



Business School / School of Economics

# UNSW Business School Working Paper

UNSW Business School Research Paper No. 2018 ECON 08

## **The Impact of Climate Change on US Agriculture: The Roles of Adaptation Techniques and Emissions Reductions**

Michael Keane  
Timothy Neal

This paper can be downloaded without charge from  
The Social Science Research Network Electronic Paper Collection:  
<https://ssrn.com/abstract=3180480>

# The Impact of Climate Change on U.S. Agriculture: The Roles of Adaptation Techniques and Emissions Reductions

By MICHAEL KEANE AND TIMOTHY NEAL\*

Draft: May 17, 2018

*We investigate the impact of climate change on U.S. agricultural productivity using county-level yield and weather data from 1950 to 2015. We present two new methods of modelling how producers adapt agricultural techniques to harsh temperatures, including a new panel data estimator that allows for two-dimensional fixed-effects in slopes. We find evidence of adaptation to geographic and temporal variation in climate, but it has stalled since 1989. We show that adaptation implies fixed-effects slope heterogeneity in the relationship between crop yield and temperature, and ignoring this leads to biased estimates of temperature sensitivity. We use our estimates to project corn yields to 2100 using a variety of climate models and emission scenarios, and find that unmitigated climate change will have severe effects on yields. Our models indicate that adaptation techniques can mitigate 10 to 45% of the damage, but significant emissions reductions can mitigate far more (i.e., 42 to 91%).*

*JEL: C23, C54, D24, Q15, Q51, Q54, Q55*

Leading scientific and environmental institutions warn of the extreme impact that future climate change may have on the productivity of agriculture and global food supply. Temperatures in excess of certain thresholds can significantly decrease crop yield either by directly damaging the plant tissue or enzymes through heat stress, or by encouraging water stress which inhibits photosynthesis. However, studies that attempt to estimate the sensitivity of agriculture to the climate, in hopes of obtaining insight into the effects of future climate change, produce very mixed results. The forecast impact of climate change

\* Keane: University of New South Wales, Sydney, NSW 2000 (email: m.keane@unsw.edu.au); Neal: Department of Economics, University of New South Wales, Sydney, NSW 2000 (email: timothy.neal@unsw.edu.au).

on U.S. agriculture in particular ranges from severe damage to productivity (e.g. Schlenker and Roberts 2009) and crop quality (Kawasaki and Shinsuke 2016) to very minor damage or even net benefits (e.g. Mendelsohn, Nordhaus and Shaw 1994). Resolving uncertainty about the relationship between crop yield and the climate is one of the top priorities for improving climate change impact assessments (Lobell and Burke 2008).

The impact of climate change on crop yield will depend critically on the degree of adaptation by agricultural producers. Adaptation to climate change may involve switching crop cultivars to ones that are more heat tolerant, increasing water retention in fields, use of irrigation and/or fertilizers, altering planting or harvesting dates, shifting the spatial distribution of agricultural production, or switching crops entirely. As we show, the very existence of adaptation implies slope heterogeneity in the historical relationship between crop yield and temperature. Furthermore, this heterogeneity is positively correlated with the extent of high temperatures in a region or time period. Such fixed effects in slopes creates important challenges for the proper econometric modelling of climate impacts.

In this article we investigate these issues using temperature and crop yield data for U.S. counties from 1950 to 2015. We focus on corn (maize) and soybean yields, as these are the two largest crops in the United States in terms of tonnage. Thus, they are important for national and indeed global food supply. Prior to our work, several authors have also investigated agricultural adaptation using panel data regressions of crop yields on temperature and precipitation (see, e.g., Burke and Emerick 2016, Butler and Huybers 2013, Schlenker and Roberts 2009), but we extend this work in two important ways:

Our first major contribution is to propose two direct methods to estimate the extent of historical adaptation to harsh temperatures. The literature contains no established method of measuring adaptation. The first method we propose involves augmenting a standard panel data crop-yield regression, where temperature bands enter linearly, with a specific non-linear functional form for harsh temperatures motivated by theory. This functional form allows the sensitivity of crop yield to harsh temperatures to decline as the frequency of harsh temperatures increases. It has no direct parallel in the previous literature on agricultural productivity, but it is related to methods of estimating adaptation to climate change in other spheres such as mortality (see Deschenes and Greenstone 2011)

and economic growth (see Dell, Jones and Olken 2012).

The second method we propose involves directly estimating the heterogeneity of the slope coefficient (on harsh temperatures) between counties and over time, which can be viewed as an extension of the approaches in Butler and Huybers (2013) and Roberts and Schlenker (2010). This method imposes less structure on the adaptation process, and thus it provides a way of validating the theoretically-motivated functional form assumed in the first approach. It is made possible by a new panel data estimator, developed in Neal (2018), that allows for two-dimensional fixed effects in slopes. This new estimator, which is applied here for first time, is described in detail in Appendix B.

The results of both methods suggest that significant adaptation occurred between counties, and over time from 1950 to 1989. However, adaptation has subsequently stalled, and average sensitivity to harsh temperatures has slightly worsened over the last four decades. This supports the conclusion of Lobell et al. (2014) that the trend towards higher sowing densities (to produce more yield) is compounding the water stress from excessive heat.

We also find that ignoring adaptation in econometric models of crop yield leads to significant underestimation of yield sensitivity to harsh temperatures. This is because the heterogeneity in the slope coefficient on harsh temperatures is positively correlated with the frequency of harsh temperatures (as hotter growing seasons generate financial incentives to adopt adaptation measures). Our results suggest that standard panel data estimators underestimate the heat sensitivity of crop yields by 35 to 40 percent.

Our second major contribution is to provide a comprehensive series of forecasts of corn and soybean yield under various climate change scenarios. The whole purpose of estimating the degree of adaptation to climatic change in the past is to gain insight into the extent to which adaptation can mitigate the damage from climate change in the future. We forecast crop yield annually from 2016 to 2100 using forecasts of temperature and precipitation from nineteen climate models under three CO<sub>2</sub> emissions growth scenarios. To our knowledge no prior study has forecasted crop yield from the present until 2100 using a range of climate model outputs across multiple emission growth pathways. The wide range of climate models that we consider allows us to quantify the variability in outcomes between models. And our econometric models allow us to isolate and compare the effectiveness of

adaptation vs. emissions reduction as ways to mitigate crop damage.

The UN-IPCC uses a set of “representative concentration pathways” (RCPs) for CO<sub>2</sub> emissions growth based on different policy scenarios. Our first prediction is that “business as usual” emissions growth (known as the RCP85 scenario), along with no adaptation to climate change, will cause very severe damage to corn yield - a 70% reduction by 2100 (80% prediction interval from 50 to 90%). Of the total cumulative damage over 2020-2100, adaptation may avert from 10 to 45%. These prediction intervals account for uncertainty across both econometric and climate models. Clearly, these results imply that the U.S. cannot rely on adaptation alone to protect corn yields from substantial impacts.

In comparison, “moderate” emissions reductions (the RCP45 scenario, which is slightly more ambitious than current government pledges) can avert 26 to 82% of the total damage without any adaptation, or 33 to 90% with adaptation. Substantial emissions reductions (the RCP26 scenario, consistent with the most ambitious targets under the Paris agreement) can avert 42 to 91% of the damage. Under all three scenarios, most of the variability in outcomes arises from disagreement between climate models, not econometric models.

Thus, we find that adaptation is less effective than emissions reductions in averting damage to crop yield, yet nevertheless that adaptation can have a meaningful impact on mitigating yield damage (contrary to the conclusions of several previous studies). Indeed, adopting both emissions reductions and adaptation will be necessary to avert most of the potential damage from climate change.

So far, we have discussed forecasts holding agricultural technology fixed at current levels (aside from adaptation). In Appendix C we report forecasts of crop yield that incorporate projections of technical progress in farm technology (based on historical trends). The results suggest that even under an optimistic scenario for technical progress, climate change still has a dramatic effect on the evolution of crop yields. Whether yields increase, stagnate, or decline over the coming decades will depend on the particular combination of emissions reductions, technological progress, and adaptation efforts that materialize.

Finally, in Appendix D we present our results for soybeans. We find that, without emissions reductions or adaptation, yields would drop by 34 to 76% by 2100. Strikingly, we find that adaptation is rather ineffective, reducing damage by only 3 to 22%, while

emissions reductions are very effective - the moderate RCP45 scenario reduces damage by 36 to 88%. Most of this prediction variability is across climate models, not econometric models. The sharp difference between the corn and soybean results reflects the different growing areas for these crops (as these areas are affected differently by climate change).

The rest of the paper is structured as follows. The first section presents a theoretical model of agricultural output with weather and adaptation, in order to provide a coherent framework for the empirical work. The second section outlines previous approaches to this topic, why pooled estimates of climate sensitivity are likely to be biased, the econometric methodology that we adopt, and finally the sources for our climate and yield data. The third section presents our econometric results for corn yield. The fourth section presents our forecasting exercise for corn yield from 2016 to 2100. The fifth section concludes, while four appendices extend the results of the article.

### I. A Model of Agricultural Output with Weather and Adaptation

To guide our analysis, we require a model of agricultural output that includes not only the standard factors of production (e.g., land, capital, labor, fertilizer), but also variables that capture the influence of climate operating through temperature and precipitation. Thus, we use a modified Cobb-Douglas production function to model crop yield:<sup>1</sup>

$$(1) \quad Y_{it}/C_{it} = A_t I_{it}^{\delta} (1 + \beta_1(GDD_{it} - GDD_{min}) + \beta_{2it}KDD_{it})$$

where  $Y_{it}$  is the output of a specific crop for farmer  $i$  at time  $t$  and  $C_{it}$  is the number of acres planted for that crop.  $A_t$  is the state of farming technology at time  $t$ , and  $I_{it}$  represents a composite of the conventional inputs into production (e.g., land, capital, labor, fertilizer), while parameter  $\delta$  captures returns to scale. The particular functional form of the aggregator  $I_{it}$  is not relevant for our purposes, so we leave it unspecified. The key variables that capture the influence of temperature are  $GDD$  and  $KDD$ . Specifically,  $GDD_{it}$ , or “growing degree days,” represents the number of hours in the growing season that the crop experiences beneficial temperature, and  $KDD_{it}$ , or “killing degree days,” is

<sup>1</sup>Using a more comprehensive production function, or incorporating more realistic features such as expectations or dynamics, will not change the basic results of this section.

the number of hours in the growing season that the crop experiences harmful temperatures.  $GDD_{min}$  represents the minimum number of hours of beneficial temperature necessary to obtain any positive yield. (Precipitation is excluded here for simplicity, but will be included in the econometric models.)

To account for the possibility of adaptation, the sensitivity of output to temperature should (at least in part) be a choice variable for the farmer. Suppose the farmer can affect the sensitivity of output to harmful temperatures,  $\beta_{2it}$  by adopting a range of adaptation techniques. We represent this adaptation effort by the continuous variable  $\alpha_{it}$ :

$$(2) \quad \beta_{2it} = s/(1 + \alpha_{it})$$

where  $s < 0$  is the effect of harmful temperatures ( $KDD_{it}$ ) on output in the absence of any adaptation.  $\alpha_{it}$  represents units of adaptation technology, which mitigate the harmful effects of  $KDD_{it}$  on crop production. Profit for farmer  $i$  at time  $t$  is:

$$(3) \quad \pi_{it} = pY_{it} - \gamma\alpha_{it} - rI_{it}$$

where  $\gamma$  is the price of a unit of adaptation technology,  $p$  is the market price for the crop, and  $r$  is the exogenous rental rate per unit of  $I_{it}$ .<sup>2</sup> In order to maximize profit, the farmer will invest in the optimal level of adaptation. Setting  $\partial\pi/\partial\alpha = 0$  we obtain:

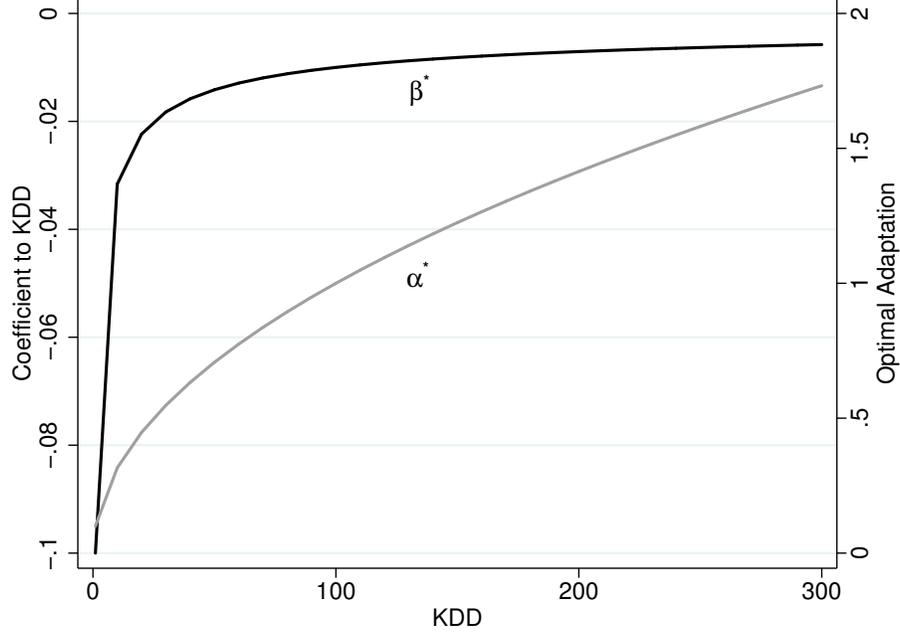
$$(4) \quad \alpha_{it}^* = \sqrt{\frac{-pC_{it}A_tI_{it}^\delta sKDD_{it}}{\gamma}} - 1$$

Clearly, as  $KDD_{it}$  increases the optimal level of adaptation also increases. Farms in hotter regions or time periods have more incentive to adapt. Figure 1 graphs the optimal level of adaptation  $\alpha_{it}^*$  and the resulting coefficient  $\beta_{2it}$  against values of  $KDD_{it}$ . As we see, the relationship between  $KDD_{it}$  and the heterogeneous coefficient  $\beta_{2it}$  closely resembles a log-linear function. Of course, this specific form may hinge on the functional form in (2). However, this prediction of the simple model is testable, and it is supported

<sup>2</sup>Recall  $I_{it}$  is a composite of inputs such as capital and labor, and for our purposes it is unnecessary to elaborate on its form. Given homotheticity,  $r$  can be interpreted as the constant marginal cost of the optimal bundle of inputs.

by our econometric analysis. More importantly, the model shows how the very existence

FIGURE 1. RELATIONSHIP BETWEEN  $\alpha_{it}^*$ ,  $\beta_{2it}^*$ , AND  $KDD_{it}$



Note: This graph presents a hypothetical scenario of  $\beta_{2it}^*$  and  $\alpha_{it}^*$  (secondary vertical axis) as a function of  $KDD_{it}$  where  $p = C_{it} = A_t = I_{it}^\delta = 1$  and  $s = -0.01$ .

of adaptation renders the coefficient on  $KDD_{it}$  heterogeneous, and that the heterogeneity will generally be correlated with the regressor itself. This has important implications for the proper econometric methods to use in estimating adaptation.

A natural way to estimate equation (1) is to take the natural logarithm of crop yield:

$$y_{it} = \ln(A_t I_{it}^\delta) + \ln(1 + \beta_1(GDD_{it} - GDD_{min}) + \beta_{2it}KDD_{it})$$

where  $y_{it} = \ln(Y_{it}/C_{it})$ . Then, using the approximation  $\ln(1 + x) \approx x$ , we obtain:

$$(5) \quad y_{it} = \ln(A_t) + (\ln(I_{it}^\delta) - \beta_1 GDD_{min}) + \beta_1 GDD_{it} + \beta_{2it} KDD_{it}$$

where it is clear that two-way fixed effects will be important to capture unobserved het-

erogeneity in  $A_t$  and  $I_{it}$ , and slope heterogeneity will be important to identify  $\beta_{2it}$ .

Below we use versions of (5) to study crop yield among U.S. counties. This a level of aggregation above the farmer model of this section. But in Appendix A we show that the main predictions of the farmer model (i.e., that  $\beta_{2it}$  is negatively correlated with  $KDD_{it}$  and the relationship is approximately log-linear) do carry over to a county-level model.

## II. Empirical Approach

### A. Previous Approaches

Several recent papers attempt to estimate the impact of climate change on crop yields (production per acre) using what we call the ‘GDD’ approach. This recognizes that temperature can be either beneficial or harmful to a crop. While moderate temperatures are beneficial, excessively high temperatures hamper photosynthesis, predominantly by increasing the need for soil water to sustain carbon assimilation, and by increasing the rate of transpiration (which drains the plant’s water supply). Both factors contribute to water stress by increasing the vapour pressure deficit (see Lobell et al. (2013) for details). High temperature can also damage plant tissue directly through heat stress. This analysis motivates separation of temperature into beneficial and harmful bands based on a threshold value. This is typically set at 29 degrees Celsius for both corn and soybeans.

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:

$$(6) \quad y_{it} = c_i + c_t + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

where  $y_{it}$  is the log of crop yield for county  $i$  in year  $t$ . The literature typically operates at the county level, as that is where data is more widely available.

To proceed, recall that  $GDD_{it}$  measures hours of beneficial temperatures, while  $KDD_{it}$  measures hours of harmful temperatures.  $PREC_{it}$  is total precipitation during the growing season.  $c_i$  captures intercept heterogeneity that is county  $i$  specific (such as soil quality), while  $c_t$  captures intercept heterogeneity that is constant across counties but varies between

years  $t$  (such as farm technology). These fixed effects are important to disentangle effects of weather from possibly omitted inputs. The key parameter of interest is  $\beta_2$ , which captures the extent to which high temperatures reduce crop yield.<sup>3</sup>

Several studies use the GDD approach to estimate the extent of adaptation to higher temperatures. The idea is that the extent of historical adaptation informs us about the scope for future adaptation. For instance, Schlenker and Roberts (2009) test for evidence of historical adaptation to high temperatures by running regressions like (6) after splitting the sample into northern and southern U.S. states and also by 1950-1977 and 1978-2005 periods. They find the coefficients do not differ significantly by region or time, which they take as evidence that historical adaptation has been very limited.

In contrast, Butler and Huybers (2013) run separate regressions by county for 1981 to 2008, and conclude from the county-specific coefficients that there is adaptation occurring between counties with different climates. When forecasting the effect of climate change on crop yields they argue that yield losses could be halved by this type of adaptation. But Roberts and Schlenker (2010) investigate whether corn and soybeans are becoming more or less sensitive to high temperatures over time by estimating a time varying parameter model, and conclude there is little evidence that heat sensitivity is trending up or down over time (i.e. little evidence for adaptation).

Burke and Emerick (2016) adopt what they call a ‘long differences’ approach to measuring the extent of historical adaptation. They estimate (6) 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:

$$(7) \quad \begin{aligned} \Delta y_{is} = & c_s + \beta_1 \Delta GDD_{is} + \beta_2 \Delta KDD_{is} + \beta_3 \Delta PREC_{is} \\ & + \beta_4 \Delta PREC_{is}^2 + \Delta \epsilon_{is} \end{aligned}$$

<sup>3</sup>An alternative way to estimate 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 climate and agricultural land value. By linking climate to land value, as opposed to crop yield, these studies can examine the entire agricultural sector. Furthermore, they can account for a broad range of adaptations and shifts in farmer behaviour, such as switching crops or changes in land use.

However, the hedonic approach does rely on some strong assumptions. It assumes output and input prices remain constant, which is problematic if there are climate-induced price changes (see Darwin 1999a). 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 can lead one to underestimate damage from climate change (Schlenker, Hanemann and Fisher 1994, Darwin 1999b).

where  $\Delta GDD_{is}$  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  $1 - \hat{\beta}_2^{LD} / \hat{\beta}_2^{FE}$ , 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.

It is important to recognize that both the GDD and long difference approaches rely on a standard panel fixed effects model to provide consistent estimates of the sensitivity of crop yield to variation in temperature. Subsection B demonstrates why this assumption is unlikely to hold in practice, while subsection C presents our alternative approach.

Finally, outside of agricultural production, there are several studies that attempt to model adaptation by interacting weather variables with the average climate of a region. This captures adaptation by allowing the marginal effect of high temperatures to depend on climate. Deschenes and Greenstone (2011) use this approach in to estimate effects of climate change on mortality, and Dell, Jones and Olken (2012) use it to estimate effects of climate change on economic growth. Neither paper finds evidence of systematic heterogeneity in the marginal effect that would be indicative of adaptation.

### B. Bias in fixed effects regressions under slope heterogeneity

In this section we demonstrate how adaptation can lead to bias in conventional estimates of crop yield equations. As we showed in Section I, the use of adaptation techniques will generate heterogeneity in the parameter  $\beta_2$  that captures sensitivity to high temperatures in (6). This heterogeneity may occur across both the  $i$  and  $t$  dimension as in (2). We can capture this multidimensional heterogeneity by applying an additive structure  $\beta_{2it} = \lambda_2 + \lambda_{2i} + \lambda_{2t}$ , and then generalize the conventional specification (6) as follows:

$$(8) \quad y_{it} = c_i + c_t + \beta_1 GDD_{it} + \beta_{2it} KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

Note that this specification allows for additive random or fixed effects in the sensitivity of crop yield to temperature, as well as unit and time fixed effects in the intercept.

To see the bias of a fixed effects estimator in a model containing adaptation through

slope heterogeneity, first simplify (8) by excluding precipitation and stacking the variables:

$$(9) \quad y_{it} = c_i + c_t + \mathbf{z}'_{it}\boldsymbol{\theta}_{it} + \epsilon_{it}$$

where  $\mathbf{z}'_{it} = (GDD_{it}, KDD_{it})$  and  $\boldsymbol{\theta}_{it} = (\beta_1, \beta_{2it})' = (\beta_1, \lambda_2 + \lambda_{2i} + \lambda_{2t})'$ . Consider a two-way within transformation of (9) to remove the fixed effects:

$$(10) \quad \tilde{y}_{it} = \tilde{\mathbf{z}}'_{it}\boldsymbol{\theta} + v_{it}$$

where  $\tilde{y}_{it} = y_{it} - N^{-1} \sum_{i=1}^N y_{it} - T^{-1} \sum_{t=1}^T y_{it} + NT^{-1} \sum_{i=1}^N \sum_{t=1}^T y_{it}$  is the two-way within transformation of  $y_{it}$ , similarly for  $\tilde{\mathbf{z}}_{it}$ ,  $\boldsymbol{\theta} = (\beta_1, \lambda_2)'$ , and  $v_{it}$  is defined as:

$$(11) \quad v_{it} = \widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_i} + \widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_t} + \tilde{\epsilon}_{it}$$

where  $\widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_i} = \mathbf{z}'_{it}\boldsymbol{\theta}_i - N^{-1} \sum_{i=1}^N \mathbf{z}'_{it}\boldsymbol{\theta}_i - T^{-1} \sum_{t=1}^T \mathbf{z}'_{it}\boldsymbol{\theta}_i + NT^{-1} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}'_{it}\boldsymbol{\theta}_i$  is the two-way within transformation of  $\mathbf{z}'_{it}\boldsymbol{\theta}_i$ , similarly for  $\widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_t}$ ,  $\boldsymbol{\theta}_i = (0, \lambda_{2i})'$ , and  $\boldsymbol{\theta}_t = (0, \lambda_{2t})'$ .

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

$$(12) \quad \hat{\boldsymbol{\theta}} = \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\tilde{\mathbf{z}}'_{it} \right)^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\tilde{y}_{it} \right)$$

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

$$(13) \quad \begin{aligned} \hat{\boldsymbol{\theta}} = & \boldsymbol{\theta} + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_i} \right) \\ & + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\widetilde{\mathbf{z}'_{it}\boldsymbol{\theta}_t} \right) + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\tilde{\epsilon}_{it} \right) \end{aligned}$$

where  $\mathbf{Q}_{zz,NT}^{-1} = \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{\mathbf{z}}_{it}\tilde{\mathbf{z}}'_{it} \right)^{-1}$ . If  $\mathbf{z}_{it}$ ,  $\boldsymbol{\theta}_i$ , and  $\boldsymbol{\theta}_t$  are not independent the two bias terms generated by the slope heterogeneity (i.e., the terms involving  $\boldsymbol{\theta}_i$  and  $\boldsymbol{\theta}_t$ ) will not vanish as sample size increases, and the fixed effects estimator will be inconsistent.

There is every reason to believe that the  $KDD_{it}$  component of  $\mathbf{z}_{it}$  is not independent of  $\boldsymbol{\theta}_i$  and  $\boldsymbol{\theta}_t$ . The theoretical model presented in Section I predicts a positive relationship

between the two, which will result in the standard FE-OLS estimator overestimating (i.e. estimates closer to zero in this case) the average sensitivity to harsh temperatures  $\beta_2$ . This is intuitive as farmers who experience hotter temperatures in their region or time period have more incentive to adopt adaptation techniques. Indeed, Butler and Huybers (2013) provide evidence that the estimate of  $\beta_2$  for an individual county is significantly positively correlated with its relative experience of heat. Our key aim is to more precisely ascertain the nature of the relationship between  $KDD_{it}$  and  $\beta_{2it}$ , and to use this information to help predict the impact of climate change and how it may be mitigated by adaptation.

### C. Econometric Methodology

In this section we propose two complementary approaches to modelling the degree of adaptation to heat that is present in the data. A useful baseline of comparison for both our new approaches is the conventional two-way fixed effects estimator (FE-OLS). Consider a fixed effects regression of crop yield as a function of temperature and precipitation:

$$(14) \quad y_{it} = c_{it} + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

where  $c_{it} = c_i + c_t$ . Since the model includes both county and year fixed effects, parameters are identified from deviations in county weather from the county average, after removing annual weather shocks and/or technology trends that are common to all counties.

In Section II.B we showed that the FE-OLS estimate of  $\beta_2$  is likely to be biased upward, because it ignores the correlation between  $KDD_{it}$  and  $\beta_{2it}$  generated by adaptation to harsh temperatures. However, it is possible to incorporate adaptation in (14). In particular, as shown in Figure 1 our simple theoretical model implies the marginal effect of harsh temperatures on crop yield will be a log-linear function of  $KDD_{it}$ . Thus, if we estimate:

$$(15) \quad y_{it} = c_{it} + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 (\ln(KDD_{it}) * KDD_{it} - KDD_{it}) + \beta_4 PREC_{it} + \beta_5 PREC_{it}^2 + \epsilon_{it}$$

by FE-OLS, then the marginal effect of  $KDD_{it}$  on  $y_{it}$  will be:

$$(16) \quad \frac{\partial y_{it}}{\partial KDD_{it}} = \hat{\beta}_2 + \hat{\beta}_3 \ln(KDD_{it})$$

so we see that the marginal effect of  $KDD_{it}$  is a log-linear function of  $KDD_{it}$  itself. This model allows us to estimate the extent of adaptation to climate, and furthermore, we can also calculate the weighted average effect of  $KDD_{it}$  on  $y_{it}$  and compare it to the estimate of  $\beta_2$  in the standard FE-OLS model in (14) to determine if there is evidence of bias.

Instead of imposing a particular relationship between  $KDD_{it}$  and its marginal effect, another approach to modelling adaptation is to estimate a model with slope heterogeneity. But in order to capture adaptation it is essential to allow the heterogeneity to be correlated with the regressors - standard estimators that only allow for random effects in slopes are not appropriate. Our estimator allows for fixed effects in intercepts and slopes, as follows:

$$(17) \quad y_{it} = c_i + c_t + \beta_{1it}GDD_{it} + \beta_{2it}KDD_{it} + \beta_{3it}PREC_{it} + \beta_{4it}PREC_{it}^2 + \epsilon_{it}$$

Estimating the heterogeneity directly provides a set of slope coefficients, and analyzing the distribution of  $\hat{\beta}_{2it}$  post-estimation will determine the nature of the relationship between  $KDD_{it}$  and  $\hat{\beta}_{2it}$  (if any). Indeed, we will fit a mapping from the regressor  $KDD_{it}$  to the heterogeneous slope, and use this to model adaptation in response to climate change in the forecasting section. The average of  $\hat{\beta}_{2it}$  will also provide further evidence on whether the estimate of  $\beta_2$  in the standard FE-OLS model is upwardly biased.

Of course, (17) involves more fixed-effects than data points. But if we assume heterogeneity is additive across county/time dimensions, such that  $\beta_{2it} = \lambda + \lambda_i + \lambda_t$  as in (8), we can estimate (17) using mean-observation OLS (or ‘MO-OLS’), introduced in Neal (2018). Crucially, MO-OLS allows the slope heterogeneity to be correlated with the regressors, and allows for county/year fixed effects. Briefly, the idea behind the estimator is to obtain a pooled estimate, say generically  $\hat{\beta}$ , then run a series of regressions for each individual to collect  $\hat{\beta}_i$ , then a series of regressions for each time period to collect  $\hat{\beta}_t$ , and then construct  $\hat{\beta}_{it} = \hat{\beta}_i + \hat{\beta}_t - \hat{\beta}$  as well as an expression to remove remaining biases. Appendix B provides a complete exposition, while asymptotic and Monte Carlo results can be found in Neal

(2018). Consistency relies on large  $N$  and  $T$ , which is appropriate given our data.

Our two approaches to modelling adaptation are complementary. The FE-OLS approach allows the marginal effect of  $KDD_{it}$  on  $y_{it}$  to vary log-linearly with  $KDD_{it}$ , while MO-OLS is less restrictive and allows the heterogeneity to also be related to unobserved variables.

Finally, it is important to be clear about the types of adaptation these approaches capture. Modelling adaptation as heterogeneity in the marginal effect of  $KDD_{it}$  on  $y_{it}$  will pick up all forms of adaptations related to the individual crop, whether short or long-term in nature. This includes irrigation, cultivar adoption, fertilizers, improving water retention, other techniques, and any technological advancement in these methods over the sample period. But our methods will not capture other forms of adaptation such as changing the planting and harvest date for the crop, crop switching, or land use changes. These forms of adaptation may prove to be important, especially when climate change becomes more severe. We leave investigation of these issues to future work.

#### *D. Data and Variable Construction*

The historical temperature and precipitation data we use in this study are from Schlenker and Roberts (2009). It contains daily observations on maximum and minimum temperature, and precipitation, on a grid across the continental U.S. from 1950 to 2015.<sup>4</sup> It also contains code to map the grid-based data onto counties, weighting grid locations by the location of agricultural production in each county. From the daily maximum and minimum observations, we can approximate the hours each day that a crop is exposed to one-degree Celsius temperature intervals using a sinusoidal function (consistent with the literature):

$$(18) \quad DD_C = \begin{cases} 0 & \text{if } C > T_{max} \\ T_{avg} - C & \text{if } C < T_{min} \\ \frac{((T_{avg}-C)\cos^{-1}(S)+(T_{max}-T_{min})\sin(S)/2)}{\pi^{-1}} & \text{otherwise} \end{cases}$$

where  $C$  is the temperature in Celsius,  $T_{max}$  and  $T_{min}$  are the supplied daily maximum

<sup>4</sup>The raw data comes from the PRISM Climate Group. Schlenker and Roberts subsequently extended the dataset to 2015 after the publication of their paper. Temperature observations prior to 1950 are available, but it is harder to convert them to national grid as fewer weather stations were operational.

and minimum temperature,  $T_{avg} = \frac{T_{max} + T_{min}}{2}$  and  $S = \cos^{-1}\left(\frac{2C - T_{max} - T_{min}}{T_{max} - T_{min}}\right)$ .

The temperature threshold that separates GDD from KDD is 29 degrees Celsius for both corn and soybeans, again following the literature.<sup>5</sup> Thus, using results from (18) we calculate daily values of  $GDD$  and  $KDD$  as follows:

$$(19) \quad GDD_{id} = DD_0 - DD_{29}, \quad KDD_{id} = DD_{29}$$

for each county  $i$  and day  $d$ . These are aggregated to yearly values of  $GDD_{it}$  and  $KDD_{it}$  by summing over the days of the growing season. For both corn and soybeans we assume this is May 1st to September 30th (in line with the literature). Allowing for growing seasons to vary between counties and over time did not meaningfully change the econometric results, and it presents complications for forecasting. Precipitation data is measured in terms of inches and is also aggregated across the growing season from the daily data.

Following the literature, we exclude states west of the 100th Meridian from the sample as they are more reliant on irrigation (although including them does not meaningfully affect parameter estimates).<sup>6</sup> Our regressions are weighted based on the average crop area in each county, and standard errors are clustered at the state level.

Crop yield data is from the United States Department of Agriculture's (USDA) National Agricultural Statistics Service. The information is at the county-level and covers the same annual sample period of 1950 to 2015, although not all counties contain crop yield data across all years (creating an unbalanced panel). The USDA also provides the average crop area of each county, which we use for regression weights.

### III. Results

#### A. Pooled Models

Table 1 presents results from corn yield regressions pooling data over counties and time. They incorporate county/time fixed effects but do not allow for heterogeneous temperature coefficients. Thus, we refer to these as FE-OLS models. The table lists four specifications,

<sup>5</sup> Butler and Huybers (2013) note that while 29 degrees may appear low as a threshold for damaging temperatures, the temperature experienced by the plant itself is higher than the air temperature above the crop canopy.

<sup>6</sup>Note: The 100th Meridian separates the Great Plains to the east from the semi-arid lands to the west.

where the first includes year fixed effects, the second and fourth include county/year fixed effects, and the third includes county fixed effects and state-specific quadratic time trends. The first three specifications do not model adaptation. The fourth specification, which is motivated by the theory in II.C, includes as an additional regressor a nonlinear function of KDD that is designed to capture adaptation to harsh temperatures (see equation (15)).

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

Specification	(1)	(2)	(3)	(4)
GDD	0.0002 (0.0001)	0.0003 (0.0001)	0.0003 (0.0001)	0.0004 (0.0001)
KDD	-0.0053 (0.0012)	-0.0063 (0.0007)	-0.0063 (0.0005)	-0.0159 (0.0024)
$\ln(\text{KDD}) * \text{KDD} - \text{KDD}$				0.0022 (0.0005)
Precipitation	0.0013 (0.0009)	0.0010 (0.0002)	0.0013 (0.0003)	0.0007 (0.0002)
Precipitation <sup>2</sup> (x 10,000)	-0.0001 (0.0001)	-0.0001 (0.0000)	-0.0001 (0.0000)	-0.0001 (0.0000)
Constant	2.7941 (0.6460)	2.7188 (0.2661)	2.4767 (0.2123)	2.1890 (0.2383)
Fixed Effects	Yr	Cty, Yr	Cty, State-Yr	Cty, Yr
Obs.	126,043			125,977

*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 and the adaptation variable is included. 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 second specification is closest to the baseline model (6) that is typically estimated in the literature. Here, the estimated coefficient on KDD is  $-0.0063$ , which implies an additional degree-day of heat above  $29^{\circ}\text{C}$  leads to a decrease in overall corn yield of 0.63 percent. The estimates of sensitivity in the first and third specifications are not significantly different, and do not change meaningfully when the regression is not weighted by harvest area or a time trend replaces the fixed time effect. These results are similar to those reported in Burke and Emerick (2016).

The fourth specification provides strong evidence of adaptation in the data. The added

regressor that captures adaptation is positive and highly significant, indicating that the level of adaptation is positively correlated with the level of KDD in a county and time period. This implies that as KDD increases the negative marginal effect of KDD on crop yield gets smaller in magnitude (i.e. closer to zero).<sup>7</sup>

According to model (4), the average marginal effect of KDD, after weighting by the amount of corn grown in each county, is -0.0086. This is significantly more negative than the average marginal effect in specifications (1)-(3), providing evidence of bias in the standard FE-OLS models. Note however, that while model (4) implies the average negative marginal effect of KDD on corn yield is 37 percent greater than the more conventional model (2), it also implies this effect is diminishing as KDD increases. The results of model (4) can be used to forecast future crop yield given climate change, taking into account the adaptation of farming practices to higher temperatures.

### *B. Heterogeneous Slope Models*

In this section we estimate a model that allows for heterogeneity in the slope coefficients on temperature and precipitation along the two dimensions of county and time. This provides a complete overview of the extent of adaptation both between counties and over the sample period. In contrast to the pooled FE-OLS model in equation (15), a heterogeneous slope model allows us to model farm adaptation behaviour without imposing a particular functional form a priori. We can use the estimates to investigate the nature of the relationship between the slope heterogeneity and KDD, and this will be informative about the nature of adaptation. We use the Mean-Observation OLS (MO-OLS) estimator to model both dimensions of heterogeneity. Table 2 presents the results for corn.

The table presents the unweighted and weighted means of the distribution of slope coefficients for each variable, as well as other features of the distribution. The mean coefficient for KDD is significantly more negative than in models (1)-(3) in Table 1, reflecting the heterogeneity bias that we outlined in Section II.B. The unweighted mean is -0.0096, implying one extra degree day of temperatures over 29°C will lead to a 0.96 percent reduction in crop yield. Weighting by the harvested area of corn in each county reduces this slightly

<sup>7</sup>Note that the coefficient on KDD itself becomes significantly more negative at -0.0159 when the nonlinear term is added. Meanwhile, the results for the other variables do not change meaningfully from the first three specifications.

TABLE 2—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 <sup>2</sup> (x10,000)	-0.0001 (0.0000)	-0.0001	-0.0001	0.0002	-0.0003	0.0001
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.

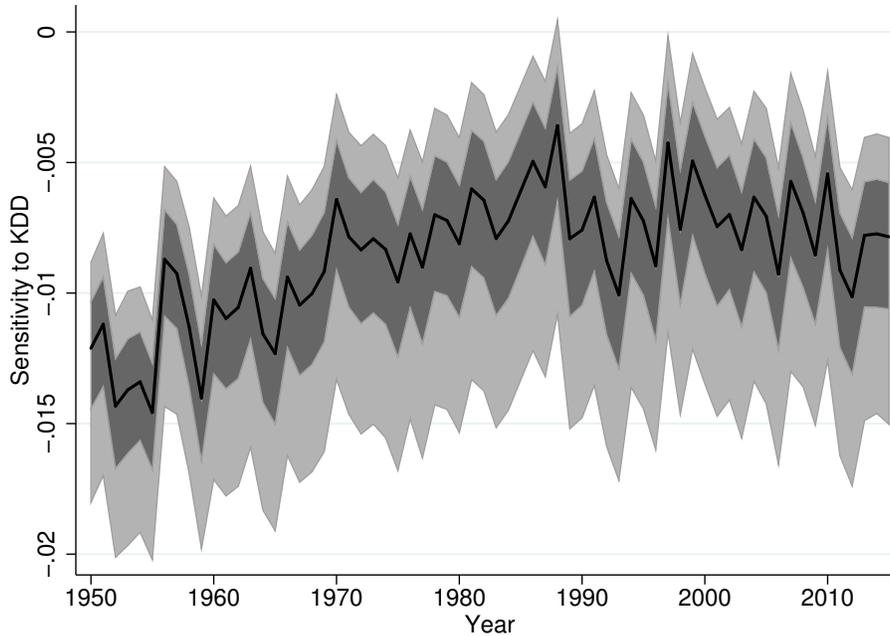
to -0.0089. This is strikingly similar to the average marginal effect of -0.0086 in Table 1 column (4) where we modelled adaptation based on the theory in Section I. It represents a 40% increase in sensitivity compared to the more conventional model (2) in Table 1.

We find substantial heterogeneity in the model coefficients. The standard deviation of the slope for KDD is 0.0068, with the 90/10 percentile range being -0.0034 to -0.0161.<sup>8</sup> Figure 2 plots the KDD coefficients between counties and over time. Clearly, there is significant heterogeneity between counties. The 90/10 percentile range of KDD coefficients is over .010 units in each year. Note that the lower light grey area, representing the 25/10 percentile range, is much wider than the 90/75 percentile range, so it is clear that the distribution of KDD coefficients has a fat left tail.

Adaptation follows a clear trend over time. The median KDD coefficient increases from 1950 until the late 1980s, and then stagnates. Table 3 presents a regression of the median coefficient for each year on a linear time trend, as well as tests for a structural trend break, following the method of Andrews (1993). The results indicate that a significant break occurred in 1989, coinciding with an extreme drought in the Midwest in 1988-89. This supports the findings of Lobell et al. (2014), who used field-level data for corn and soybeans and found that while average yields increased between 1995 and 2010, the sensitivity of

<sup>8</sup>It is reassuring that the 90th percentile coefficient is well below zero in almost all years.

FIGURE 2. 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 2, 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.*

crops to droughts or high heat increased because of the agronomic trend toward higher sowing densities. While this may lead to higher average yield, it can accentuate problems that arise when crops experience water stress.

Table 3 also presents a regression that accounts for the structural break by including a post-1988 dummy and its interaction with a time trend. Compared to the simple linear trend regression, this trend-break regression implies the time trend was twice as large up until the structural break, after which it turns slightly 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 main prediction of the theoretical model of Section I is that sensitivity to KDD will decrease as KDD increases, giving farmers more incentive to adopt adaptation techniques. In fact, the correlation between  $\hat{\beta}_{2it}$  and  $KDD_{it}$  is 0.43, supporting this prediction. Furthermore, as we showed in Section II.B, such a positive correlation leads to the upward

TABLE 3—ANALYZING THE MEDIAN COEFFICIENTS TO KDD FOR CORN

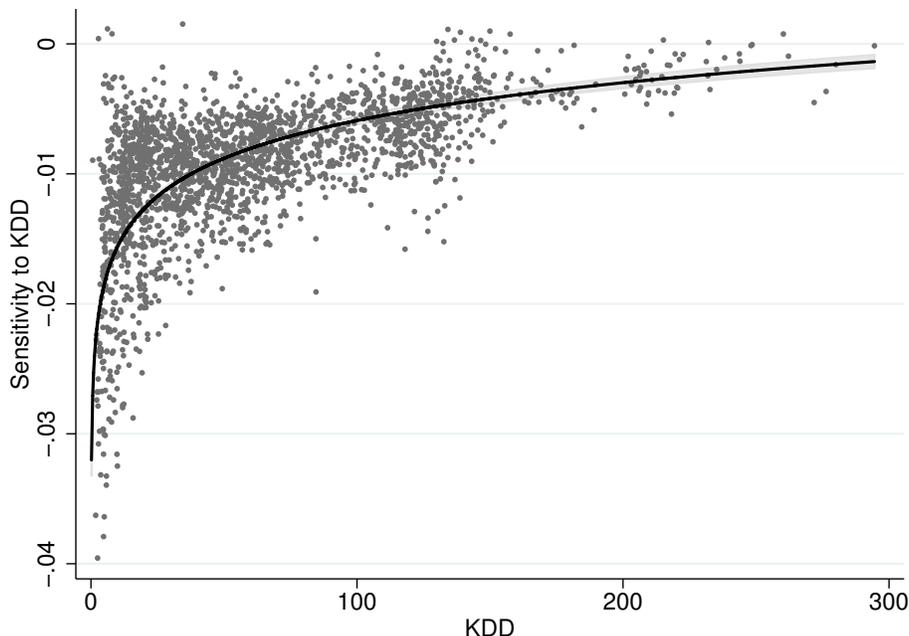
Regression Results	$\beta$	Std. Err.	$\beta$	Std. Err.
$t$	0.0001	0.0000	0.0002	0.0000
Constant	-0.0123	0.0005	-0.0142	0.0006
$d_{t>break}$			0.0067	0.0020
$t^*d_{t>break}$			-0.0002	0.0000
Structural Break Test	Statistic	p-value		
Supremum Wald	42.16	0.00		
Average Wald	27.57	0.00		
Supremum LR	34.24	0.00		
Average LR	24.00	0.00		

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

bias in the KDD coefficient, which is what we observe in pooled regressions that ignore adaptation (see Table 1 columns (1)-(3)).

Figure 3 plots the average values of  $\hat{\beta}_{2it}$  and  $KDD_{it}$  for each county, as well as a non-linear regression of  $\hat{\beta}_{2it}$  on  $KDD_{it}$ . The graph reveals a clear positive relationship, and the best fitting line to approximate this relationship is log-linear, consistent with the prediction of the simple theoretical model. Thus, the FE-OLS and MO-OLS approaches to modelling adaptation both generate log-linear relationships between  $\hat{\beta}_{2it}$  and  $KDD_{it}$ .

It is interesting to compare the relationship between  $\hat{\beta}_{2it}$  and  $KDD_{it}$  as calculated from the marginal effect of  $KDD_{it}$  on crop yield in column (4) of Table 1, with that obtained from the fitted log-linear relationship in Figure 3. Figure 4 plots the two curves together, and confirms that the MO-OLS log-linear curve lies within a 95 percent confidence interval of the FE-OLS log-linear curve. The MO-OLS curve implies a slightly more negative marginal effect across all levels of  $KDD_{it}$ , but the difference is not significant.

FIGURE 3. RELATIONSHIP BETWEEN  $\hat{\beta}_{2it}$  AND  $KDD_{it}$  FOR U.S. CORN

Note: This graph is a scatter plot of the average coefficient on  $KDD$  in each county reported in Table 2 column (4) and the average measurement of the  $KDD$  variable for the same county. The fitted line was obtained from the regression  $\hat{\beta}_{2it} = \alpha_1 + \alpha_2 \ln(KDD_{it})$ . The estimates are  $\alpha_1 = -0.0183$  and  $\alpha_2 = 0.0025$  and the 95% confidence interval for the curve is shaded.

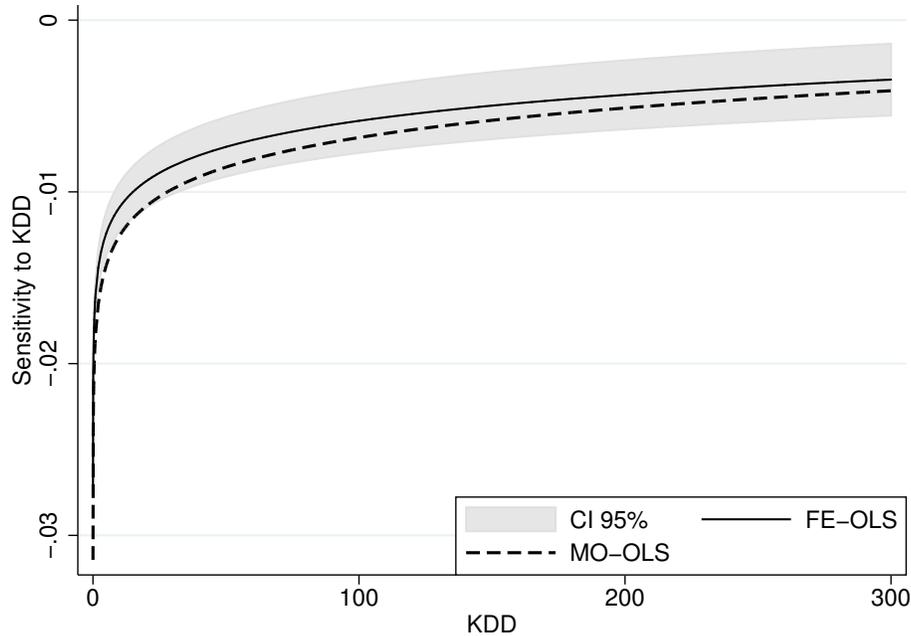
#### IV. Forecasting the Effect of Climate Change on Crop Yield

##### A. Approach to Forecasting

Here we forecast effects of climate change on U.S. corn yields over the 21st century, and examine the effectiveness of adaptation in mitigating these effects. We present annual forecasts from 2016 to 2100 using the FE-OLS and MO-OLS models with and without adaptation. We provide answers to several key questions: (i) the extent of damage that climate change may have on corn yield, (ii) the effectiveness of adaptation to mitigate this damage, (iii) the relative effectiveness of adaptation vs. emissions reductions as ways to avert damage, (iv) the distribution of losses between counties, and (v) the extent of heterogeneity in forecasts between climate models and/or econometric methods.

In order to forecast future corn yields, we need predictions of temperature (both KDD

FIGURE 4. COMPARISON OF LOG-LINEAR RELATIONSHIPS DERIVED UNDER MO-OLS AND FE-OLS FOR CORN



*Note: This graph compares the fitted log-linear relationships between  $\hat{\beta}_{2it}$  and  $KDD_{it}$  obtained from (i) the county/time specific coefficients estimated with MO-OLS (see Figure 3) vs. (ii) the FE-OLS regression results with adaptation, presented in Table 1 column 4.*

and GDD) and precipitation in each corn growing county in the U.S. through the end of the century. This requires predictions from a climate model, which further requires that we input a CO2 emission scenario. To gauge the robustness of our predictions to the scientific modelling assumptions, we use a large set of climate models and three emissions scenarios. To our knowledge no prior study has forecast crop yield using a suite of climate models and multiple emission scenarios. We now describe the procedure in detail:

We begin by taking the raw outputs of nineteen general circulation models (GCMs or simply ‘climate models’)<sup>9</sup> from the Coupled Model Intercomparison Project v5 (‘CMIP5’).<sup>10</sup> The core of every GCM is a series of equations that describe the behaviour of rotating

<sup>9</sup>The list of models are: ACCESS 1.0, BNU-ESM, CANESM2, CCSM4, CESM1(CAM5), CSIRO-Mk3.6.0, EC-EARTH, FGOALS-g2, FIO-ESM, GFDL-CM3, GFDL-ESM2G, GISS-E2-R, HadGEM2-ES, IPSL-CM5A-LR, IPSL-CM5A-MR, MIROC-ESM, MIROC-ESM-CHEM, MPI-ESM-LR, and NorESM1-M.

<sup>10</sup>The CMIP protocol was introduced by The World Climate Research Program as part of the Working Group on Coupled Modelling. It collects notable climate model outputs and ensures that their outputs are comparable on a consistent basis, allowing scientists to analyze the outputs of GCMs in a systematic fashion. The fifth version of CMIP is part of the broader effort for the IPCC Fifth Assessment Report.

spheres of continuous fluid that simulate behaviour of the Earth’s atmosphere and oceans. Climate models differ in how they represent a number of processes, such as chemical reactions, cloud formation, and vegetation growth, and, as we will see, they can generate rather different predictions of how the climate will respond to CO<sub>2</sub> emission growth.

Each GCM or climate model requires as an exogenous input the atmospheric concentration of greenhouse gases for each time period the model is run. In order to compare the output of climate models consistently, every model is run using four standard emissions scenarios known as representative concentration pathways (‘RCPs’). CMIP5 uses four RCPs: RCP26, RCP45, RCP60, and RCP85.<sup>11</sup> Each RCP contains assumptions regarding the future trajectory of population growth, technological development, and government policies, to produce predictions of global emissions.

We will henceforth refer to RCP85 as the “business as usual” scenario in the sense that very little action is taken by the global community to curb emissions. Under this scenario, emissions and consequently atmospheric concentrations of CO<sub>2</sub> continue to grow at present rates until the end of the forecast horizon.

We will refer to RCP45 as the “moderate” emissions reductions scenario. It corresponds to a case that is slightly more ambitious than current government pledges. Under this scenario, atmospheric CO<sub>2</sub> concentrations continue to climb until the end of the century, but the growth rate significantly declines after 2060.

Finally, the RCP26 scenario represents “substantial” emission reductions. It is consistent with the most ambitious targets set under the Paris agreement. Under this scenario, CO<sub>2</sub> concentrations peak in 2040 and then slowly decline until 2100.<sup>12</sup> This requires that actual emissions peak as early as 2020 and then begin to substantially decline, which would require significant and immediate action on curbing emissions by the global community.

The nineteen climate models that we utilize take the predictions of atmospheric CO<sub>2</sub> concentrations from the three RCP scenarios and use them to produce forecasts of daily min/max temperature, as well as precipitation, across a 12x12km grid of the contiguous

<sup>11</sup>The RCP scenario numbers correspond to the radiative forcing values in 2100. For instance, RCP26 results in a forcing value of +2.6W/m<sup>2</sup> above pre-industrial levels at 2100. We do not consider RCP60 in this article due to its similarity to RCP45.

<sup>12</sup>To be precise, atmospheric concentrations of CO<sub>2</sub> (and all other forcing agents converted to CO<sub>2</sub> equivalence) reach 1240 parts per million by 2100 under the “business as usual” scenario, 575 pp million under the “moderate” scenario, and 435 ppm under the “ambitious” (or “substantial” reduction) scenario.

U.S. Next, these grid level forecasts must be converted to county level forecasts. This interpolation procedure is known as “bias-correction and spatial disaggregation” (BCSD). We take the BCSD output for our nineteen climate models from the United States Geological Survey Geo Data Portal, who in turn rely on the Bureau of Reclamation (2013).

Finally, we convert the daily temperature and precipitation forecasts into growing season specific values of  $GDD_{itrm}$ ,  $KDD_{itrm}$ , and  $PREC_{itrm}$  for county  $i$ , year  $t$ , RCP scenario  $r$  and climate model  $m$ . This is done using the sinusoidal function in (18) and (19).

Each climate model in CMIP5 was run from 2006 to 2100, so county-specific forecast errors for 2006-15 are observable. We center the forecasts so the average county-specific forecast error of each model for 2006 to 2015 is zero. This corrects for level biases that a particular climate model may have for specific counties. This provides a more accurate picture of the future paths that would be optimally predicted by each model, as optimal predictions should take into account already observed errors. Thus, for KDD we have:

$$(20) \quad K\tilde{D}D_{itrm} = KDD_{itrm} + (K\bar{D}D_i - K\bar{D}D_{irm})$$

for  $t = 2016, 2017, \dots, 2100$  where  $K\bar{D}D_{irm} = 10^{-1} \sum_{t=2006}^{2015} KDD_{itrm}$  and  $K\bar{D}D_i = 10^{-1} \sum_{t=2006}^{2015} KDD_{it}$ . We use the same procedure for  $G\tilde{D}D_{itrm}$  and  $P\tilde{R}E\tilde{C}_{itrm}$ .

We are now in a position to use the econometric models in Section III, paired with each climate model and CO2 emissions scenario, to forecast future corn yields with and without adaptation. First, we forecast yields using the conventional FE-OLS model from Table 1 column (2) that does not account for adaptation, which we denote model 1. This gives:

$$(21) \quad \hat{y}_{1itrm} = 2.7188 + \hat{c}_i + 1.3385 + 0.0003(G\tilde{D}D_{itrm}) - 0.0063(K\tilde{D}D_{itrm}) + \\ 0.0010(P\tilde{R}E\tilde{C}_{itrm}) - 0.00000092(P\tilde{R}E\tilde{C}_{itrm}^2)$$

where all parameters can be read from Table 1, except for  $\hat{c}_i$ , which is the estimate of the county-specific fixed effect, and 1.3385, which is the 2006-15 mean of the time fixed effect.

Next, we use the FE-OLS specification in Table 1 column (4) that allows for adaptation:

$$(22) \quad \begin{aligned} \hat{y}_{2itrm} = & 2.1890 + \hat{c}_i + 1.3574 + 0.0004(G\tilde{D}D_{itrm}) - 0.0159(K\tilde{D}D_{itrm}) + \\ & 0.0022(\ln(K\tilde{D}D_{itrm}) * K\tilde{D}D_{itrm} - K\tilde{D}D_{itrm}) + 0.0007(P\tilde{R}E\tilde{C}_{itrm}) \\ & - 0.00000067(P\tilde{R}E\tilde{C}_{itrm}^2) \end{aligned}$$

where all parameters can be read from Table 1, except  $\hat{c}_i$  and 1.3574, which is the 2006-15 average of the time fixed effect. We call this model 2.

The MO-OLS model can provide yield forecasts with or without adaptation. First, it is useful to define  $\bar{\beta}_{ki}$  as the mean of  $\beta_{kit}$  over the 2016-15 period for  $k = 1, \dots, 4$ , and similarly for  $\bar{c}_t$ . Then, a yield forecast that does not allow for adaptation can be obtained as:

$$(23) \quad \hat{y}_{3itrm} = \bar{c}_t + \bar{\beta}_{1i}G\tilde{D}D_{itrm} + \bar{\beta}_{2i}K\tilde{D}D_{itrm} + \bar{\beta}_{3i}P\tilde{R}E\tilde{C}_{itrm} + \bar{\beta}_{4i}P\tilde{R}E\tilde{C}_{itrm}^2$$

Here the marginal effect of each variable is fixed at the 2006-2015 mean for each county, so there is no adaptation. We call this model 3.

Finally, the MO-OLS model can allow for adaptation to high temperatures by setting the slope coefficient on the  $K\tilde{D}D_{itrm}$  variable equal to the log-linear fitted relationship shown in Figure 3. That is, we set  $\hat{\beta}_{2it} = \alpha_1 + \alpha_2 \ln(KDD_{it})$  where  $\hat{\alpha}_1 = -0.0183$  and  $\hat{\alpha}_2 = 0.0025$ . We refer to this as model 4:

$$(24) \quad \begin{aligned} \hat{y}_{4itrm} = & \hat{c}_i + \bar{c}_t + \bar{\beta}_{1i}G\tilde{D}D_{itrm} + (\log(K\tilde{D}D_{itrm}) * 0.0025 - 0.0183)K\tilde{D}D_{itrm} \\ & + \bar{\beta}_{3i}P\tilde{R}E\tilde{C}_{itrm} + \bar{\beta}_{4i}P\tilde{R}E\tilde{C}_{itrm}^2 \end{aligned}$$

There is an asymmetry in (24) in that we allow the KDD coefficient to adapt over time in response to high temperatures, but we do not allow other aspects of technology to adapt.<sup>13</sup> This is because we find no evidence that other parameters adapt to high temperatures in our sample. For instance, the correlations of  $\hat{c}_{it}$  and  $\hat{\beta}_{1it}$  with  $KDD_{it}$

<sup>13</sup>Schlenker, Roberts and Lobell (2013) hypothesize, in a comment to Butler and Huybers (2013), that a model where time effects and the  $GDD_{it}$  coefficient also change as a function  $KDD_{it}$  would be more realistic, as adaptation techniques may lead to a decrease in overall farming productivity (a phenomenon known as 'mal-adaptation').

are only 0.02 and -0.04, respectively. This is in sharp contrast to the strong relationship between  $\hat{\beta}_{2it}$  and  $KDD_{it}$  depicted in Figure 3. It is crucial to bear in mind the distinction between adaptation per se (which we can forecast from historical data) and more general forms of technical progress. Accounting for changes in  $\hat{c}_t$  and  $\beta_{1it}$  that may occur due to general technical progress is a more speculative exercise that we take up in Appendix C.

Our county-level yield forecasts are aggregated to the national level as follows:

$$(25) \quad \hat{y}_{trm} = \left( \sum_{n=1}^N w_i \right)^{-1} \left( \sum_{n=1}^N \hat{y}_{itrn} w_i \right)$$

where  $w_i$  is the average crop area of county  $i$ . For each econometric model and emissions scenario  $r$ , we will report the mean forecast across all climate models,  $\hat{y}_{tr}$ . We also report prediction intervals around this mean derived from the standard deviation of  $\hat{y}_{trm}$  across models  $m$ . This allow us to assess the extent of disagreement between climate models.<sup>14</sup>

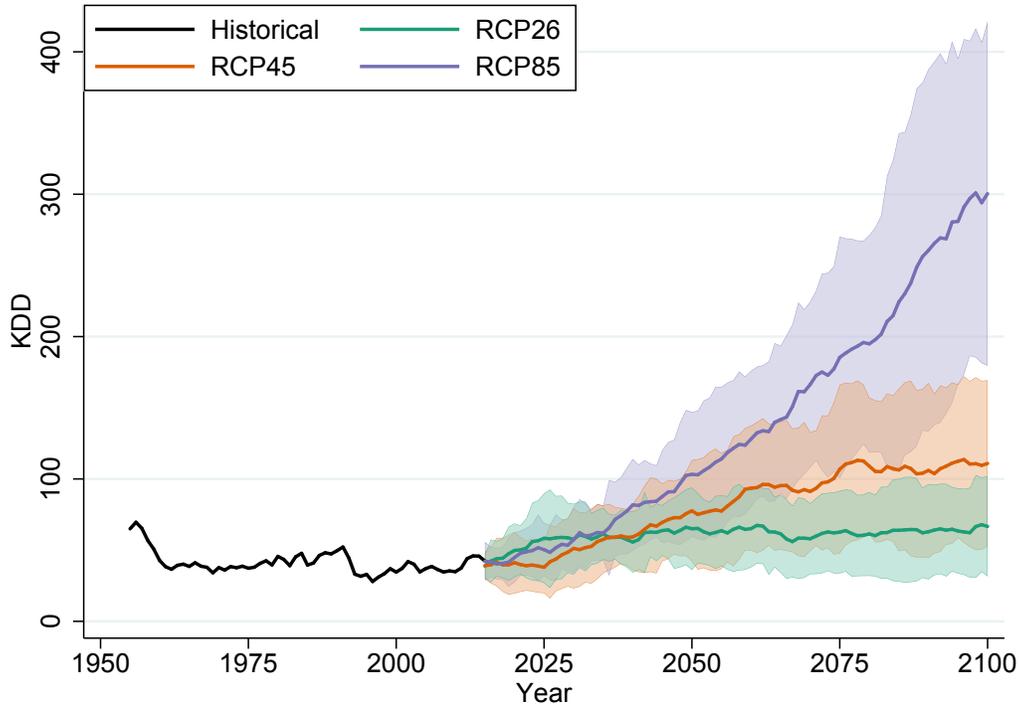
We present the predictions as percentage changes relative to the 2006-2015 historical average yield. Finally, we apply a five-period moving average to the ensemble average prediction and prediction intervals, simply to reduce noise so as to help visualize trends.

### B. Forecast Results

This section presents forecasts of corn yield from 2015 to 2100 using four econometric model, thee RCP scenarios and nineteen climate models from CMIP5. This gives  $4 \cdot 19 \cdot 3 = 228$  different forecast scenarios. Before presenting the yield forecasts, it is useful to examine the climate models' projections of both KDD and GDD in the corn-growing regions of the U.S., and their variability across climate models and RCP scenarios. Figure 5 presents the average forecast of KDD across the ensemble of models for three RCPs, where the solid line is the average forecast and the shaded areas are the 80% prediction intervals.

Over the historical period (1950-2015) it is clear that the moving average of KDD has moved little, except for a few abrupt shifts. In the RCP26 scenario (green) the ensemble average prediction is that KDD will increase slowly until 2050 and then plateau. The increase is sufficiently minor that the level at 2100 is not above the levels of the mid-

<sup>14</sup> Reporting results as an ensemble average of climate models and reporting the variability in impacts between models is an approach that is recommended in Auffhammer et al. (2013).

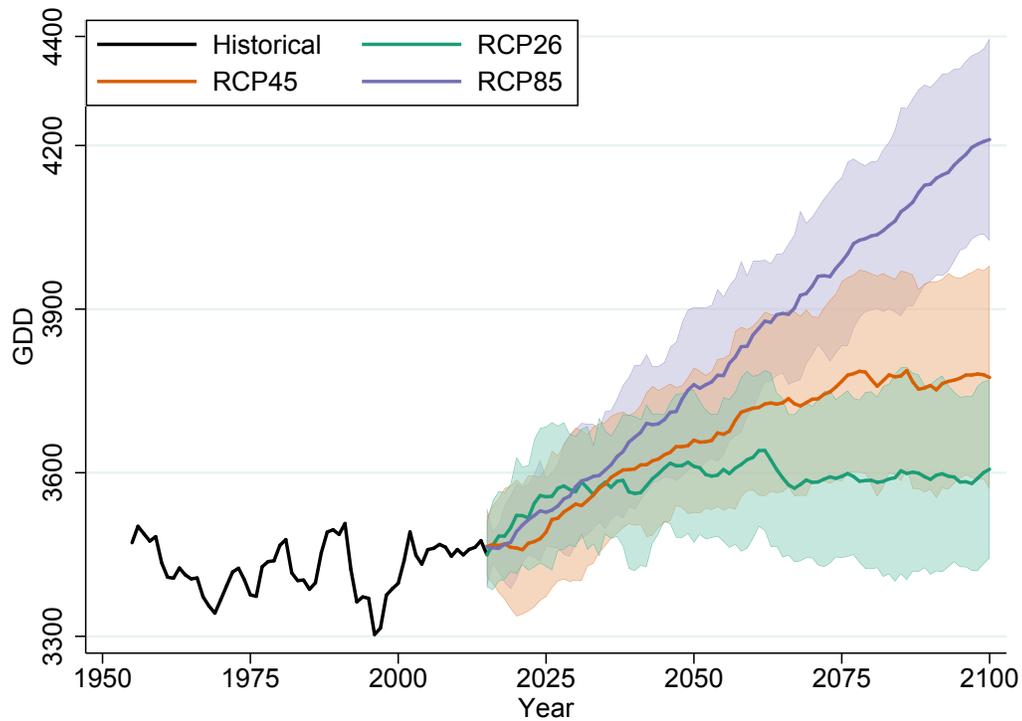
FIGURE 5. FORECAST  $KDD_{it}$  BY REPRESENTATIVE CONCENTRATION PATHWAY

*Note: This graph presents average forecasts of KDD (weighted by corn acreage) across three representative concentration pathways, where the solid line is the average forecast across 19 CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.*

1950s. Under the RCP45 scenario (orange), KDD grows at a similar pace until 2050, but then continues to grow until around 2080. In 2100 the average forecast implies a level of KDD more than double that we observed over the last five decades. But there is substantial disagreement between climate models, with some predicting much more modest changes.

Finally, in the RCP85 scenario (purple) KDD rapidly diverges from the other two scenarios after around 2040, and despite wide prediction intervals, every climate model predicts severe increases in KDD relative to historical experience.

Figure 6 presents the forecasts for GDD. The predictions for the three RCPs are broadly consistent with those for KDD, with RCP85 leading to the largest upward shift in the distribution of temperature. Thus, the climate models predict that both KDD and GDD will increase in the absence of emissions reductions. As a result, cooler corn-growing regions

FIGURE 6. FORECAST  $GDD_{it}$  BY REPRESENTATIVE CONCENTRATION PATHWAY

*Note: This graph presents average forecasts of GDD (weighted by corn acreage) across three representative concentration pathways, where the solid line is the ensemble average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.*

will experience less temperatures close to freezing over the growing season. Accordingly, the effect of climate change on yields will be the net balance between the negative effect of additional KDD and the positive effect of additional GDD.

Table 4 summarizes our main results for corn yield. We report forecast results for the three RCP scenarios at four points in time: 2030, 2050, 2080, and 2100. The effect of climate change on corn yield is expressed as a percentage change relative to the 2006-2015 average. The table also separates forecasts by the four econometric methodologies, which are FE-OLS with and without adaptation, and MO-OLS with and without adaptation. We present both the ensemble average forecast and the 80% prediction interval (in brackets).

As we see in Table 4, we predict extreme reductions in corn yield in a scenario without emissions reductions (RCP85) or adaptation. The model ensemble average reduction is

TABLE 4—THE EFFECTS OF CLIMATE CHANGE ON CORN YIELD (PCT)

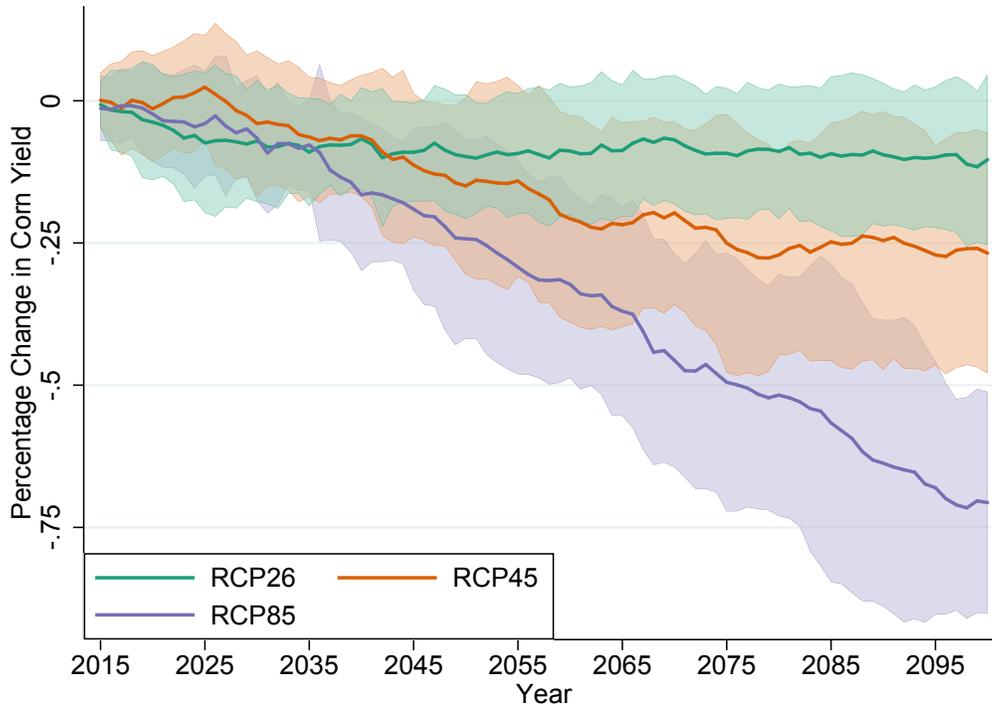
Year	No Future Adaptation		Future Adaptation	
	FE-OLS	MO-OLS	FE-OLS	MO-OLS
<u>RCP85</u>				
2030	-08 (-18, 02)	-07 (-18, 04)	-04 (-16, 07)	-06 (-14, 02)
2050	-25 (-41, -09)	-24 (-42, -07)	-21 (-36, -05)	-15 (-25, -05)
2080	-51 (-70, -32)	-52 (-73, -30)	-42 (-59, -26)	-27 (-38, -16)
2100	-69 (-88, -50)	-71 (-90, -51)	-57 (-73, -42)	-37 (-48, -26)
<u>RCP45</u>				
2030	-05 (-17, 07)	-04 (-17, 09)	-02 (-15, 11)	-04 (-13, 04)
2050	-16 (-30, -02)	-15 (-31, 01)	-12 (-26, 02)	-11 (-20, -01)
2080	-27 (-44, -10)	-27 (-45, -09)	-23 (-38, -07)	-17 (-26, -07)
2100	-27 (-47, -08)	-27 (-48, -06)	-22 (-41, -04)	-16 (-28, -05)
<u>RCP26</u>				
2030	-09 (-19, 02)	-07 (-18, 04)	-06 (-16, 05)	-07 (-14, 00)
2050	-11 (-22, 00)	-10 (-21, 01)	-08 (-19, 03)	-08 (-14, -02)
2080	-10 (-22, 01)	-09 (-21, 03)	-07 (-19, 06)	-08 (-15, 00)
2100	-12 (-26, 02)	-10 (-25, 05)	-08 (-23, 06)	-08 (-18, 01)

*Notes:* Results are expressed in terms of percentage change from the 2006-2015 historical weighted average crop yield. Each number represents the ensemble average over nineteen climate models, while the numbers in brackets represents the 80 percent prediction interval.

about 70% by 2100. Notice the predictions based on FE-OLS vs. MO-OLS are almost identical. Even with “moderate” emissions reductions (the RCP45 scenario) we still observe significant reductions in yield, with an ensemble average loss of 27% by 2100. But with significant immediate action on emissions (as in RCP26) we see only a 10% model average reduction in yield (using MO-OLS), and the 80% prediction interval indicates that some climate models even predict small gains.

Figure 7 charts forecasts of corn yield under the three RCP scenarios using the MO-OLS estimator with no future adaptation. All three model ensemble average scenarios show modest declines in corn yield between now and 2035-40, but then they start to diverge appreciably. By 2100, RCP26 shows only a mild decline in yield. RCP45 shows a larger mean decline, but with large disagreement between climate models (e.g., the prediction interval overlaps greatly with RCP26). The decline under RCP85 is so severe that by the

FIGURE 7. THE EFFECT OF CLIMATE CHANGE ON CORN YIELD BY RCP (MO-OLS WITH NO FUTURE ADAPTATION)



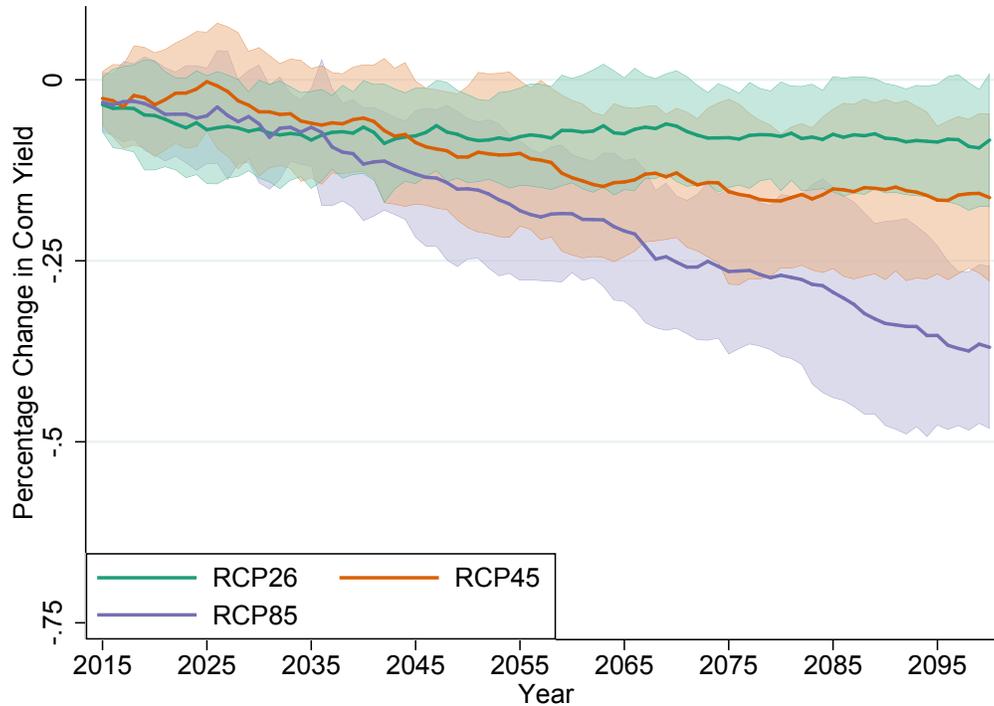
*Note: This graph presents forecasts of the percentage change in corn yield (relative to the 2006-2015 historical average) across three RCPs, where the solid line is the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.*

end of the century the 80% prediction intervals for RCP85 and RCP45 do not overlap.

The wide prediction intervals in Figure 7 reveal that climate models differ greatly in their predictions of future conditions in the corn-growing states, so it is clearly important to use a large ensemble of GCMs to predict future yields. In contrast, in the case of no adaptation, the FE-OLS and MO-OLS econometric models generate very similar forecasts. This may seem surprising given our earlier finding that FE-OLS with no adaptation gives upward biased estimates of the KDD coefficient. However, FE-OLS also gives a smaller estimate of the GDD coefficient. GDD and KDD both increase over time, so these differences between the two models tend to cancel out, giving similar forecasts.

We now turn to forecasts that incorporate predicted future adaptation, which our models infer from the historical adaptation patterns documented in Figure 4. Results are reported

FIGURE 8. THE EFFECT OF CLIMATE CHANGE ON CORN YIELD BY RCP (MO-OLS WITH FUTURE ADAPTATION)



*Note: This graph presents forecasts of the percentage change in corn yield (relative to the 2006-2015 historical average) across three RCPs, where the solid line is the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80% prediction intervals.*

in Table 4. Adaptation does appreciably mitigate the damage from climate change in the RCP85 scenario. For example, using the FE-OLS model with adaptation, the average decline in corn yield in 2100 drops from 69% to 57%. The MO-OLS approach implies that adaptation is even more effective at mitigating damage. In 2100 the reduction in crop yield decreases from 71% to “only” 37%.

Of course, with greater emissions reductions KDD increases are smaller, and the scope for adaptation is reduced. For example, given the RCP45 scenario and the MO-OLS model, the mean predicted drop in yield in 2100 is 27% without and 16% with adaptation. And under RCP26 the analogous figures are 10% and 8%, so adaptation is almost irrelevant.

A key takeaway from Table 4 is that adaptation may substantially mitigate damage in the “business as usual” scenario, but damage remains severe. In contrast, even moderate

emission reductions are quite effective at reducing damage - e.g., for MO-OLS, compare the 37% mean drop in yield under RCP85 with adaptation with the 27% drop in yield achieved under RCP45 even without adaptation (falling to 16% when adaptation is included).

Figure 8 charts forecasts of corn yield under the three RCP scenarios using the MO-OLS estimator with adaptation. Comparison with Figure 7 reveals how adaptation reigns in the most negative forecasts and brings the distribution of outcomes under the three RCP scenarios closer together. Adaptation also reduces the disagreement between climate models (which is also clear from Table 4). However, the extent of disagreement remains large. While the mean prediction of yield damage by 2100 under MO-OLS with adaptation and given current government emissions pledges (RCP45) may be viewed as fairly modest at 16%, the 80% prediction interval still includes a much more severe 28% decline.

As we noted earlier, the MO-OLS estimates imply a greater scope for adaptation to mitigate damage to yields than the FE-OLS estimates. This is likely due to differences in how the effect of  $KDD_{it}$  on  $y_{it}$  changes as  $KDD_{it}$  changes. Under the MO-OLS approach in (24), the average effect of  $KDD_{it}$  is always equal to the marginal effect, so adaptation shifts both equally. But in the FE-OLS approach to modelling adaptation in (22), the average effect of  $KDD_{it}$  on  $y_{it}$  is given by  $\hat{\beta}_2 + \hat{\beta}_3(\ln(KDD_{it}) - 1)$ , which is always less than the marginal effect given in (16). Thus, adaptation in the FE-OLS approach leads to a much slower decrease in the average effect relative to the MO-OLS approach. The MO-OLS approach is perhaps more intuitive here, as it implies that adaptation techniques alter the impact of all units of  $KDD_{it}$  on the crop, not simply the additional units.

### C. Cumulative Losses from 2020-2100

Next we examine the effectiveness of adaptation and alternative emissions reduction scenarios as ways to avert damage from climate change over the whole forecast horizon. We first calculate the total damage to future crop yield as:

$$(26) \quad Loss_{r,ad} = M^{-1} \sum_{m=1}^M \sum_{t=2020}^{2100} (\hat{y}_{trm,ad} - \bar{y})$$

where  $\bar{y}$  is historical average crop yield from 2006 to 2015,  $r$  is the RCP scenario and  $ad$

is the adaptation scenario. The proportion of damage averted under emissions and adaptation scenario ( $RCPn, ad = k$ ) for  $n = 26, 45, 85$  and  $k = 0, 1$  is then  $1 - \frac{Loss_{RCPn, ad=k}}{Loss_{RCP85, ad=0}}$ , where ( $RCP85, ad = 0$ ) is the baseline of “business as usual” and no adaptation.<sup>15</sup>

TABLE 5—PROPORTION OF CLIMATE CHANGE DAMAGE AVERTED (PCT)

Estimator	RCP85 + Adapt	RCP45	RCP45 + Adapt.	RCP26
MO-OLS	33 (21, 45)	55 (28, 82)	57 (36, 78)	67 (43, 91)
FE-OLS	18 (10, 25)	52 (26, 78)	62 (33, 90)	64 (42, 87)

Table 5 presents the proportion of damage averted under four scenarios. Recall that RCP85 is “business as usual”, RCP45 entails emissions reductions slightly greater than current government pledges, and RCP26 is consistent with the most ambitious targets under the Paris agreement. According to Table 5 column (3), the MO-OLS model predicts that shifting from RCP85 to RCP45 without any adaptation will avert 55% of damage on average, with the 80% prediction interval ranging from 28% to 82%. If we also account for adaptation the mean damage reduction improves slightly to 57%, and the 80% prediction interval narrows slightly to 36% to 78%. Thus, our point forecast is that current policies combined with adaptation will avert more than half of climate change induced damage to corn yields, but the extent of disagreement across climate models is substantial. Indeed, if we look at the more ambitious RCP26 scenario, the mean damage abatement is 67%, but the 80% prediction interval ranges from only 43% to a massive 91%. Note that FE-OLS gives similar point forecasts and similarly large prediction intervals. Thus, most of the forecast uncertainty is across climate models, not econometric methods.

According to Table 5 column (2), MO-OLS predicts that adaptation alone can avert 33% of total damage to yield (on average),<sup>16</sup> while FE-OLS gives a much smaller mean damage abatement of only 18%. This is a substantial difference, yet both econometric

<sup>15</sup>We could of course discount future losses in (26), but the proper way to discount losses of future generations is highly controversial. Many have argued against discounting on ethical grounds. In our case, discounting scales down losses, but has almost no effect on relative losses across scenarios, which is what we focus on.

<sup>16</sup>As the MO-OLS model predicts a mean yield reduction of 71% in 2100 without adaptation and 37% with adaptation (see Table 4), one might conclude that adaptation averts about half of the damage from climate change. However, the results in Table 5 are based on all years from 2020 to 2100, not merely 2100.

approaches agree that adaptation alone cannot avert the majority of the severe damage to yield that is forecast to occur with climate change. Emissions reductions are necessary to avert substantial reductions in future yields.

It is also clear from our results that it is virtually impossible to avert all of the damage from climate change. Even under the most optimistic RCP26 scenario with adaptation, both models predict about a 30% reduction in yields (on average). Unfortunately, past emissions growth will lead to increases in atmospheric CO<sub>2</sub> concentrations irrespective of ambitious efforts to reduce emissions growth in the near future.

In Appendix C we present forecasts of corn yield that also include forecasts of future technological growth. These results suggest that the positive influence of technical progress may counter much of the adverse effect of climate change, but only if emission reductions are undertaken (at least to the extent of RCP45).

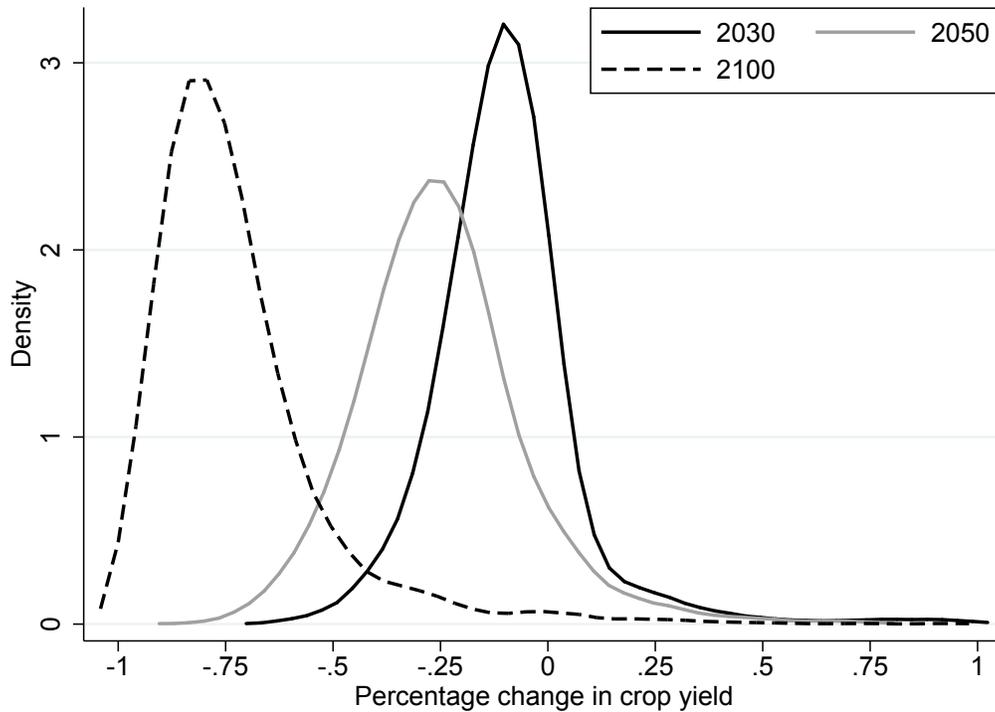
Finally, we present our forecast results for soybeans in Appendix D. The results suggest that soybean yield is also likely to decline significantly as a result of climate change. Our models are less optimistic about the scope of adaptation to mitigate this damage.

#### *D. Distribution of Losses Across U.S. Counties*

Here we examine the distribution of losses across counties due to future climate change. Figure 9 plots kernel densities of the percentage change in corn yield (relative to the 2006-2015 historical average) across all corn-growing counties at three points in time. All three plots are based on the MO-OLS forecasts without adaptation, and are the ensemble average of the nineteen climate models. In 2030 the mean change in yield is negative, yet a fair number of counties do receive benefits climate change. Severe losses, say over 30 percent, are not very common. But in 2050 about experience losses of over 30 percent, and few counties have gains. By 2100 almost all counties suffer severe drops in yield, and almost no counties experience gains. While not definitive, these results cast doubt on the notion that corn production can be shifted to colder corn-growing counties as a way to avert significant damage to yields.

Figure 10 graphs the percentage of counties that experience losses as a function of time, RCP, and adaptation scenario. Consistent with Figure 9, the fraction that experience

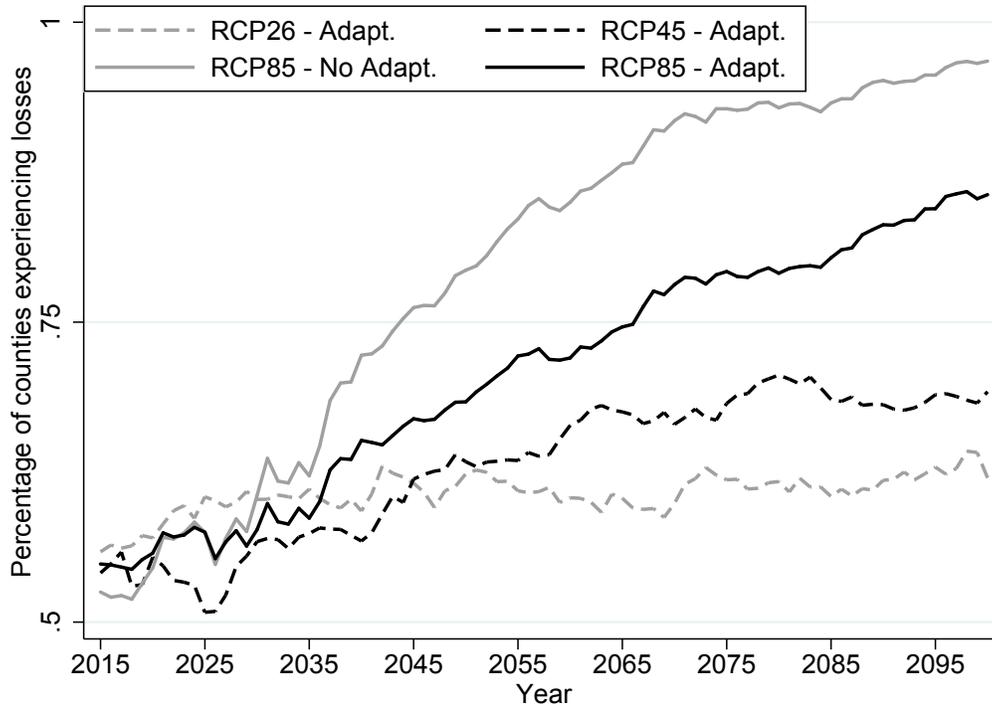
FIGURE 9. THE DISTRIBUTION OF LOSSES OVER COUNTIES BY YEAR



*Note: This graph presents the distribution of losses and gains for each county (defined as the percentage change in corn yield relative to the 2006-2015 historical average) over three time periods. Estimates use the MO-OLS econometric model with no future adaptation, and are the ensemble average across nineteen climate models.*

losses approaches one under RCP85 without adaptation by the end of the century. Allowing for adaptation reduces this fraction, but not substantially, as over 80 percent of counties still experience losses by 2100. Under more ambitious RCP scenarios we still find that the proportion of counties that suffer losses increases over time, but at a much slower rate. In the best case scenario (RCP26 with adaptation) roughly 60% of counties experience losses from about 2040 onward, with about 40% experiencing gains. This suggests that shifting production to colder counties may be a useful strategy if combined with substantial emissions reductions and adaptation.

FIGURE 10. PROPORTION OF CORN-GROWING COUNTIES EXPERIENCING LOSSES FROM CLIMATE CHANGE



*Note: This graph presents the percentage of corn-growing counties that experiencing losses from climate change under four combinations of RCP emissions scenario and the MO-OLS model: i) RCP26 with future adaptation, ii) RCP45 with future adaptation, iii) RCP85 without future adaptation, and iv) RCP85 with future adaptation.*

## V. Conclusion

We make two contributions to the literature on forecasting the effect of climate change on crop yield. First, we propose two new methods to directly model adaptation in the agricultural production function. The first method is motivated by a simple theory that predicts the marginal effect of high temperatures should decrease as harsh temperatures become more common. The theory generates a simple function form that is not rejected by the data. The second method directly estimates two dimensions of fixed-effects slope heterogeneity in the yield equation, using the newly developed “mean observation OLS” estimator due to Neal (2016). This is the first panel data estimator to allow for two-dimensional fixed effects in slopes. Ex post we find significant heterogeneity in the tem-

perature effect. It tends to decrease as harsh temperatures are more common, again consistent with the simple theory. Both methods generate strong evidence for adaptation across hot vs. cold counties, but no evidence to suggest that the average level of adaptation has increased since 1989. We also find that estimators that do not model adaptation produce biased estimates of the effects of harsh temperatures.

Our second contribution is to apply these econometric models to generate comprehensive county-level forecasts of corn and soybean yields from 2016 to 2100 using three standard emission growth pathways and nineteen climate models. We are not aware of a previous forecasting effort that encompasses such a wide range of econometric methods, climate models and emission scenarios.

Our results imply several major conclusions regarding the impact of climate change on corn yield: First, absent emissions reductions or adaptation, we predict very severe effects on yield, with an average prediction across climate models of -70% by 2100. Second, adaptation methods are predicted to avert roughly 18 percent to 33 percent of that damage, depending on the econometric method. Thus adaptation is important, but at best will only avert a fraction of the damage to yields. Third, we predict that moderate emissions reductions similar to current government pledges can be expected to avert roughly 60 percent of the damage to yields, rising to roughly 70 percent under the more ambitious Paris targets. These results imply that plausible emissions reductions can avert a large fraction of the damage from climate change. Fourth, however, even under the most optimistic scenario (e.g., Paris accords plus adaptation) U.S. corn yields are expected to drop by about 30% (averaging across climate models). Thus, given past CO<sub>2</sub> emissions, it may well be too late to avert some serious damage to yields even given immediate action.

We also attempted to forecast future technical progress based on past trends (admittedly a rather speculative exercise). We predict that technical progress and adaptation alone (absent emissions reductions) will lead to yield increases that lag far behind the population growth rate. But, an optimistic prediction of technical change, combined with moderate to substantial emissions reductions and adaptation can, together, achieve yield growth roughly in line with population growth according the mean climate model forecast. Still, these figures deteriorate quickly under slightly less optimistic technology projections.

A striking feature of our results is the wide variability of forecasts across climate models. Indeed, we have focussed on mean predictions in this conclusion to avoid drowning the reader in a morass of prediction intervals (all of which are presented in detail in the text). Suffice it to say that even our more optimistic emissions/technology/adaptation scenarios put non-negligible mass on rather adverse outcomes. Furthermore, our forecasts for the second largest U.S. crop, soybeans, are generally a bit more pessimistic. So it is fair to say that climate change poses a substantial *risk* to U.S. agricultural yields, even under the more benign scenarios where our point forecasts of yield losses are moderate.

#### REFERENCES

- Andrews, D.** 1993. “Tests for Parameter Instability and Structural Change with Unknown Change Point.” *Econometrica*, 61(4): 821–856.
- Auffhammer, M., S. Hsiang, W. Schlenker, and A. Sobel.** 2013. “Using Weather Data and Climate Model Output in Economic Analyses of Climate Change.” *Review of Environmental Economics and Policy*, 7(2): 181–198.
- 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.
- Dell, M., B. Jones, and B. Olken.** 2012. “Temperature Shocks and Economic Growth: Evidence from the Last Half Century.” *American Economic Journal: Macroeconomics*, 4(3): 66–95.

- 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.
- Deschenes, O., and M. Greenstone.** 2011. “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US.” *American Economic Journal: Applied Economics*, 3(4): 152–185.
- 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., G. Hammer, G. McLean, C. Messina, M. Roberts, and W. Schlenker.** 2013. “The critical role of extreme heat for maize production in the United States.” *Nature climate change*, 3: 497–501.
- 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.
- Lobell, D., M. Roberts, W. Schlenker, N. Braun, B. Little, R. Rejesus, and G. Hammer.** 2014. “Greater Sensitivity to Drought Accompanies Maize Yield Increase in the U.S. Midwest.” *Science*, 344(6183): 516–519.
- 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.** 2018. “Multidimensional Slope Heterogeneity in Panel Data Models.” *UNSW Research Working Paper 2016-15A*. Available online: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2853992](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2853992).
- Reclamation.** 2013. “Downscaled CMIP3 and CMIP5 Climate Projections: Release of Downscaled CMIP5 Climate Projections, Comparison with Preceding Information, and

Summary of User Needs.” U.S. Department of the Interior, Bureau of Reclamation, Technical Service Center, Denver, Colorado, 116 p., available at: [http://gdo-dcp.ucllnl.org/downscaled\\_cmip\\_projections/techmemo/downscaled\\_climate.pdf](http://gdo-dcp.ucllnl.org/downscaled_cmip_projections/techmemo/downscaled_climate.pdf).

**Roberts, M., and W. Schlenker.** 2010. “Is Agricultural Production Becoming More or Less Sensitive to Extreme Heat? Evidence from U.S. Corn and Soybean Yields.” *NBER Working Paper*, , (16308).

**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.

**Schlenker, W., M. Roberts, and D. Lobell.** 2013. “US maize adaptability.” *Nature Climate Change*, 3: 690–691.

**United Nations Population Division.** 2017. “2017 Revision of World Population Prospects.” Can be found online at <https://esa.un.org/unpd/wpp/>.

#### APPENDIX A: AGGREGATION OF THE FARM PRODUCTION FUNCTION

Here we verify that the main predictions of farm-level model in Section I aggregate to the collection of farmers within each county. We directly observe crop yield for counties:

$$(A1) \quad y_{ct} = \frac{\sum_{i=1}^N Y_{it}}{\sum_{i=1}^N C_{it}}$$

for county  $c$  at time  $t$ , where  $Y_{it}$  and  $C_{it}$  are output and acreage of farm  $i$ , respectively. Thus, the county-level sensitivity to  $KDD_{it}$  is an average of the farm-level sensitivities weighted by the size of the farm in terms of land and other inputs:

$$(A2) \quad \beta_{2ct}^* = \frac{\sum_{i=1}^N C_{it} I_{it}^\delta \beta_{2it}^*}{\sum_{i=1}^N C_{it}} = \frac{\sum_{i=1}^N C_{it} I_{it}^\delta \left( \frac{-p C_{it} A_t I_{it}^\delta s KDD_{ct}}{\gamma} \right)^{-1/2}}{\sum_{i=1}^N C_{it}}$$

Assume that  $KDD_{it}$  is common to farms within counties, giving  $KDD_{ct}$ . Since  $s < 0$ ,  $C_{it} > 0$ , and  $I_{it}^\delta > 0$  for all  $i$  and  $t$ , we can verify that the county sensitivity to harsh temperature  $\beta_{2ct}^*$  is strictly concave in  $KDD_{ct}$ , just as  $\beta_{2it}^*$  is strictly concave in  $KDD_{it}$ :

$$\frac{\partial \beta_{2ct}^*}{\partial KDD_{ct}} = \frac{\sum_{i=1}^N}{\sum_{i=1}^N C_{it}} \left( \frac{s^2 A_t C_{it} (I_{it}^\delta)^2 p}{2\gamma \left( \frac{-sp A_t C_{it} I_{it}^\delta KDD_{ct}}{\gamma} \right)^{3/2}} \right) > 0 \quad \forall KDD_{ct}$$

$$\frac{\partial^2 \beta_{2ct}^*}{\partial KDD_{ct}^2} = \frac{\sum_{i=1}^N}{\sum_{i=1}^N C_{it}} \left( \frac{3s^3 A_t^2 C_{it}^2 (I_{it}^\delta)^3 p^2}{4\gamma^2 \left( \frac{-sp A_t C_{it} I_{it}^\delta KDD_{ct}}{\gamma} \right)^{5/2}} \right) < 0 \quad \forall KDD_{ct}$$

We can also show that  $\beta_{2ct}^*$  is a log-linear function of  $KDD_{ct}$ . Thus, the two main predictions of the theoretical model carry over to the county-level (even if farm-level production functions cannot be aggregated into a ‘representative producer’).

#### APPENDIX B: A DESCRIPTION OF MO-OLS

In this appendix we describe the MO-OLS estimator due to Neal (2016), which is the first panel data estimator to allow for two-dimensional fixed effects in slopes (as well as intercepts). Consider the following model, with fixed effects in intercepts and slopes:

$$(B1) \quad y_{it} = \alpha_{it} + \beta'_{it} \mathbf{x}_{it} + e_{it}$$

where  $\mathbf{x}_{it} = (x_{1it}, \dots, x_{Kit})$  is a  $K \times 1$  vector of regressors,  $\beta_{it} = (\beta_{1it}, \dots, \beta_{Kit})$  is a  $K \times 1$  vector of coefficients that vary across individuals and over time, and  $e_{it}$  is an error term satisfying  $E[e_{it}|x_{it'}] = 0, \forall t, t', E[e_{it}e_{jt}] = 0, i \neq j$ . We assume each heterogeneous parameter consists of three components: (i) one that is common across all observations, (ii) an individual-specific fixed-effect, and (iii) a time-specific effect. Thus we have  $\alpha_{it} = \alpha + c_i + c_t$  and  $\beta_{it} = \beta + \lambda_i + \lambda_t$ , where  $\beta = (\beta_1, \dots, \beta_K)$ ,  $\lambda_i = (\lambda_{1i}, \dots, \lambda_{Ki})$ , and  $\lambda_t = (\lambda_{1t}, \dots, \lambda_{Kt})$ . Finally, define  $\theta = (\alpha, \beta)$  as the mean effects,  $\theta_i = (c_i, \lambda_i)$  as the individual effects,  $\theta_t = (c_t, \lambda_t)$  as the time effects, and  $\mathbf{z}_{it} = (1, \mathbf{x}_{it})$  as the regressors (including a constant, so fixed effects in levels are subsumed in the analysis).

We seek to obtain consistent estimates of the individual coefficients  $\theta_{it} = (\alpha_{it}, \beta_{it})$  as well as the average coefficient over the sample  $\bar{\theta} = (\alpha + E(c_i) + E(c_t), \beta + E(\lambda_i) + E(\lambda_t))$ .

MO-OLS does this by combining three types of regressions: pooled, person-specific, and time-specific. First, consider the pooled OLS estimator of  $\boldsymbol{\theta}$ . Equation (B1) can be written:

$$y_{it} = \mathbf{z}'_{it}\boldsymbol{\theta} + v_{it}$$

$$v_{it} = \mathbf{z}'_{it}\boldsymbol{\theta}_i + \mathbf{z}'_{it}\boldsymbol{\theta}_t + e_{it}$$

Then the pooled OLS estimator of  $\boldsymbol{\theta}$  is:

$$(B2) \quad \hat{\boldsymbol{\theta}} = \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it} \right)^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}y_{it} \right)$$

Expanding on  $y_{it}$  and simplifying yields:

$$(B3) \quad \begin{aligned} \hat{\boldsymbol{\theta}} = & \boldsymbol{\theta} + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it}\boldsymbol{\theta}_i \right) \\ & + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it}\boldsymbol{\theta}_t \right) + \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}e_{it} \right) \end{aligned}$$

where  $\mathbf{Q}_{zz,NT}^{-1} = \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it} \right)^{-1}$ .

Second, consider a set of regressions for each individual  $i$ . Equation (B1) can be written:

$$y_{it} = \mathbf{z}'_{it}(\boldsymbol{\theta} + \boldsymbol{\theta}_i) + v_{it}$$

$$v_{it} = \mathbf{z}'_{it}\boldsymbol{\theta}_t + e_{it}$$

Then the individual specific OLS regressions will yield:

$$(B4) \quad \hat{\boldsymbol{\theta}}_i = \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it} \right)^{-1} \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it}y_{it} \right)$$

Expanding on  $y_{it}$  and noting that  $\boldsymbol{\theta}_i$  is now a scalar vector:

$$(B5) \quad \hat{\boldsymbol{\theta}}_i = \boldsymbol{\theta} + \boldsymbol{\theta}_i + \mathbf{Q}_{zz,T}^{-1} \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it}\mathbf{z}'_{it}\boldsymbol{\theta}_t \right) + \mathbf{Q}_{zz,T}^{-1} \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it}e_{it} \right)$$

where  $\mathbf{Q}_{zz,T}^{-1} = \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \right)^{-1}$ .

Third, consider a set of regressions for each time period  $t$ . Equation (B1) can be written:

$$y_{it} = \mathbf{z}'_{it}(\boldsymbol{\theta} + \boldsymbol{\theta}_t) + v_{it}$$

$$v_{it} = \mathbf{z}'_{it}\boldsymbol{\theta}_i + e_{it}$$

The time-specific regressions yield:

$$(B6) \quad \hat{\boldsymbol{\theta}}_t = \boldsymbol{\theta} + \boldsymbol{\theta}_t + \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_i \right) + \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} e_{it} \right)$$

where  $\mathbf{Q}_{zz,N}^{-1} = \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \right)^{-1}$ .

We construct a preliminary estimate of  $\boldsymbol{\theta}_{it}$  by combining (B3), (B5), and (B6) as follows:

$$(B7) \quad \begin{aligned} \hat{\boldsymbol{\theta}}_{it}^{Prel} &= \hat{\boldsymbol{\theta}}_i + \hat{\boldsymbol{\theta}}_t - \hat{\boldsymbol{\theta}} = \\ &\boldsymbol{\theta} + \boldsymbol{\theta}_i + \mathbf{Q}_{zz,T}^{-1} \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_t \right) + \mathbf{Q}_{zz,T}^{-1} \left( \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} e_{it} \right) + \\ &\boldsymbol{\theta} + \boldsymbol{\theta}_t + \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_i \right) + \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} e_{it} \right) \\ &- \boldsymbol{\theta} - \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_i \right) \\ &- \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_t \right) - \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it} e_{it} \right) \end{aligned}$$

This simplifies to:

$$(B8) \quad \hat{\boldsymbol{\theta}}_{it}^{Prel} = \boldsymbol{\theta} + \boldsymbol{\theta}_i + \boldsymbol{\theta}_t + (\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT}) + (\mathbf{Q}_{ze,N} + \mathbf{Q}_{ze,T} - \mathbf{Q}_{ze,NT})$$

where  $\mathbf{R}_N = \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_i \right)$ ,  $\mathbf{R}_{t,NT} = \mathbf{Q}_{zz,NT}^{-1} \left( \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \boldsymbol{\theta}_t \right)$  and similarly for  $\mathbf{R}_T$  and  $\mathbf{R}_{i,NT}$ , and also  $\mathbf{Q}_{ze,N} = \mathbf{Q}_{zz,N}^{-1} \left( \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} e_{it} \right)$  and similarly for  $\mathbf{Q}_{ze,T}$  and  $\mathbf{Q}_{ze,NT}$ .

The expression in (B8) can be decomposed into three parts. First, there are the true

observation-level coefficients  $\boldsymbol{\theta} + \boldsymbol{\theta}_i + \boldsymbol{\theta}_t$ , then the biases originating from any correlation between the regressors and the heterogeneity (including the fixed effect in the intercept)  $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$ , and terms involving the errors  $(\mathbf{Q}_{ze,N} + \mathbf{Q}_{ze,T} - \mathbf{Q}_{ze,NT})$ . These  $Q$  terms vanish asymptotically given our assumptions on the  $\mathbf{z}_{it}$  and  $\mathbf{e}_{it}$ .

Crucially, as shown in Neal (2016), the bias terms  $\mathbf{R}$  can be calculated to arbitrary accuracy and eliminated from (B8), leaving a consistent estimator of  $\boldsymbol{\theta}_{it}$ .<sup>17</sup> We now explain the procedure: MO-OLS uses  $\hat{\boldsymbol{\theta}}_i$  as a first stage approximation for  $\boldsymbol{\theta}_i$  in  $(\mathbf{R}_N - \mathbf{R}_{i,NT})$  to form  $\hat{\mathbf{R}}_N$  and  $\hat{\mathbf{R}}_{i,NT}$ , and also uses  $\hat{\boldsymbol{\theta}}_t$  as a first stage approximation for  $\boldsymbol{\theta}_t$  in  $(\mathbf{R}_T - \mathbf{R}_{t,NT})$  to form  $\hat{\mathbf{R}}_T$  and  $\hat{\mathbf{R}}_{t,NT}$ . Inserting (B5) and (B6) into these parts of (B8) yields:

(B9)

$$\begin{aligned} (\hat{\mathbf{R}}_N - \hat{\mathbf{R}}_{i,NT}) + (\hat{\mathbf{R}}_T - \hat{\mathbf{R}}_{t,NT}) &= (\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT}) + \mathbf{Q}_{zz,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \mathbf{R}_T \\ &+ \mathbf{Q}_{zz,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \mathbf{R}_N - \mathbf{Q}_{zz,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{z}_{it} \mathbf{z}'_{it} \mathbf{R}_T + \mathbf{z}_{it} \mathbf{z}'_{it} \mathbf{R}_N) + o_p(1) \end{aligned}$$

This expression is equal to the original heterogeneity bias  $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$ , plus a set of additional terms. Thus, the use of (B9) to remove the bias in (B8) eliminates the original bias while introducing a new set of bias terms.<sup>18</sup> Importantly, Neal (2016) shows that these new bias terms must be smaller in magnitude than the original. They can be again approximated by using  $\hat{\mathbf{R}}_N$ ,  $\hat{\mathbf{R}}_{i,NT}$ ,  $\hat{\mathbf{R}}_T$ , and  $\hat{\mathbf{R}}_{t,NT}$  in the remaining biases in (B9). This in turn produces new bias terms that are even smaller in magnitude. This process can be repeated  $L$  times to render the bias arbitrarily small, forming the bias

<sup>17</sup>If the slope heterogeneity is independent of the regressors, then a two-way within transformation prior to estimation gives a consistent estimator of the average coefficient  $\bar{\boldsymbol{\theta}}$ , while the preliminary estimate of  $\boldsymbol{\theta}_{it}$  in (B7) will give a consistent estimate of the observation-level coefficients. However, absent independence, removing this bias is required to provide consistent estimates of the observation-level coefficients and the mean coefficient.

<sup>18</sup>Some of the new bias terms relate to the idiosyncratic error term (which are  $o_p(1)$  under these assumptions) while others relate to the slope heterogeneity.

removed estimates:

$$(B10) \quad \hat{\theta}_{it} = \hat{\theta}_i + \hat{\theta}_t - \hat{\theta} + \sum_{\ell=0}^L (-1)^{\ell+1} \left( \mathbf{Q}_{zz,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{1,\ell} + \mathbf{Q}_{zz,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{2,\ell} - \mathbf{Q}_{zz,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{1,\ell} + \mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{2,\ell}) \right)$$

where  $\Gamma_{1,\ell} = \mathbf{Q}_{zz,T}^{-1} (\frac{1}{T} \sum_{t=1}^T \mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{2,\ell-1})$  and  $\Gamma_{2,\ell} = \mathbf{Q}_{zz,N}^{-1} (\frac{1}{N} \sum_{i=1}^N \mathbf{z}_{it} \mathbf{z}'_{it} \Gamma_{1,\ell-1})$  when  $\ell > 0$ ,  $\Gamma_{1,0} = \hat{\theta}_i$ , and finally  $\Gamma_{2,0} = \hat{\theta}_t$ . This is a Cauchy sequence in  $\ell$ , so a suitable  $L$  can be determined endogenously by programming the sum to stop once the procedure converges to a given level of tolerance. In practice, in all examples we have considered the bias becomes negligible for reasonable values of  $L$ .

Finally, the average coefficients  $\bar{\theta} = (\alpha + E(c_i) + E(c_t), \beta + E(\lambda_i) + E(\lambda_t))$  can be easily constructed by taking a simple average over the sample:

$$(B11) \quad \hat{\theta}_{MO} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\theta}_{it}$$

Please see Neal (2018) for proofs of consistency and asymptotic normality when  $(N, T) \xrightarrow{j} \infty$ , more details on the assumptions needed to maintain these results, and a Monte Carlo analysis of the estimator.

#### APPENDIX C: FORECASTING CORN YIELD WITH TECHNICAL CHANGE

Here we report forecasts of corn yield that incorporate projected technical progress. In the main text we reported predicted changes in yield that account for adaptation to high temperatures, but that hold other aspects of technology fixed. We have a firm historical basis for predicting adaptation, based on comparing historical production function parameters for hot vs. cold counties and time periods. But forecasting more general forms of technical changed into the future is a more speculative exercise.

The econometric models in the main text do provide estimates of technological progress over the 1950-2015 period. In the FE-OLS models, neutral technical progress is captured by the time fixed effect, while in the MO-OLS model technology is more complex, as it may

alter the time fixed effect or the time effects in the slope coefficients for the temperature and precipitation inputs. Accordingly, it is possible to extrapolate trends in to these time effects into the future to forecast future technical progress.

Of course, our sample period of 1950 to 2015 saw dramatic improvements in agricultural technology from the spread of machines including tractors, cutters, harvesters, planters and trucks, commercial fertilizer, insect and weed-resistant hybrid seeds, computers and satellite technology, and genetic modification. It may be optimistic to assume that advances in agricultural technology can continue at such an impressive pace through 2100. But that is in effect what we do here by extrapolating these historical trends.

Here we focus on the MO-OLS model, and forecast both the fixed effect and the coefficient on  $GDD_{it}$  into the future using a VAR(1) system of two equations. We assume that time effects in the KDD coefficient are already captured by our adaptation process, and we ignore time effects in the PREC coefficients because we find no significant time trends on these coefficients in the historical data.

An important issue in specifying the VAR is nature of the time trends. As equation (8) is for log yield, we rule out using trends in  $t$  or  $t^2$  as this permits exponential yield growth in levels. Accordingly, we consider two specifications for the trend terms: The first is to include just  $\log(t)$ , which we call the ‘pessimistic’ scenario of technical progress. The second is to include both  $\log(t)$  and  $\sqrt{\log(t)}$ , which we call the ‘optimistic’ scenario. This is because a negative coefficient on  $\sqrt{\log(t)}$  allows diminishing returns to technology to set in more slowly than when using  $\log(t)$  alone. It is our intent to present both these scenarios of future technology, without commenting on their likelihood. To conserve on space we do not present the estimates of the VAR models. We note, however, that the ‘optimistic’ model had the best in-sample fit of several models we examined. Of course, this does not necessarily make it preferable for out-of-sample forecasts.

Table C1 presents the results of our forecasting exercise broken down by technology scenario, adaptation scenario and RCP scenario. We present mean forecasts across 19 climate models, as well as an 80% prediction interval, at four points in time. Figures C1 to C2 present some key results visually.

Turning to our forecasts, consider first the worst case scenario of pessimistic technology

TABLE C1—THE EFFECTS OF CLIMATE CHANGE ON CORN YIELD (MO-OLS)

Year	Pessimistic Tech Growth		Optimistic Tech Growth	
	No Adapt.	Adapt.	No Adapt.	Adapt.
<u>RCP85</u>				
2030	159 (142, 177)	158 (145, 171)	164 (146, 181)	162 (149, 175)
2050	148 (114, 182)	162 (142, 181)	162 (128, 196)	178 (159, 197)
2080	110 ( 62, 158)	160 (134, 185)	139 ( 84, 195)	206 (179, 232)
2100	73 ( 25, 121)	149 (121, 176)	103 ( 42, 164)	217 (189, 245)
<u>RCP45</u>				
2030	163 (142, 185)	159 (146, 173)	168 (147, 189)	164 (150, 178)
2050	165 (136, 195)	170 (152, 189)	179 (149, 209)	185 (165, 204)
2080	165 (125, 206)	184 (162, 206)	196 (156, 235)	220 (201, 239)
2100	180 (128, 233)	201 (172, 229)	222 (169, 276)	251 (224, 278)
<u>RCP26</u>				
2030	158 (140, 176)	156 (145, 168)	163 (145, 181)	161 (150, 172)
2050	176 (154, 197)	176 (164, 188)	188 (168, 208)	189 (179, 200)
2080	207 (181, 233)	206 (189, 222)	233 (208, 258)	233 (215, 251)
2100	221 (183, 259)	222 (198, 247)	259 (221, 298)	262 (235, 289)

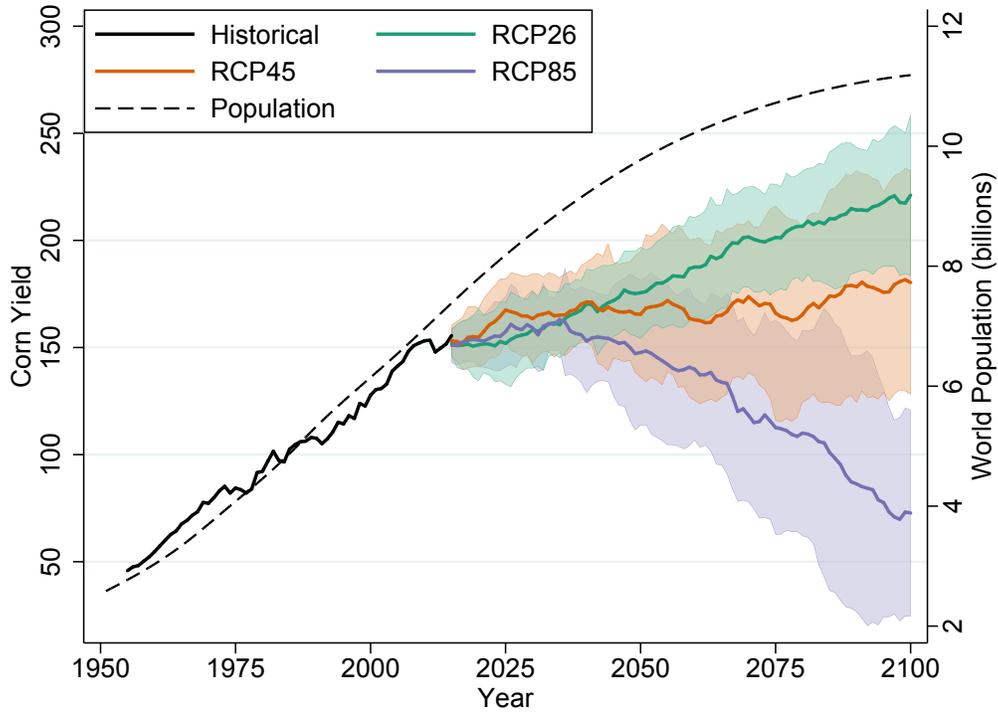
*Notes:* Results are expressed in terms of actual crop yield. Each number represents a model average over nineteen climate models, while the numbers in brackets represents the 80 percent prediction interval of that ensemble average.

growth and no adaptation. As we see in Table C1, under the the RCP85 scenario the mean forecast is a catastrophic drop in corn yield from roughly 150 today to 73 in 2100. With adaptation yield stagnates and is about the same in 2100 as today.

It is important to understand that a stagnation of yield would itself be a catastrophic outcome. As we see in Figure C1 historical corn yield tripled from roughly 50 in the 1950s to roughly 150 in 2015, so rapid growth was the norm for the past 65 years. Furthermore, Figure C1 plots historical world population along with the median forecast from the United Nations Population Division (2017). Note that U.S. corn yields closely track world population growth from 1950 to 2015, as population also roughly tripled from 2.5 to 7.3 billion. The U.S. provides more than a third of the world's corn exports, so a stagnation of U.S. yields would have a devastating effect on world supply.

Next consider the RCP26 scenario of ambitious emissions reductions, while continuing

FIGURE C1. EFFECT OF CLIMATE CHANGE ON CORN YIELD WITH PESSIMISTIC TECH GROWTH AND NO ADAPTATION



*Note: This graph presents forecasts of corn yield by RCP emission scenario using the MO-OLS model where future adaptation is not modelled and technology is forecast using a VAR(1) with a  $\log(t)$  trend term. The solid lines are the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.*

to maintain the pessimistic technology scenario. As we see in Table C1, the mean forecast yield in 2100 is 221 or 222, depending on whether we include adaptation.<sup>19</sup> But as we see in Figure C1, yield growth in this scenario is not nearly as rapid as population growth.

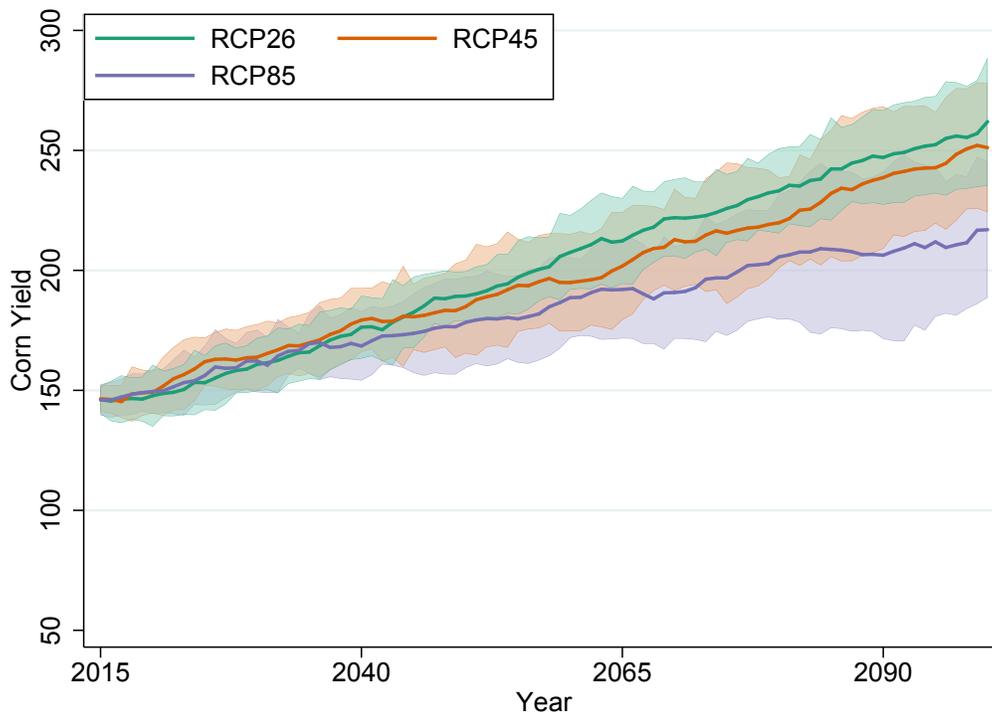
If we turn to the ‘optimistic’ technology assumption, we see in Table C1 that yield forecasts increase rather substantially under all scenarios. In fact, the mean prediction under RCP85 of a 217 yield in 2100 is similar to the prediction under pessimistic technology and RCP26. Thus, being ‘optimistic’ about technology while assuming ‘business as usual’ emissions leads to similar predictions as being ‘pessimistic’ about technology while assuming ‘ambitious’ emissions reductions. Importantly, however, yield in either case falls

<sup>19</sup>Adaptation makes little difference for the mean forecast in this case, because temperature increases are moderate. But it does reduce variance.

well short of keeping pace with population growth.

Finally, consider the best case scenario of optimistic technology growth and ambitious emissions reductions (RCP26). As we see in Table C1, the mean forecast yield in 2100 is 259 or 262, depending on whether we include adaptation. This is roughly what is necessary for yields to increase at a rate that is commensurate with population growth. Even here however, the 80% prediction interval extends down to 235, which is well short of keeping pace with population growth. Thus, even in a best case scenario, climate change creates an environment of considerable *risk* with respect to U.S. agricultural yields.

FIGURE C2. THE EFFECT OF CLIMATE CHANGE ON CORN YIELD WITH OPTIMISTIC TECH GROWTH AND ADAPTATION



Note: This graph presents forecasts of corn yield by RCP emission scenario using the MO-OLS model where future adaptation is not modelled and technology is forecast using a VAR(1) with  $\log(t)$  and  $\sqrt{\log(t)}$  as trend terms. The solid lines are the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.

## APPENDIX D: MODEL AND FORECAST RESULTS FOR SOYBEANS

The case of soybeans offers an interesting comparison to corn. This is the second largest U.S. crop, and both its geographic distribution and growing characteristics are different. Here, we apply the same econometric and forecasting methods as in the main text, except with soybean yield as the dependent variable.

Table D1 presents the pooled panel results, organized in the same manner as Table 1 for corn. We see that the sensitivity of soybeans to high temperatures is lower than it is for corn. In the fourth specification, which models adaptation as a log-linear relationship with  $KDD_{it}$ , the coefficient on  $KDD_{it}$  is more negative while the coefficient on the nonlinear term is positive and statistically significant. This indicates, as it did with corn, that there is heterogeneity in the marginal effect of  $KDD_{it}$  and it has a positive correlation with the level of  $KDD_{it}$ .

TABLE D1—POOLED PANEL ESTIMATES OF THE IMPACTS OF TEMPERATURE ON SOYBEAN YIELDS

Specification	(1)	(2)	(3)	(4)
GDD	0.0003 (0.0001)	0.0005 (0.0001)	0.0004 (0.0000)	0.0006 (0.0000)
KDD	-0.0057 (0.0008)	-0.0048 (0.0003)	-0.0043 (0.0003)	-0.0090 (0.0005)
ln(KDD)*KDD - KDD				0.0009 (0.0001)
Precipitation	0.0017 (0.0004)	0.0012 (0.0001)	0.0012 (0.0002)	0.0013 (0.0000)
Precipitation <sup>2</sup>	-1.3e-06 (3.1e-07)	-9.1e-07 (1.1e-07)	-9.2e-07 (1.4e-07)	-9.1e-07 (3e-08)
Constant	1.6434 (0.4169)	0.9127 (0.2702)	1.2853 (0.1461)	0.5788 (0.0667)
Fixed Effects	Yr	Cty, Yr	Cty, State-Yr	Cty, Yr
Obs.	87,767			
$R^2$	0.6237	0.7523	0.7509	0.3959

*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, and whether the adaptation variable  $\ln(\text{KDD}) * \text{KDD} - \text{KDD}$  is included. Regressions are weighted by average county harvest area for each crop. Standard errors are in parentheses, and are clustered at the state level.

Table D2 presents the results for the MO-OLS model. With soybeans, the mean coefficient for  $KDD_{it}$  (whether unweighted or weighted) is very similar to the pooled model estimates. Moreover, with a standard deviation of 0.0037 there is significantly less heterogeneity in the distribution than with corn. Figure D1 illustrates this, where the 10th to 90th percentiles of the distribution of KDD slope coefficients comprise a much tighter range than they did with corn.<sup>20</sup> This is evidence for comparatively little historical adaptation between counties. This result is intuitive given the much greater success that corn has had in developing seed hybrids that are more resistant to heat.

TABLE D2—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 <sup>2</sup>	-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					

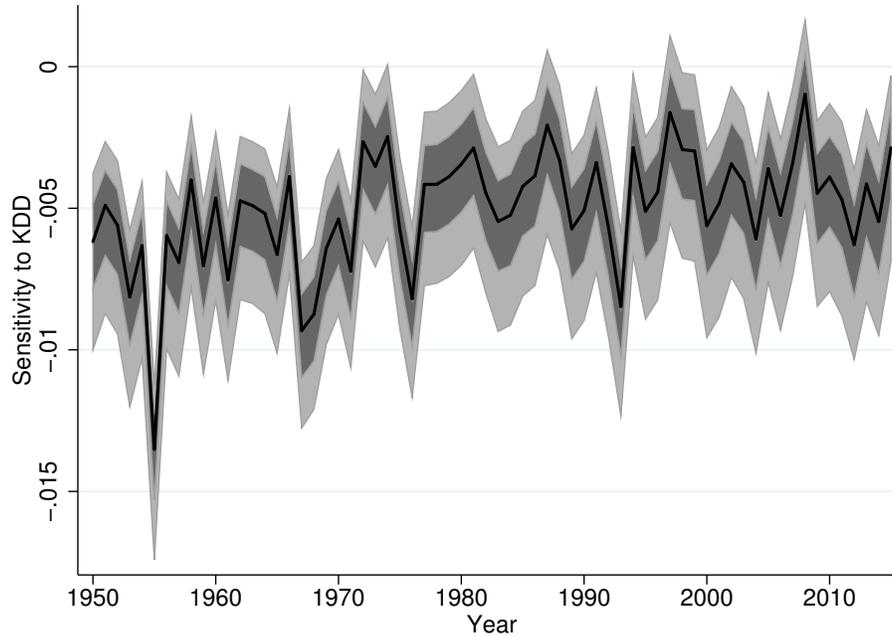
*Notes:* Results exclude states west of the 100th Meridian line. The sample range is 1950-2015, and standard errors are reported in parentheses.

The time trend of aggregate adaptation is somewhat difficult to interpret due to the significant degree of random fluctuation in the median coefficient between years. An unreported structural break test could not reject the null hypothesis that there is no structural break in the parameters of this regression, and fitting a linear trend to the median coefficient over time resulted in an extremely small positive trend.

Figure D2 plots the average values of  $\hat{\beta}_{2it}$  and  $KDD_{it}$  for each county. The correlation between  $\hat{\beta}_{2it}$  and  $KDD_{it}$  is 0.25, which is weaker than for corn, yet still positive. We fit a log-linear curve to these points to obtain the function we use to model adaptation.

<sup>20</sup>The relatively small amount of heterogeneity for soybeans presumably explains why the MO-OLS and pooled

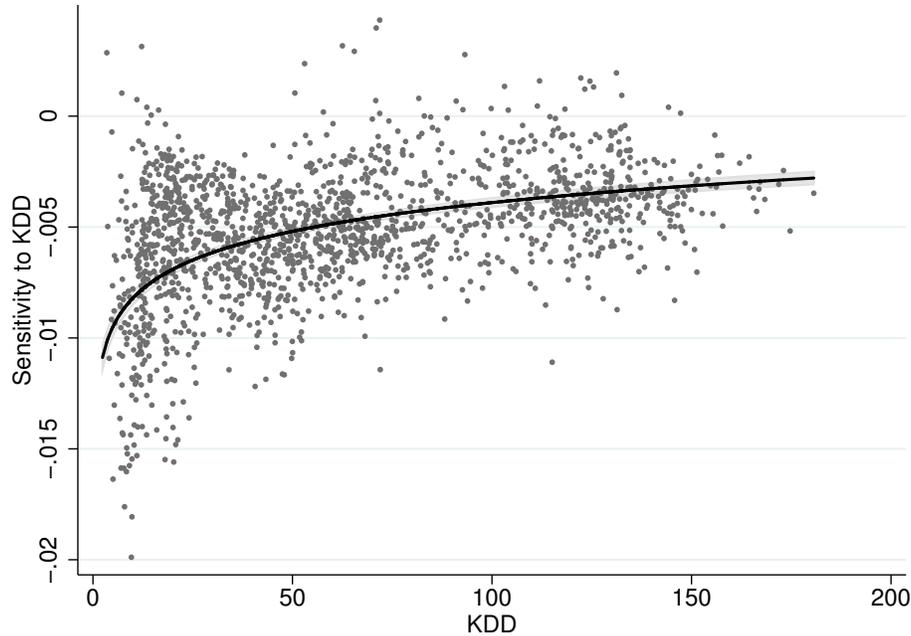
FIGURE D1. DISTRIBUTION OF KDD SLOPE COEFFICIENTS ACROSS TIME AND COUNTIES FOR U.S. SOY



Note: The black line plots the median coefficient of the  $KDD$  variable that is reported in Table D2, 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.

Figure D3 compares the log-linear relationships that are obtained using the FE-OLS and MO-OLS models. Fortunately, the relationship derived from MO-OLS sits comfortably within the 95% confidence interval for the FE-OLS relationship, indicating that they are not significantly different from each other. In contrast to what we find for corn, the MO-OLS relationship here implies less sensitivity to  $KDD_{it}$  at every level of  $KDD_{it}$ . We can use these results to forecast soybean yield into the future under climate change.

Turning to the forecasts, Table D3 presents a summary of the main results for soybean yield with and without adaptation. The format is identical to Table 4 in the text. Without emissions reductions or adaptation, both econometric models predict large reductions in yield. In 2100 the ensemble average is a 51 or 56% reduction (with 34 to 76% prediction interval). However, the RCP45 and RCP26 scenarios illustrate the effectiveness of emissions reductions in reducing the impact on soybean yields, particularly in RCP26 where results are not very different. This is in sharp contrast to the bias we observe in corn.

FIGURE D2. RELATIONSHIP BETWEEN  $\hat{\beta}_{it}$  AND  $KDD_{it}$  FOR U.S. SOYBEANS

*Note: This graph is a scatter plot of the average coefficient to the KDD variable for soybeans in each county that is reported in Table D2 and the average measurement of the KDD variable for the same county. The fitted line was obtained from the regression  $\hat{\beta}_{it} = \alpha_1 + \alpha_2 \ln(KDD_{it})$ .*

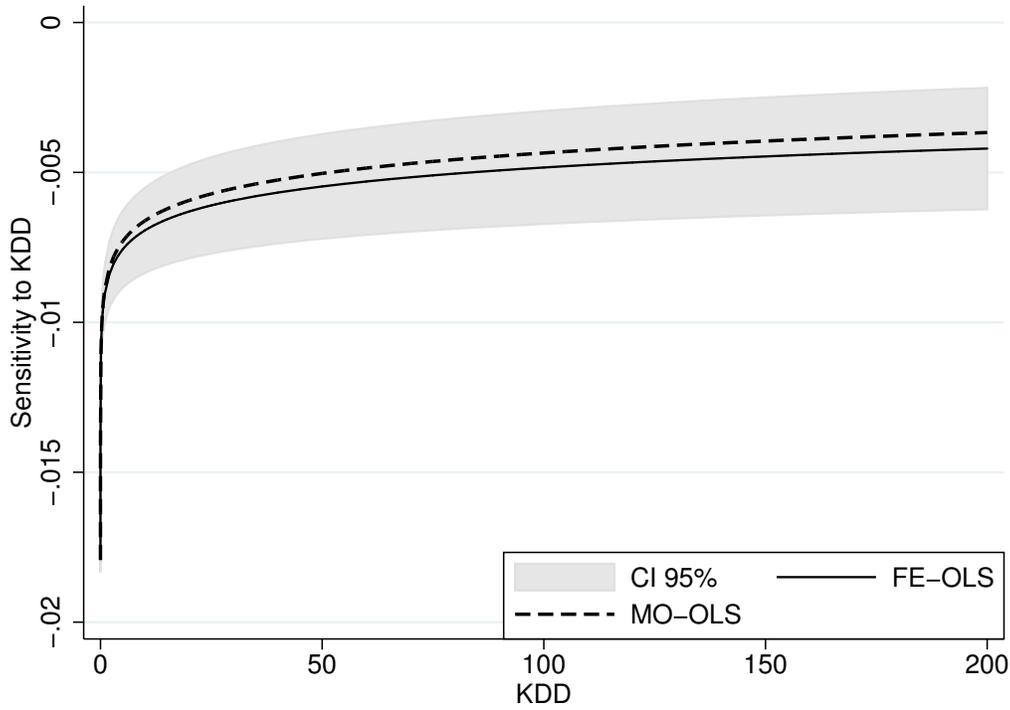
the prediction interval covers the possibility of no damage over the entire forecast horizon.

Figure D4 shows the trajectory of soybean yield under the three RCP scenarios using the MO-OLS estimator with no further adaptation. The RCP85 scenario diverges from RCP45 and RCP26 around 2035, leading to a decline in soybean yield of roughly 50 percent (ensemble average) by 2100. The prediction intervals for both the RCP85 and RCP45 scenarios are fairly large, especially in the latter half of the century, indicating that the climate models offer varied predictions of growing conditions under these two emission pathways. Meanwhile, as we see in Table D3, the two econometric models give very similar forecasts across all three RCP scenarios with no adaptation.<sup>21</sup>

We now turn to the forecasts that incorporate predicted future adaptation. Results are reported in Table D3 and Figure D5. Similar to the case of corn yield, MO-OLS offers a

<sup>21</sup>The one exception is under RCP85 towards the latter half of the century, where MO-OLS becomes slightly more optimistic in its forecasts.

FIGURE D3. COMPARISON OF LOG-LINEAR RELATIONSHIPS DERIVED UNDER MO-OLS AND FE-OLS FOR SOYBEANS

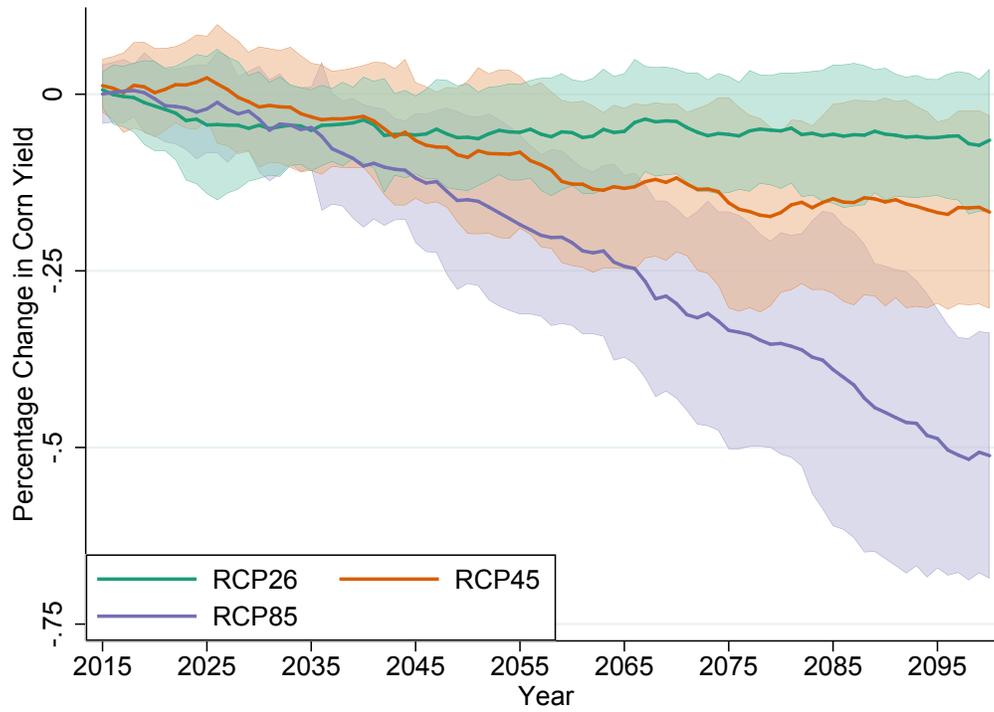


*Note: This graph compares the fitted log-linear relationships between  $\hat{\beta}$  and  $KDD_{it}$  obtained from the county/time specific coefficients estimated with MO-OLS (see Figure D2) vs. the regression results in Table D1 column (4).*

more optimistic forecast than FE-OLS of the effectiveness of adaptation to avert climate change damage. Nevertheless, both estimators predict that the benefits of adaptation are much smaller for soybeans than they are for corn. This is intuitive since the econometric results found much less heterogeneity in the slope coefficient on KDD for soybean yield relative to corn yield. Consider first the RCP85 scenario. Using the FE-OLS approach adaptation reduces the damage to yield from 56 to 51% in 2100. Using the MO-OLS approach, the reduction in yield decreases from 51 to 41% in 2100. Of course, the benefits of adaptation become even smaller under the RCP45 and RCP26 scenarios.

Lastly, Table D4 examines the effectiveness of adaptation and alternative emissions reduction scenarios as ways to avert damage from climate change over the whole forecast horizon. To construct the percent of damage averted we use the same methodology that we

FIGURE D4. THE EFFECT OF CLIMATE CHANGE ON SOYBEAN YIELD BY RCP AND MODEL



*Note: This graph presents forecasts of the percentage change in soybean yield (relative to the 2006-2015 historical average) from the MO-OLS model by RCP and whether future adaptation is modelled. The solid line is the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80 percent prediction intervals.*

used to construct Table 5 in the main text. The results suggest that adaptation can avert only a small proportion of the total damage, between 13 and 22% according to MO-OLS and 3 to 9% according to FE-OLS. In contrast, shifting from the RCP85 to the RCP45 emissions path offers significantly more damage mitigation, with between 36 to 85 percent of the damage averted. The RCP26 path averts more damage still, with between half to all of the total damage being averted. It is clear, relative to the results for corn, that the potential for adaptation is far lower for soybeans, while emissions reductions are more effective. Thus, emissions reductions are clearly essential to prevent significant harm to soybean yields.

TABLE D3—THE EFFECTS OF CLIMATE CHANGE ON SOYBEAN YIELD (PCT)

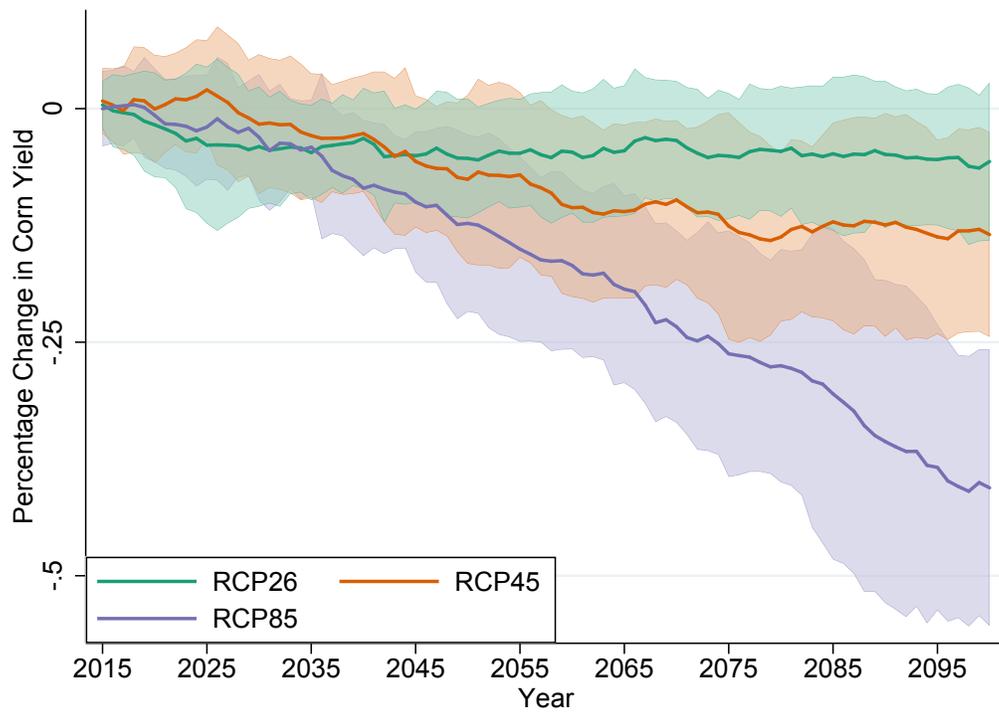
Year	No Future Adaptation		Future Adaptation	
	FE-OLS	MO-OLS	FE-OLS	MO-OLS
<u>RCP85</u>				
2030	-04 (-12, 05)	-04 (-11, 04)	-04 (-13, 06)	-03 (-10, 04)
2050	-16 (-29, -03)	-15 (-27, -03)	-15 (-28, -02)	-12 (-22, -03)
2080	-37 (-55, -19)	-35 (-51, -20)	-34 (-51, -18)	-28 (-40, -15)
2100	-56 (-76, -36)	-51 (-69, -34)	-51 (-69, -33)	-41 (-55, -26)
<u>RCP45</u>				
2030	-02 (-12, 08)	-02 (-10, 07)	-02 (-13, 09)	-02 (-09, 06)
2050	-09 (-22, 03)	-09 (-20, 02)	-09 (-22, 04)	-08 (-17, 01)
2080	-17 (-30, -05)	-17 (-28, -05)	-17 (-29, -05)	-14 (-23, -05)
2100	-18 (-33, -02)	-17 (-30, -03)	-17 (-31, -02)	-13 (-24, -03)
<u>RCP26</u>				
2030	-05 (-14, 03)	-04 (-12, 03)	-06 (-15, 03)	-04 (-10, 02)
2050	-07 (-15, 01)	-06 (-14, 02)	-07 (-15, 02)	-05 (-12, 01)
2080	-06 (-15, 03)	-05 (-13, 03)	-06 (-15, 04)	-05 (-12, 02)
2100	-07 (-19, 04)	-07 (-17, 04)	-08 (-19, 04)	-06 (-14, 03)

*Notes:* Results are expressed in terms of percentage change from the 2006-2015 historical weighted average crop yield. Each number represents the ensemble average over nineteen climate models, while the numbers in brackets represents the 80 percent prediction interval.

TABLE D4—PROPORTION OF CLIMATE CHANGE DAMAGE AVERTED FOR SOYBEANS (PCT)

Estimator	RCP85 + Adapt.	RCP45	RCP45 + Adapt.	RCP26
MO-OLS	17 (13, 22)	60 (36, 85)	64 (42, 86)	78 (53, 103)
FE-OLS	6 (3, 9)	63 (38, 88)	63 (38, 88)	78 (53, 103)

FIGURE D5. THE EFFECT OF CLIMATE CHANGE ON SOYBEAN YIELD BY RCP (MO-OLS WITH FUTURE ADAPTATION)



*Note: This graph presents forecasts of the percentage change in soybean yield (relative to the 2006-2015 historical average) across three RCPs, where the solid line is the average forecast across nineteen CMIP5 climate models, and the shaded areas are the 80% prediction intervals.*