: 68–72.
- Darwin, R.** 1999*a*. “A FARMer’s view on the Ricardian approach to measuring agricultural effects of climatic change.” *Climatic Change*, 41: 371–411.
- Darwin, R.** 1999*b*. “The impact of global warming on agriculture: a Ricardian analysis: Comment.” *American Economic Review*, 89(4): 1049–1052.
- Deschenes, O., and M. Greenstone.** 2007. “The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather.” *American Economic Review*, 97(1): 354–396.
- Kawasaki, K., and U. Shinsuke.** 2016. “Quality Matters more than quantity: asymmetric temperature effects on crop yield and quality grade.” *American Journal of Agricultural Economics*, 98(4): 1195–1209.
- Lobell, D., and M. Burke.** 2008. “Why are agricultural impacts of climate change so uncertain? The importance of temperature relative to precipitation.” *Environmental Research Letters*, 3(3): 1–8.

- Lobell, D., M. Banziger, C. Magorokosho, and B. Vivek.** 2011. “Nonlinear heat effects on African maize as evidenced by historical yield trials.” *Nature climate change*, 1: 42–45.
- Mendelsohn, R., W. Nordhaus, and D. Shaw.** 1994. “The impact of global warming on agriculture: a Ricardian analysis.” *American Economic Review*, 84(4): 753–771.
- Neal, T.** 2016. “Multidimensional Parameter Heterogeneity in Panel Data Models.” *UNSW Research Working Paper No. 2016-15*.
- Pesaran, M. H., and R. Smith.** 1995. “Estimating long-run relationships from dynamic heterogeneous panels.” *Journal of Econometrics*, 68: 79–113.
- Schlenker, W., and M. Roberts.** 2009. “Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change.” *Proceedings of the National Academy of Sciences*, 106(37): 15594–15598.
- Schlenker, W., M. Hanemann, and A. Fisher.** 1994. “Will U.S. Agriculture really benefit from global warming? Accounting for irrigation in the Hedonic Approach.” *American Economic Review*, 84(4): 753–771.