

Climate Change and U.S. Agriculture: Accounting for Multidimensional Slope Heterogeneity in Panel Data

By MICHAEL KEANE AND TIMOTHY NEAL*

Draft: April 10, 2020

We study potential impacts of future climate change on U.S. agricultural productivity using county-level yield and weather data from 1950 to 2015. To account for adaptation of production to different weather conditions, it is crucial to allow for both spatial and temporal variation in the production process mapping weather to crop yields. We present a new panel data estimation technique, called mean observation OLS (MO-OLS) that allows for spatial and temporal heterogeneity in all regression parameters (intercepts and slopes). Both forms of heterogeneity are important: We find strong evidence that production function parameters adapt to local climate, and also that sensitivity of yield to high temperature declined from 1950-89. We use our estimates to project corn yields to 2100 using 19 climate models and three greenhouse gas emission scenarios. We predict unmitigated climate change will greatly reduce yield. Our mean prediction (over climate models) is that adaptation alone can mitigate 36% of the damage, while emissions reductions consistent with the Paris targets would mitigate 76%.

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

Leading scientific and environmental institutions warn that future climate change may severely impact agricultural productivity and global food supply (Porter et al. 2014). But studies that estimate the sensitivity of agriculture to climate, in hopes of obtaining insight into effects of future climate change, have given mixed results. Projected impacts of climate change on U.S. agriculture in particular range from severe damage to productivity (e.g. Schlenker and Roberts 2009) to negligible damage (e.g. Butler and Huybers 2013).

* 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).

Resolving this uncertainty 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 adaptation by agricultural producers. This may entail using more heat tolerant hybrids, improved water retention in fields, irrigation, adjusting sowing density, etc. The existence of adaptation in the historical data implies spatial and temporal heterogeneity in the production process mapping weather to crop yields. Thus, we would expect the production function to exhibit region and time fixed effects in both intercepts and slopes, where the fixed effects are correlated with regional and temporal variation in climate. This complex heterogeneity structure creates important challenges for the proper econometric modelling of climate impacts - challenges that have not been fully addressed in prior literature.

We investigate these issues using weather and crop yield data for U.S. counties from 1950 to 2015. We focus on corn, as it is historically the largest crop in the U.S. in terms of tonnage. Prior to our work, several authors investigated agricultural adaptation using panel data regressions of yield on temperature and precipitation (see, e.g., Schlenker and Roberts 2009, Burke and Emerick 2016). Particularly relevant are Butler and Huybers (2013), who allow for heterogeneity in heat sensitivity across counties, and Roberts and Schlenker (2012), who allow for variation over time (common to all counties). But prior work has not accommodated adaptation across both regions and time in a flexible way.

Our first main contribution is to present a new panel data method that allows us to flexibly estimate the extent of historical adaptation to high temperatures. As in the literature cited above, we define “adaptation” broadly as including all sources of co-variation between heat and heat sensitivity.¹ As variation in heat sensitivity can occur across space *and* over time, we develop a method we call ‘mean observation OLS’ (MO-OLS) that is able to feasibly estimate in large panel datasets a model that contains both county and time fixed effects in both the intercept and slope coefficients. It is a more flexible approach for modelling adaptation than in prior studies, and we find that it offers substantial improvements in fit over existing econometric models of yield.²

¹Adaptation as defined here includes farmers’ active adaptation of growing techniques to temperature, as well as any other factors (not under farmers control) that cause yields to be less heat sensitive under hotter conditions.

²In a follow up article, Keane and Neal (2020) show a MO-OLS model can produce more accurate out-of-sample forecasts of corn yield than a deep neural net using the same inputs.

The MO-OLS approach we develop here should be useful in contexts beyond the present application. There are many cases where both cross-sectional heterogeneity and temporal parameter variation are of interest, so a method that can handle both in a computationally feasible manner should have broad applicability.³ For example, in macroeconomics, time-varying parameter models are often used to study how key economic relationships - such as fiscal multipliers or the Beveridge curve - change over time in response to government policy or changes in the economic environment. Modeling unit-specific slope heterogeneity at the same time would allow one to expand this analysis to account for heterogeneity in those relationships by country, region or industry.⁴

Returning to our study of corn yields, we also consider a simple parametric alternative to MO-OLS where, in the yield regression, we let the coefficient on high temperatures be a theoretically motivated non-linear function of temperature itself – allowing the sensitivity of yield to high temperatures to decline as the frequency of high temperatures increases. We find a fairly close agreement between the parametric estimates of temperature sensitivity and the MO-OLS (fixed-effects) estimates. Both approaches imply that the heat sensitivity of corn yields declines with temperature according to a log-linear relationship similar to that implied by our simple theory.⁵

Using both MO-OLS and the parametric approach, we find that significant adaptation occurred across hot/cool counties, as well as over time from 1950 to 1989.⁶ This is contrary to earlier conclusions of Schlenker and Roberts (2009) and Burke and Emerick (2016) that there is little evidence for adaptation in corn yields. Yet, we also find that adaptation stalled after 1989. This is consistent with a trend towards higher sowing densities, which

³Notably, it may be important to model slope heterogeneity both across units and over time even if one is only interested in average marginal effects. Appendix D shows how conventional panel data methods like fixed-effects or mean-group regression can give severely biased estimates of average marginal effects in environments where both dimensions of heterogeneity exist.

⁴For example, Auerbach and Gorodnichenko (2012) study how fiscal multipliers vary over business cycles, while Auerbach et al. (2020) study heterogeneity across regions/industries, but neither paper considers both. Benati and Lubik (2014) study how the Beveridge curve changes over time, while many papers study how it varies by country/region, but we are unaware of work that studies both. Johnson and Papageorgiou (2020) review work on cross-country growth regressions, including work that models either cross-sectional heterogeneity or time-variation in convergence rates, but not both. For additional examples, see Hsiao and Pesaran (2008) who survey econometric models with slope heterogeneity across either time or space.

⁵Our approach is related to Butler and Huybers (2013), who let the sensitivity of yield to high temperature depend on average temperature in a county. But they only allow for heterogeneity in heat sensitivity across counties and not over time. We show that time effects are also very important.

⁶We also show, in Appendix B, that ignoring adaptation across counties (i.e., ignoring fixed effects in slopes) in econometric models of crop yield leads to underestimation of yield sensitivity to high temperatures. Our estimates imply this underestimation is on the order of 60 to 85 percent.

generate higher yield in good years, but worsen heat stress in drought years (see Lobell et al. 2014). The substantial expansion of crop insurance in the early 1990s may have encouraged this trend, by reducing the incentives for farmers to adapt to extreme heat (see Annan and Schlenker 2015).

Our second main contribution is to use our econometric models to project future corn yields annually until 2100, using temperature and precipitation projections from 19 climate models under three greenhouse gas emissions scenarios. We provide projections both with and without adaptation, utilizing our estimated historical relationship between temperature and heat sensitivity to predict adaptation. This approach relies on the admittedly strong assumption that future adaptation to high temperatures can be forecast based on historical adaptation patterns. But we argue this is an important benchmark for assessing how much damage mitigation may be plausibly be expected from adaptation. Given this assumption, our econometric models allow us to compare the effectiveness of adaptation vs. emissions reduction as ways to mitigate crop damage. By using 19 climate models we also quantify the variability in forecast outcomes across models.

The UN-IPCC uses a set of “representative concentration pathways” (RCPs) for greenhouse gas emissions based on different policy scenarios. We predict “business as usual” emissions growth (the RCP85 scenario), combined with no adaptation to climate change, will cause catastrophic damage to corn yield. Our central projection is a 70% reduction in yield by 2100, with an 80% prediction interval from 51 to 89%.⁷ We predict adaptation consistent with that observed in the historical data may avert from 29 to 44% of the total damage over 2020-2100. Thus, while adaptation may be important, it seems implausible that the U.S. can rely on adaptation alone to protect corn yields from substantial impacts.

Next, we consider “moderate” emissions reductions, consistent with the UN-IPCC’s “RCP45” scenario, which is somewhat more ambitious than current government pledges. Our central projection is that this would avert 55% of the total damage to corn yield without any adaptation, or 61% if we factor in adaptation. Finally, we consider substantial emissions reductions, consistent with the RCP26 scenario, which follows from the most

⁷The wide prediction interval stems from substantial disagreement across climate models about future weather conditions in the corn growing counties of the US. In the results section we’ll show that uncertainty about future yields is due more to uncertainty across climate models than across econometric models.

ambitious targets under the Paris agreement. We predict this would avert 76% of damage, even without adaptation. Thus, while adaptation has the potential to avert a meaningful fraction of yield damage, it seems that substantial emissions reductions will be necessary to avert most of the potential damage from climate change.

The paper is structured as follows: Section I presents a simple model of agricultural yield with adaptation, to provide a coherent framework for the empirical work. Section II discusses our econometric methods, including the MO-OLS estimator. Section III describes our data. Section IV presents our main econometric results for corn yield. Section V presents our projections for corn yield through to 2100. Section VI concludes. A Mathematical Appendix provides proofs of MO-OLS properties. Eight online appendices extend the results of the article, including Monte Carlo simulations and results for soybeans.

I. A Simple Model of Agricultural Yield with Weather and Adaptation

There is a long tradition in agricultural economics of estimating production functions for corn yield. Starting from the classic work by Wallace (1920), researchers have estimated regressions for annual yield as a function of temperature and precipitation during the growing season. Recent work uses modern panel-data techniques to control for county fixed effects (e.g., to account for soil quality) and common time effects. Here, we present a simple model that (i) rationalizes the conventional econometric specification and (ii) shows how adaptation generates fixed-effects heterogeneity in slopes.

We start with a production function for corn that incorporates measures of temperature:

$$(1) \quad Y_{it}/C_{it} = A_t \mu_i I_{it} (1 + \beta_1 (GDD_{it} - GDD_{min}) + \beta_2 KDD_{it})$$

Here Y_{it} is output of corn for farmer i at time t and C_{it} is the number of acres planted, so the dependent variable is yield. A_t is total factor productivity at time t and μ_i is an area effect (e.g., soil quality). The term I_{it} is a composite of conventional inputs; e.g., capital, labor, fertilizer. We assume the functional form of I_{it} is common for all i .

The variables GDD_{it} and KDD_{it} capture effects of temperature. ‘Growing degree days,’ or GDD_{it} , is the total hours in the growing season that the crop experiences beneficial temperature. ‘Killing degree days,’ or KDD_{it} , is the total hours of harmful temperatures.

GDD_{min} is the minimum level of GDD needed to obtain a positive yield. The percent shift in yield due to temperature factors is $x = \beta_1(GDD_{it} - GDD_{min}) + \beta_2KDD_{it}$. We omit precipitation for simplicity, but it is included in the econometric models.

Now we show how the simple model in (1) can rationalize the yield models estimated in the literature. Taking the log of crop yield and using the approximation $\ln(1+x) \approx x$, which is accurate, as values of x outside the +20% to -20% range are rare, we obtain:

$$(2) \quad y_{it} = \ln(A_t) + (\ln(\mu_i) - \beta_1 GDD_{min}) + \ln(I_{it}) + \beta_1 GDD_{it} + \beta_2 KDD_{it}$$

where $y_{it} = \ln(Y_{it}/C_{it})$. Equation (2) is akin to that estimated in several recent papers. Typically, these papers use fixed effects over i and t to capture the A_t , μ_i and I_{it} terms.

This approach is valid if we can write $\ln(I_{it}) = f_i + f_t + \epsilon_{it}$, where f_i and f_t denote variation in inputs over farms/time that are captured by unit and time fixed effects,⁸ while ϵ_{it} is an idiosyncratic factor uncorrelated with time t weather shocks. Then, ϵ_{it} provides the econometric error for estimation of (2). As a practical matter, research on corn yield emphasizes the role of weather in the production function (using i and t dummies to capture other inputs) for two reasons: variation in yield over time - beyond what is explained by farm/time effects - is well explained by variation in weather (Tannura et al. 2008, Westcott and Jewison 2013, Wang et al. 2012), and modification of conventional inputs after weather shocks are revealed has only minor effects on yield.⁹

Next, we extend this simple model to account for adaptation. Assume that by bearing a cost farmers can reduce sensitivity of yield to high temperature (e.g., paying a premium for drought resistant seed). To capture this, let the KDD coefficient be $\beta_{2it} = s/(1 + \alpha_{it})$, where $s < 0$ is the effect of high temperatures on yield absent any adaptation, while α_{it} denotes units of adaptation purchased by farmer i in period t . Letting γ denote the price of adaptation, profit for farmer i at time t is $\pi_{it} = p_t Y_{it} - \gamma \alpha_{it} - r_t I_{it}$ where p_t is the price of the crop, and r_t is the rental rate per unit of I_{it} .¹⁰ To maximize profit, farmers

⁸For instance, if farmers use a common technology and face common input and output prices, then the time factor f_t will capture proportional year-to-year shifts in inputs across all farmers in response to changes in input or expected output prices.

⁹In contrast, if changes in I_{it} could mitigate unanticipated shocks to KDD , then β_2 would be biased towards zero. The traditional agricultural economics literature (implicitly) rules this out when it estimates versions of (2). But (2) could still be interpreted as a reduced form, provided input prices are captured by the fixed effects.

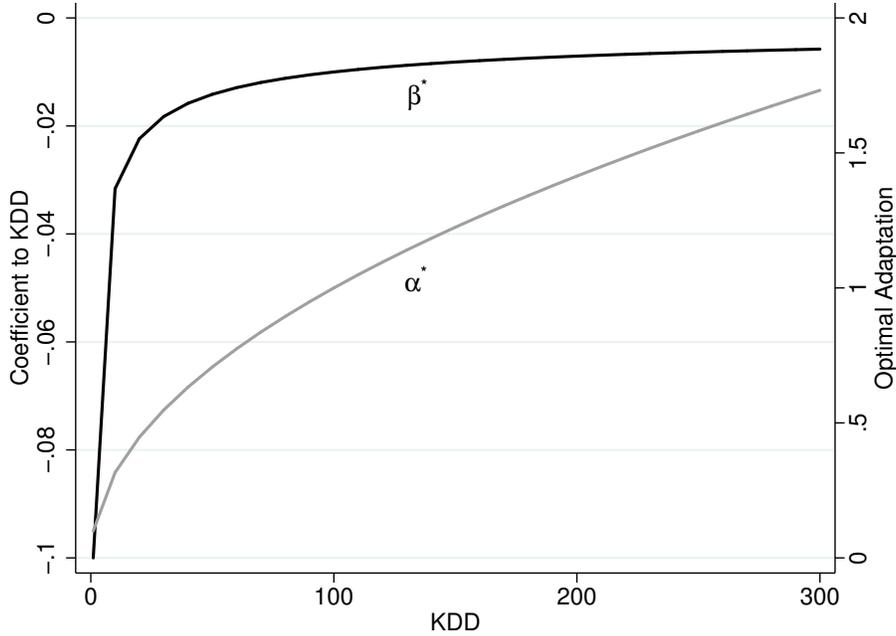
¹⁰Recall I_{it} is a composite of conventional inputs such as capital and labor. Given a common homothetic

purchase the optimal level of adaptation. Setting $\partial\pi/\partial\alpha = 0$ we obtain:

$$(3) \quad \alpha_{it}^* = \sqrt{\frac{p_t(C_{it}A_t\mu_i I_{it})(-s)KDD_{it}}{\gamma}} - 1$$

Thus as KDD_{it} increases the optimal level of adaptation increases. Farms in hotter counties/time periods have more incentive to adapt. Figure 1 plots the optimal level of adaptation α_{it}^* and the implied coefficient β_{2it} against KDD_{it} . As we see, the relationship between KDD_{it} and the heterogeneous coefficient β_{2it} closely resembles a log-linear function. Of course, this depends on the functional form in (2). However, this prediction of the simple model is testable, and it is supported by our estimates. More importantly,

FIGURE 1. RELATIONSHIP BETWEEN α_{it}^* , β_{2it}^* , AND KDD_{it}



Note: This graph presents optimal values of β_{2it}^ and α_{it}^* (secondary vertical axis) as a function of KDD_{it} . We use a normalization of p such that $pC_{it}A_tI_{it} = 1$ and $s = -0.01$, which produce values of β_{2it}^* that is within a range consistent with our econometric analysis.*

the model illustrates how adaptation renders the coefficient on KDD_{it} heterogeneous, and

technology (to rule out scale effects) and common factor prices, all farmers will use inputs in the same proportions. Then r_t can be interpreted as the constant unit price of the optimal bundle of inputs.

the heterogeneity will be correlated with the regressor itself. To capture this requires an econometric method that allows for fixed effects in slopes (not just intercepts).

Letting *both* the KDD_{it} and GDD_{it} coefficients be heterogeneous across i and t , and using time and area fixed-effects (which we denote by τ_t and c_i) to pick up the A_t, μ_i, f_t and f_i terms, we obtain a modified version of equation (2) of the form:

$$(4) \quad y_{it} = \tau_t + c_i + \beta_{1it}GDD_{it} + \beta_{2it}KDD_{it} + \epsilon_{it}$$

In Section II.C we explain how MO-OLS makes it feasible to estimate models like (4) with fixed effects in intercepts and slopes in large panels. In Section IV we use versions of (4) to study the effects of temperature on crop yield across U.S. counties and over time.

Note that our simple model is at the farmer level, while in Section IV we will estimate models at the county level. In Appendix A we show that the main predictions of the farmer-level model (i.e., that β_{2it} is positively correlated with KDD_{it} and the relationship is approximately log-linear) carry over to a county-level model.

II. Econometric Methods

A. Previous Approaches to Modelling Crop Yield

Several recent papers estimate effects of temperature on crop yield using the “degree day” approach. This recognizes that moderate temperatures are beneficial for yield, while high temperatures cause damage. For instance, Schlenker and Roberts (2009) estimate county-level panel-data yield regressions of the following form:

$$(5) \quad y_{it} = c_i + \tau_t + \sum_{j=0,3,\dots}^{39} \beta_j(DD_{j,it} - DD_{j+3,it}) + \beta_{40}PREC_{it} + \beta_{41}PREC_{it}^2 + \epsilon_{it}$$

where y_{it} is log yield for county i , year t . The degree day measure $DD_{j,it}$ is total time over the growing season that the crop experiences temperatures above $j^\circ\text{C}$. This model allows the effects of temperature to differ across 3° degree bands. $PREC_{it}$ is total precipitation during the growing season. The c_i and τ_t are county and time fixed effects.

Other authors simplify this model by splitting degree days into those above and below

29°C, which is considered a critical threshold for corn.¹¹ Cumulative beneficial temperatures (“growing degree days”) are given by $GDD_{it} = DD_{0,it} - DD_{29,it}$ and harmful temperatures (“killing degree days”) are given by $KDD_{it} = DD_{29,it}$. For instance, Lobell et al. (2011) and Burke and Emerick (2016) estimate equations similar to the following:

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

This is basically equation (2) of Section I, with county and time effects used to capture the A_t , μ_i and I_{it} terms, and precipitation added. In equations (5) and (6) the key parameter of interest is $\beta_2 < 0$, which captures the extent to which high temperatures reduce crop yield. The county fixed effects c_i capture intercept heterogeneity that is county i specific (such as soil quality), while the time effects τ_t capture changes in total factor productivity that are common across counties but vary between years t .

Several studies use the degree day approach to estimate the extent of adaptation to high temperatures, motivated by the idea 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 (5) and (6), splitting the sample into northern and southern U.S. states and also by 1950-1977 and 1978-2005 periods. Surprisingly, they find coefficients do not differ significantly by region or time, which they take as evidence that historical adaptation has been very limited.

Burke and Emerick (2016) note that, in conventional county fixed-effects models with homogeneous slopes, effects of temperature are identified off of short-run (i.e., annual) deviations of temperature from county means. They adopt a “long difference” approach to try to estimate the response of yields to long-term changes in temperature. Specifically, they estimate a “long difference” regression similar to:

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

where ΔX_{is} is defined as the change in the average value X in county i of State s from

¹¹At temperatures below 29°C the corn plant can be viewed as a machine for converting heat, water and nutrients into corn. But temperatures above 29°C 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.

1998-2002 to 1978-82. Comparing estimates of (7) with a conventional panel-data fixed effects model (6) estimated on annual county-level from 1980 to 2000, they find similar coefficients on KDD . Thus, they infer that adaptation was fairly minor over this period.¹²

It is important to recognize that all the approaches we have discussed rely on fixed-effects models with homogeneous slopes to provide consistent estimates of the sensitivity of crop yield to temperature. Appendix B demonstrates that this assumption is unlikely to hold in practice, as correlation between the (heterogeneous) slope coefficients and temperature, arising due to adaptation, leads to bias in estimating temperature effects.

In contrast to studies we have discussed so far, Butler and Huybers (2013) run separate regressions by county for 1981 to 2008, and conclude from the county-specific coefficients that substantial adaptation has occurred across counties with different climates. In projecting effects of climate change on crop yields, they argue that losses could be halved by adaptation. The use of county-specific regressions avoids the criticism that the coefficients on temperature are identified from short-run variation (i.e., weather rather than climate). Roberts and Schlenker (2012) study whether corn yield has become less sensitive to high temperatures over time, by estimating time-varying regression parameters, and conclude it has not. These approaches are related to ours, in that slope parameters are either county-specific or time-varying. We generalize these approaches by allowing for *both* county and time effects in slopes, thus allowing for adaptation *both* across counties and over time.

Finally, outside of agriculture, Deschênes and Greenstone (2011) model adaptation by interacting current temperature with mean temperature of a region in order to estimate effects of climate change on mortality. Dell et al. (2012) use this approach 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. Two Approaches to Modelling Adaptation

Here we propose two complementary approaches to modelling crop yield that account for potential adaptation to high temperatures. A useful starting point is the conventional

¹²Liang et al. (2017) and Ortiz-Bobea et al. (2018) consider a related question of whether total factor productivity of U.S. agriculture (across multiple crops) has become more sensitive to the climate over time.

two-way fixed effects specification (FE-OLS) that we repeat for convenience:

$$(8) \quad y_{it} = c_i + \tau_t + \beta_1 GDD_{it} + \beta_2 KDD_{it} + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

In Appendix B we show how, given heterogeneity in the KDD coefficient (β_{2it}) induced by adaptation, the FE-OLS estimate of β_2 in (8) is likely to be an upward biased (i.e., toward 0) estimate of the mean KDD coefficient.¹³

In Section I we presented a simple model that rationalized the widely-used specification in (8). When we extended the model to include adaptation, it implied that the effect of high temperatures on crop yield is log-linear function of KDD_{it} . Thus, one way to capture adaptation is the parametric specification:

$$(9) \quad y_{it} = c_i + \tau_t + \beta_1 GDD_{it} + \beta_{20} KDD_{it} + \beta_{21} (\ln(KDD_{it}) * KDD_{it} - KDD_{it}) \\ + \beta_3 PREC_{it} + \beta_4 PREC_{it}^2 + \epsilon_{it}$$

which implies that the marginal effect of KDD_{it} on y_{it} is the log-linear function:

$$(10) \quad \frac{\partial y_{it}}{\partial KDD_{it}} = \hat{\beta}_{20} + \hat{\beta}_{21} \ln(KDD_{it})$$

The model in equation (9) can be simply estimated using FE-OLS. If $\hat{\beta}_{21} > 0$ it implies that the adverse KDD effect is smaller in hotter counties or time periods.¹⁴ But this approach relies on the parametric assumptions in Section I being correct.

Our second - and more novel - approach to estimating adaptation is to estimate a model with both spatial and temporal heterogeneity in the slope coefficients, as in:

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

This approach is more flexible, as we do not have to specify a particular form of nonlinearity for the KDD coefficient. Instead, we allow the slope heterogeneity to be correlated with the regressors. Then, by analyzing the distribution of the estimates $\hat{\beta}_{2it}$ post-estimation,

¹³This is because FE-OLS ignores the correlation between KDD_{it} and β_{2it} generated by adaptation.

¹⁴Alternatively, we could allow the KDD coefficient to be a log-linear function of a measure of average temperature in a county, as in Butler and Huybers (2013). But this does not capture adaptation over time within counties.

we can determine the nature of the relationship between KDD_{it} and $\hat{\beta}_{2it}$ (if any). Of course, estimating (11) without further restrictions would result in more unknown parameters than data points. To achieve identification we restrict attention to additive heterogeneity across the county/time dimensions, as in:

$$(12) \quad \beta_{kit} = \beta_k + \lambda_{ki} + \theta_{kt}, \quad k = 1, \dots, 4$$

This set up captures adaptation across *both* counties and time periods.¹⁵ The implication of additive fixed effects in slopes is that each county’s *relative* sensitivity to weather is fixed over time. Time effects shift all county’s sensitivities up or down to the same degree.

The direct way to estimate (11)-(12) is via OLS, where each regressor is interacted with a full set of dummies for each i and t . We call this the ‘brute force’ approach. In a panel with large N and T this is infeasible as the regressor matrix grows extremely large.

Recall that in large panels it is standard to estimate models with fixed effects in intercepts by demeaning the data for each unit prior to running OLS. Frisch and Waugh (1933) showed this simple procedure gives the fixed effects estimator. Similarly, one can estimate models with unit fixed effects in both intercepts and slopes by running OLS regressions at the unit level (Pesaran and Smith 1995). But this approach is not possible for models with time effects, which exhaust all degrees of freedom at the unit level. As a result, the literature lacks a computationally practical approach to estimate models with heterogeneous slopes in large panels in the presence of time effects.

Accordingly, in section II.C we present a computationally practical method to estimate models with additive slope heterogeneity over two dimensions that may be correlated with the regressors, as in (11)-(12). We call this the ‘mean observation OLS’ estimator (MO-OLS), and it is numerically equivalent to the ‘brute force’ OLS approach.

Importantly, the data variation that identifies the KDD coefficient(s) in (11)-(12), using the MO-OLS approach, is fundamentally different from that in the FE-OLS model (8). In FE-OLS, slopes are identified from the response of yield to idiosyncratic variation in the regressors (i.e., local weather shocks). But, as we show in Appendix E, if KDD has a permanent/transitory structure, and we use MO-OLS, then the λ_i will be identified from

¹⁵Note the three terms β_k , λ_{ki} and θ_{kt} are only separately identified given location normalizations. The reason for writing β_{kit} as the sum of these three terms will become apparent in Section II.C.

the response of yield to permanent differences in counties' climates, while the θ_t will be identified from responses to aggregate time effects in U.S. weather. Thus, MO-OLS can identify long-term adaptation by farmers (in the form of slope heterogeneity) that is driven by county level climate, or by common time effects.

It is important to be clear about the types of adaptation our approach captures. We define "adaptation" broadly as including all factors that alter sensitivity of yield to weather, which we model as heterogeneity across counties/time in marginal effect of KDD on crop yield. This captures several forms of active farmer adaptation to weather, including adoption of heat resistant seed hybrids, irrigation, improved water retention in the fields, and sowing density. But it does not capture other adaptations such as crop switching, changes to the growing season, or land use changes.¹⁶ Heterogeneity in slope parameters will also capture "natural" adaptation, by which we mean any inherent nonlinearity in the relation between temperature and yield generated by plant biology. For the purpose of obtaining unbiased projections of effects of climate change on crop yield, a model should account for both farmer and natural adaptations. But a limitation of our approach is we may subsume cross-county variation in heat sensitivity due to factors that cannot be altered further.

C. The MO-OLS Algorithm

Here we present the mean observation OLS (MO-OLS) procedure that we use to estimate (11)-(12). This is the first computationally feasible panel-data estimator that allows for fixed-effects slope heterogeneity over space and time. Consider the following generic model that includes fixed effects in both intercept and slopes:

$$(13) \quad y_{it} = \boldsymbol{\beta}'_{it} \mathbf{x}_{it} + u_{it}$$

for units $i = 1, \dots, N$ and time periods $t = 1, \dots, T$, where $\mathbf{x}_{it} = (1, x_{1it}, \dots, x_{Kit})'$ is a $(K + 1) \times 1$ vector of regressors, $\boldsymbol{\beta}_{it} = (\beta_{0it}, \beta_{1it}, \dots, \beta_{Kit})'$ is a $(K + 1) \times 1$ vector of coefficients that vary across individuals and over time, and u_{it} is the idiosyncratic error term. Note that \mathbf{x}_{it} includes a constant term, which accordingly allows for intercept

¹⁶However, in Section IV.D we examine heterogeneity in the effects of climate change across counties. This sheds light on the potential for land use changes to mitigate damage.

heterogeneity across i and t . \mathbf{x}_{it} may also include lags of the dependent variable or any of the regressors as needed. We assume **A.1**: u_{it} is iid, **A.2**: the regressors are weakly exogenous $E(x_{kist}u_{it}) = 0 \forall k$ for $s \geq t$, and **A.3**: $E(u_{it}^2|x_{kit}) < \infty \forall k$.

We further assume **A.4**: the coefficient heterogeneity is additively separable, such that $\beta_{it} = \beta + \lambda_i + \theta_t$, where $\beta = (\beta_0, \beta_1, \dots, \beta_K)'$ is the constant effect across all observations, $\lambda_i = (\lambda_{0i}, \lambda_{1i}, \dots, \lambda_{Ki})'$ are the individual effects that vary across every unit in the panel, and $\theta_t = (\theta_{0t}, \theta_{1t}, \dots, \theta_{Kt})'$ are time effects that vary between each time period.

A ‘brute force’ approach to estimating (13) is to run OLS on a model that includes: (i) dummy variables for each i and t (to capture fixed effects in the intercept), and (ii) a complete set of interaction terms between the regressors and the i and t dummies (to capture unit/time fixed effects in slopes). This is computationally infeasible in medium to large panels, as it involves $(N + T)(K + 1)$ regressors, making it impractical to store and invert $X'X$, or to solve the linear system $(X'X)\beta = X'Y$.¹⁷ See Appendix C for details.

Instead, MO-OLS constructs consistent estimates of β_{it} by running a series of feasible regressions and then removing the resulting bias. MO-OLS does this by combining three types of regressions: pooled, i -specific, and t -specific. First, rewrite (13) as:

$$y_{it} = \mathbf{x}'_{it}\beta + v_{it}$$

$$v_{it} = \mathbf{x}'_{it}\lambda_i + \mathbf{x}'_{it}\theta_t + u_{it}$$

Consider the pooled OLS estimator of β :

$$(14) \quad \hat{\beta} = \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it}\mathbf{x}'_{it} \right)^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it}y_{it} \right)$$

Expanding on y_{it} and simplifying yields:

$$(15) \quad \begin{aligned} \hat{\beta} = & \beta + \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it}\mathbf{x}'_{it}\lambda_i \right) \\ & + \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it}\mathbf{x}'_{it}\theta_t \right) + \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it}u_{it} \right) \end{aligned}$$

¹⁷In panels with very large N , such as marketing datasets with many products, even models with fixed effects in the intercept alone are computationally daunting, unless one uses the Frisch and Waugh (1933) ‘within’ transformation.

where $\mathbf{Q}_{xx,NT}^{-1} = \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \right)^{-1}$. Next, consider the unit-specific regressions:

$$y_{it} = \mathbf{x}'_{it}(\boldsymbol{\beta} + \boldsymbol{\lambda}_i) + v_{it}$$

$$v_{it} = \mathbf{x}'_{it}\boldsymbol{\theta}_t + u_{it}$$

The unit-specific OLS regressions yield:

$$(16) \quad \hat{\boldsymbol{\beta}}_i = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \right)^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} y_{it} \right)$$

Expanding on y_{it} and simplifying yields:

$$(17) \quad \hat{\boldsymbol{\beta}}_i = \boldsymbol{\beta} + \boldsymbol{\lambda}_i + \mathbf{Q}_{xx,T}^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \boldsymbol{\theta}_t \right) + \mathbf{Q}_{xx,T}^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} u_{it} \right)$$

where $\mathbf{Q}_{xx,T}^{-1} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \right)^{-1}$. Next, consider the set of time-specific regressions:

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

$$v_{it} = \mathbf{x}'_{it}\boldsymbol{\lambda}_i + u_{it}$$

The time-specific OLS regressions yield:

$$(18) \quad \hat{\boldsymbol{\beta}}_t = \boldsymbol{\beta} + \boldsymbol{\theta}_t + \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \boldsymbol{\lambda}_i \right) + \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} u_{it} \right)$$

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

Given the pooled, unit-specific and time-specific regression results, we construct a preliminary estimate of $\boldsymbol{\beta}_{it}$ as follows:

$$(19) \quad \hat{\boldsymbol{\beta}}_{it}^{Prel} = \hat{\boldsymbol{\beta}}_i + \hat{\boldsymbol{\beta}}_t - \hat{\boldsymbol{\beta}}$$

This preliminary estimate is biased, but we can understand the nature of the bias by substituting (15), (17), and (18) into (19) to obtain the expression:

$$\begin{aligned}
\hat{\beta}_{it}^{Prel} &= \hat{\beta}_i + \hat{\beta}_t - \hat{\beta} = \\
&\beta + \lambda_i + \mathbf{Q}_{xx,T}^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \theta_t \right) + \mathbf{Q}_{xx,T}^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} u_{it} \right) + \\
(20) \quad &\beta + \theta_t + \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \lambda_i \right) + \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} u_{it} \right) \\
&- \beta - \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \lambda_i \right) \\
&- \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \theta_t \right) - \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it} u_{it} \right)
\end{aligned}$$

This can be written more compactly as:

$$(21) \quad \hat{\beta}_{it}^{Prel} = \beta + \lambda_i + \theta_t + (\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT}) + (\mathbf{Q}_{xu,N} + \mathbf{Q}_{xu,T} - \mathbf{Q}_{xu,NT})$$

where $\mathbf{R}_N \equiv \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \lambda_i \right)$, $\mathbf{R}_{t,NT} \equiv \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \theta_t \right)$ and $\mathbf{Q}_{xu,N} \equiv \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} u_{it} \right)$, and \mathbf{R}_T , $\mathbf{R}_{i,NT}$, $\mathbf{Q}_{xu,T}$ and $\mathbf{Q}_{xu,NT}$ are defined similarly.

The expression in (21) can be decomposed into three parts: First, the true observation (i, t) level coefficients $\beta + \lambda_i + \theta_t$, second, the bias term $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$ that arises from correlation between the regressors and the heterogeneity (including the fixed effects in the intercept), and finally terms $(\mathbf{Q}_{xu,N} + \mathbf{Q}_{xu,T} - \mathbf{Q}_{xu,NT})$ involving the errors. The latter vanish asymptotically given our assumptions on the \mathbf{x}_{it} and u_{it} .

Crucially, the bias term $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$ can be calculated to arbitrary accuracy and eliminated from (21), leaving a consistent estimator of β_{it} .¹⁸ We now explain the procedure: MO-OLS uses $\hat{\beta}_i$ as a first stage approximation for λ_i in $(\mathbf{R}_N - \mathbf{R}_{i,NT})$ to form $\hat{\mathbf{R}}_N$ and $\hat{\mathbf{R}}_{i,NT}$, and also uses $\hat{\beta}_t$ as a first stage approximation for θ_t in $(\mathbf{R}_T - \mathbf{R}_{t,NT})$ to form $\hat{\mathbf{R}}_T$ and $\hat{\mathbf{R}}_{t,NT}$. Substituting the definitions of $\hat{\beta}_i$ and $\hat{\beta}_t$ given by equations (17)

¹⁸In the special case that the slope heterogeneity is independent of the regressors, the FE-OLS estimator of equation (8), obtained via a two-way within transformation followed by OLS estimation, gives a consistent estimator of the mean coefficient vector $\bar{\beta}$, while the preliminary estimate of β_{it} in (20) will give a consistent estimate of the observation-level coefficients. However, as we discuss in Appendix B, we would expect adaptation to generate correlation between the slope heterogeneity and the regressors (in particular, a positive correlation between KDD and β_{2it}). In that case FE-OLS will generally deliver inconsistent estimates of the mean coefficient vector $\bar{\beta}$.

and (18) into the bias term $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$ we obtain:

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

This expression is equal to the original bias $(\mathbf{R}_N - \mathbf{R}_{i,NT}) + (\mathbf{R}_T - \mathbf{R}_{t,NT})$, plus additional bias terms that relate to \mathbf{Q}_{xu} (which is $o_p(1)$ under these assumptions) as well as the slope heterogeneity. By subtracting (22) from (21) we eliminate the original bias from $\hat{\beta}_{it}^{Pre\ell}$, while introducing a new bias. Importantly, the new bias must be smaller in magnitude than the original bias. This is stated as Lemma 1 of the Mathematical Appendix, where we prove the result.

We can repeat this process, again using the $\hat{\mathbf{R}}_N$, $\hat{\mathbf{R}}_{i,NT}$, $\hat{\mathbf{R}}_T$, and $\hat{\mathbf{R}}_{t,NT}$ to approximate the new bias term in (22). As we show in Lemma 1, this in turn produces a new bias term that is even smaller in magnitude. Thus, this process can be iterated L times to render the bias arbitrarily small, forming the bias removed estimates:

$$\begin{aligned}
(23) \quad \hat{\beta}_{it} = & \hat{\beta}_i + \hat{\beta}_t - \hat{\beta} + \sum_{\ell=0}^L (-1)^{\ell+1} \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{1,\ell} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{2,\ell} \right. \\
& \left. - \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{1,\ell} + \mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{2,\ell}) \right)
\end{aligned}$$

where $\Gamma_{1,\ell} = \mathbf{Q}_{xx,T}^{-1} (\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{2,\ell-1})$ and $\Gamma_{2,\ell} = \mathbf{Q}_{xx,N}^{-1} (\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Gamma_{1,\ell-1})$ when $\ell > 0$, and where $\Gamma_{1,0} = \hat{\beta}_i$ and $\Gamma_{2,0} = \hat{\beta}_t$. This is a Cauchy sequence in ℓ , so a suitable L can be chosen by terminating the sequence once it converges to a desired tolerance. In practice, small values of L are usually adequate. Equation (23) is simple to construct, as it is a function of only the preliminary estimates $(\hat{\beta}_i, \hat{\beta}_t, \hat{\beta})$ and the covariates \mathbf{x} .

Theorem 1 states consistency of MO-OLS estimates of the observation-level coefficients β_{it} as L goes to infinity, and then N and T jointly go to infinity. As MO-OLS is numerically equivalent to ‘brute force’ OLS, the proof is relegated to the Mathematical Appendix.

Theorem 1: Consistency of $\hat{\beta}_{it}$

For the model in (13), with **A.1-A.4**, if $L \rightarrow \infty$ and subsequently $(N, T) \xrightarrow{j} \infty$, then:

$$(24) \quad \hat{\beta}_{it} - \beta_{it} \xrightarrow{p} 0$$

Proof: See Mathematical Appendix

Remark: Recall that **A.4** imposes that $\beta_{it} = \beta + \lambda_i + \theta_t$. As in OLS, it does not make sense to discuss consistent estimates of the separate components $(\beta, \lambda_i, \theta_t)$ as they are only identified up to location normalizations.

Given consistent estimates of $\beta_{it} = \beta + \lambda_i + \theta_t$, a researcher may study them ex-post as desired. In some cases, a researcher may be interested in the mean coefficient vector $\bar{\beta} = (\beta + E(\lambda_i) + E(\theta_t))$. Given consistent estimates of the observation-level coefficients, the mean coefficient vector can be estimated by taking a simple average:

$$(25) \quad \hat{\beta}_{MO} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\beta}_{it}$$

We refer to this as the Mean Observation OLS (MO-OLS) estimate, as it averages the observation-level coefficients. The following theorems provide results for consistency and asymptotic normality. The proofs are given in the Mathematical Appendix.

Theorem 2: Consistency of $\hat{\beta}_{MO}$

For the model in (13), with **A.1-A.4**, if $L \rightarrow \infty$ and subsequently $(N, T) \xrightarrow{j} \infty$, then:

$$(26) \quad \hat{\beta}_{MO} - \bar{\beta} \xrightarrow{p} 0$$

Proof: See Mathematical Appendix

Theorem 3: Asymptotic Normality of $\hat{\beta}_{MO}$

For the model in (13), with **A.1-A.4**, if $L \rightarrow \infty$ and subsequently $(N, T) \xrightarrow{j} \infty$ such that $N/T \rightarrow \chi$ and $0 < \chi < \infty$, then:

$$(27) \quad \sqrt{NT}(\hat{\beta}_{MO} - \bar{\beta}) \xrightarrow{d} N(0, \Sigma_{MO})$$

where $\Sigma_{MO} = Var(\lambda_i) + Var(\theta_t)$. The asymptotic variance can be consistently estimated

nonparametrically by:

$$(28) \quad \hat{\Sigma}_{MO} = \frac{1}{NT-1} \sum_{i=1}^N \sum_{t=1}^T ((\hat{\beta}_{it} - \hat{\beta}_{\bar{t}})(\hat{\beta}_{it} - \hat{\beta}_{\bar{t}})' + (\hat{\beta}_{it} - \hat{\beta}_{\bar{i}})(\hat{\beta}_{it} - \hat{\beta}_{\bar{i}})')$$

where $\hat{\beta}_{\bar{i}} = \frac{1}{T} \sum_{t=1}^T \hat{\beta}_{it}$ and $\hat{\beta}_{\bar{t}} = \frac{1}{N} \sum_{i=1}^N \hat{\beta}_{it}$.¹⁹

Proof: See Mathematical Appendix

Remark: MO-OLS extends the ‘mean group regression’ (MG-OLS) of Pesaran and Smith (1995) by providing a consistent estimate of the average effect $\bar{\beta}$ in the presence of time fixed effects. Time fixed effects would render the MG-OLS estimate inconsistent.

In this article, we are not primarily interested in the average effect of weather on crop yield. Rather, we are primarily interested in the observation-level coefficients $\hat{\beta}_{it}$. We examine the distribution of the $\hat{\beta}_{it}$ to determine if they are correlated with KDD_{it} , and to determine if there are trends in the mean of $\hat{\beta}_{it}$ over time. We use these patterns to quantify and understand historical adaptation.

In Appendix D we present Monte Carlo simulation results where we compare the performance of MO-OLS against several traditional panel data estimators in an environment with multidimensional slope heterogeneity that is correlated with the regressors. Our results show that MO-OLS is able to consistently and efficiently estimate the coefficients in this environment at reasonable sample sizes. Our results also reveal how poorly traditional panel data estimators can perform in this environment. Fixed effects and mean group estimators generate biased estimates of average effects, and depending on the underlying structure of the heterogeneity the bias can be very severe.

III. Data and Variable Construction

We use temperature and precipitation data from Schlenker and Roberts (2009). These data contain daily observations on maximum and minimum temperature, and precipitation, on a grid across the continental U.S., from 1950 to 2015.²⁰ We map the grid-based

¹⁹Restrictions on the relative rate of convergence of N and T are required due to the small sample time series bias $O(T^{-1})$, noted by Hurwicz (1950) in the case of weakly exogenous regressors (such as a lagged dependent variable), and to prevent an incidental parameter problem. Thus, the estimator is not appropriate for panels with small T .

²⁰Schlenker and Roberts subsequently extended the dataset to 2015 after the publication of their paper. They used daily weather observations from stations to construct the dataset. Temperature observations prior to 1950 are available, but it is harder to convert them to a national grid as fewer weather stations were operational.

data onto counties, weighting grid locations by the location of corn production in each county. Using the daily max/min temperatures, we approximate the hours each day that a crop is exposed to one-degree Celsius temperature intervals using a sinusoidal function:

$$(29) \quad DD_C = \begin{cases} 0 & \text{if } C > T_{max} \\ T_{avg} - C & \text{if } C < T_{min} \\ \frac{((T_{avg}-C)\cos^{-1}(S)+(T_{max}-T_{min})\sin(S)/2)}{\pi^{-1}} & \text{otherwise} \end{cases}$$

where C is the temperature in Celsius, T_{max} and T_{min} are the daily max/min temperature, $T_{avg} = \frac{T_{max}+T_{min}}{2}$ and $S = \cos^{-1}(\frac{2C-T_{max}-T_{min}}{T_{max}-T_{min}})$. This is consistent with the literature.

The threshold separating GDD from KDD is $29^\circ C$, while the minimum GDD threshold is set at $0^\circ C$, following the literature.²¹ Using results from (29), we calculate daily degree day values by forming $GDD_{id} = DD_0 - DD_{29}$ and $KDD_{id} = DD_{29}$ for each county i and day d . We aggregate to annual values of GDD_{it} and KDD_{it} by summing over the days of the growing season, which we assume is May 1st to September 30th (in line with the literature).²² Precipitation is measured as total inches of rain over the growing season.

Our county-level yield data is from the U.S. Department of Agriculture (USDA) National Agricultural Statistics Service. It covers the 1950 to 2015 period. Following the literature, we exclude counties west of the 100th Meridian.²³ We also exclude counties with less than 30 years of data. This gives $N = 2209$ corn growing counties with an average of 57 observations per county. 32% of counties have full coverage over the sample period. The 15% with less than 45 years of data are mostly found in the northeast and southern states.

IV. Results

A. Conventional Fixed Effects Models (FE-OLS)

Table 1 presents results from corn yield regressions that incorporate county and/or time fixed effects but do not allow for heterogeneous slope coefficients. As in Section II, we

²¹ 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 measured air temperature above the crop canopy.

²²When we let the growing season vary by county and over time it did not meaningfully change the results.

²³The 100th Meridian separates the Great Plains to the east from the semi-arid lands to the west. The western counties are much more reliant on irrigation. We tried including them but it did not meaningfully affect our results.

TABLE 1—FE-OLS ESTIMATES OF THE IMPACTS OF TEMPERATURE ON U.S. CORN YIELDS

Specification	(1)	(2)	(3)
GDD	0.0002 (0.0001)	0.0003 (0.0001)	0.0004 (0.0001)
KDD	-0.0052 (0.0006)	-0.0054 (0.0006)	-0.0158 (0.0023)
ln(KDD)*KDD - KDD			0.0022 (0.0005)
Precipitation	0.0007 (0.0002)	0.0010 (0.0002)	0.0007 (0.0002)
Precipitation ² (x 1000)	-0.0006 (0.0002)	-0.0007 (0.0001)	-0.0006 (0.0002)
Constant	2.6755 (0.2564)	2.3552 (0.2509)	2.1977 (0.2399)
Fixed Effects	Cty, Yr	Cty	Cty, Yr
Quad. Time Trend	N/A	State-specific	N/A
R-squared	0.82	0.82	0.83
Obs.	126,373	126,373	126,373

Notes: Results exclude counties west of the 100th Meridian. The sample period is 1950-2015, and $N = 2209$. Models (1)-(3) differ by type of fixed effects and whether the adaptation variable is included. Standard errors are reported in parentheses, and are clustered at the state level.

refer to these as FE-OLS models. The model in column (1) is similar to ones in the extant literature (see equation (8), Lobell et al. 2011, Burke and Emerick 2016). The estimated KDD coefficient is -0.0052 , implying an additional degree-day above 29°C leads to a decrease in overall corn yield of 0.52 percent. This effect is precisely estimated with standard errors that are clustered at the State level.²⁴ The model in column (2) restricts time effects to be quadratic trends at the State level (as in Schlenker and Roberts 2009, Roberts and Schlenker 2012). But the estimate of KDD sensitivity is not significantly different. Neither model (1) or (2) includes adaptation.

The model in column (3), which is motivated by the simple theory in Section I, includes as an additional regressor a nonlinear function of KDD designed to capture adaptation to harsh temperatures (see equation (9)). The added regressor is positive and highly

²⁴It is also possible to adopt spatial standard errors as in Conley (1999), where the correlation between county errors are assumed to decline with distance. Since doing this leads to less conservative standard errors we report the standard errors that are clustered by state in the results.

significant, implying that as KDD increases the negative marginal effect of KDD on crop yield gets smaller.²⁵ This may be due to both farmer and natural adaptations, and we argue that to predict future corn yield it is important to capture both.

According to model (3), the *average* marginal effect of KDD is -0.0082. This is significantly more negative than the average marginal effect in specifications (1)-(2), illustrating the bias in standard FE-OLS models that ignore adaptation (see Appendix B). However, while model (3) implies the average marginal effect of KDD on corn yield is 58% more negative than in model (1), it also implies the effect diminishes as KDD increases.

B. Heterogeneous Slope Models (MO-OLS)

In this section we use MO-OLS to estimate a model that allows for both county and time fixed effects in coefficients on temperature and precipitation, as in equations (11) and (12). In contrast to the FE-OLS model in (9), the heterogeneous slope model allows us to model adaptation 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 , which is informative about the nature of adaptation. Table 2 presents the results. It gives unweighted and weighted means of the estimated slope coefficients (using average crop acreage of each county as weights), as well as other moments.

The unweighted mean coefficient on KDD is -0.0096, implying one extra degree day of temperatures over 29°C causes a 0.96% reduction in crop yield. This effect is slightly larger than the average marginal effect of -0.0082 we obtained Table 1 column (3) when we modelled adaptation parametrically using the nonlinear KDD coefficient. Furthermore, our MO-OLS estimate of the mean KDD coefficient is about 80% more negative than we obtained using the conventional FE-OLS models in Table 1 columns (1)-(2). This illustrates the substantial bias in estimators that ignore slope heterogeneity.

The MO-OLS estimates imply substantial heterogeneity in the model coefficients. The standard deviation of the KDD coefficient is 0.0068, with a 90/10 percentile range of -0.0034 to -0.0161. Notice R^2 improves from .82 to .88 when slope heterogeneity is include in the model. Figure 2 plots the distribution of the KDD coefficients across counties and

²⁵Note that the coefficient on KDD itself becomes significantly more negative at -0.0158 when the nonlinear term is added. The results for the other variables do not change meaningfully from the first two specifications.

TABLE 2—MO-OLS ESTIMATES OF THE IMPACTS OF TEMPERATURE ON U.S. CORN YIELDS

	Mean	Weighted Mean	Median	Standard Deviation	10th Percentile	90th Percentile
GDD	0.0005 (0.0000)	0.0005	0.0005	0.0005	-0.0001	0.0011
KDD	-0.0096 (0.0003)	-0.0089	-0.0085	0.0068	-0.0161	-0.0034
Precipitation	0.0011 (0.0002)	0.0015	0.0010	0.0028	-0.0020	0.0044
Precipitation ² (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
R-squared	0.88					
Obs.	126,373					

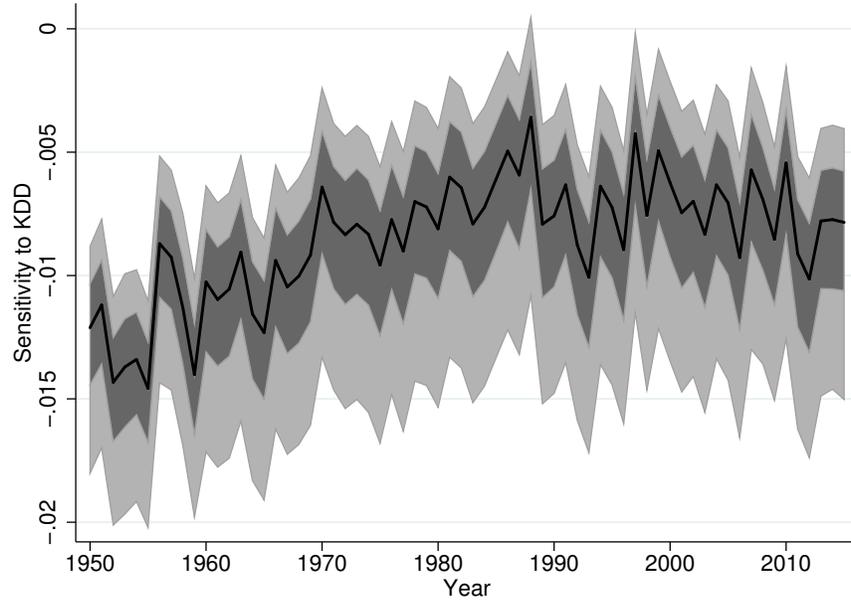
Notes: Results exclude counties west of the 100th Meridian. The sample period is 1950-2015, and $N = 2209$. Standard errors are reported in parentheses.

over time. Clearly, there is significant heterogeneity across counties. The 90/10 percentile range of KDD coefficients is over .010 units in each year. The 25/10 percentile range (i.e., the lower light grey area), is much wider than the 90/75 range, so the distribution has a fat left tail. Figure 2 also shows that the median KDD coefficient follows a clear trend over time. It 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, using the method of Andrews (1993). The results imply a positive trend in the KDD coefficient until 1988, implying reduced heat sensitivity of yield - which we view as evidence of adaptation. However, a significant break occurred in 1989, coinciding with an extreme drought in the Midwest in 1988-89. After that, the time trend on the KDD coefficient becomes small and insignificant. The median KDD coefficient in 2015 is similar to that in the 1970s, suggesting a lack of sustained progress in adaptation to high temperatures over the last four decades.

The stalling of adaptation may be explained by the results of Lobell et al. (2014), who find that while *average* corn yield increased from 1995 to 2010, sensitivity to droughts and high heat increased because of a trend toward higher sowing densities. This, in turn, may be due to an expansion of crop insurance, see Annan and Schlenker (2015). Our results do

FIGURE 2. DISTRIBUTION OF KDD SLOPE COEFFICIENTS ACROSS TIME AND COUNTIES FOR U.S. CORN



Note: The black line plots the median *KDD* coefficient from the MO-OLS model in Table 2. The dark (light) grey areas represent the 25th to 75th, and 10th to 90th percentiles, respectively.

differ from Lobell et al. (2014), as they find heat sensitivity increased in the latter period, while we only find that the trend toward reduced heat sensitivity stalled.

We emphasize, however, that Figure 2 and Table 3 provide clear evidence of adaptation of corn yield in the U.S. to heat between the 1950s and 1980s, due to the strong upward trend in the *KDD* coefficient over the first half of the sample period. This contradicts the conclusion of Schlenker and Roberts (2009) that there has been no significant adaptation. Notably, however, if we implement their simple test for adaptation, by splitting the sample in half (by time) and estimating the two-way fixed effects model of Table 1 column 1 for each subsample, we find the *KDD* coefficient decreases significantly from -0.0055 to -0.0046 , giving evidence of adaptation between the first and second half of the sample.²⁶

Next we consider the geographic pattern of corn's sensitivity to heat. This is mapped in Figure 3, which presents the average value of the *KDD* coefficient from the MO-OLS model

²⁶Schlenker and Roberts (2009) conclude there is no evidence of a changing relationship between yield and temperature in the second half of their sample period (1978-2005) relative to the first half (1950-1977). But Figures A11 and A12 of their supplementary appendix do show a significant decline in heat sensitivity for corn using their 'piecewise linear' model, which is similar to the two-way fixed effects model in our Table 1 column 1.

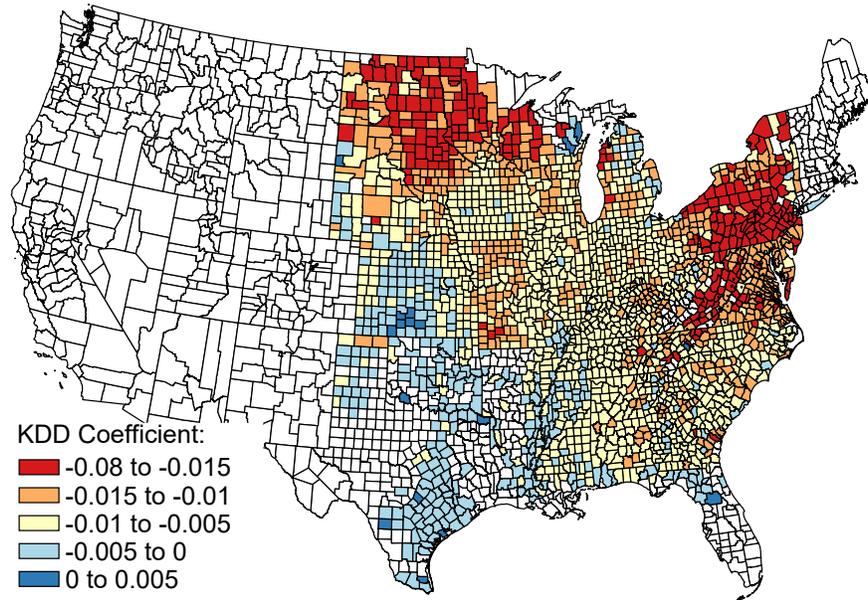
TABLE 3—ANALYZING TRENDS IN THE MEDIAN KDD COEFFICIENT FOR CORN OVER TIME

Regression Results	β	Std. Err.	β	Std. Err.
$t/10$	0.0008	0.0001	0.0020	0.0002
Constant	-0.0112	0.0005	-0.0133	0.0005
$(t/10)*d_{t>break}$			-0.0022	0.0004
$d_{t>break}$			0.0070	0.0020

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.

FIGURE 3. DISTRIBUTION OF KDD SLOPE COEFFICIENTS ACROSS COUNTIES FOR U.S. CORN



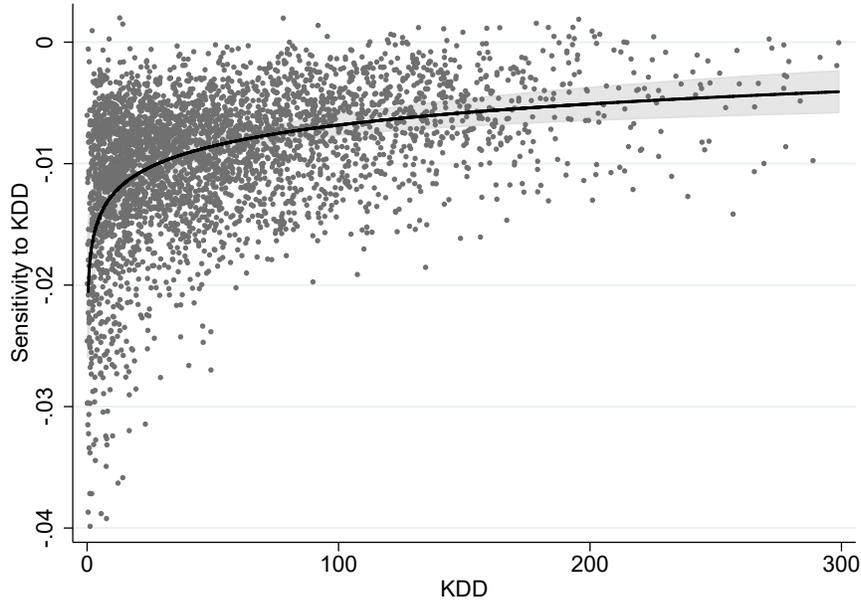
Note: This graph maps for each county of the United States in the estimation sample the average value of β_{2it} from the MO-OLS model (see Table 2).

for each county in the estimation sample. The highest category of sensitivity, shown in red, is common in North and South Dakota, Minnesota, and the northeastern region of the country. Most of the Corn Belt, along with much of the southeast, contains counties with

moderate sensitivities between -0.01 and -0.005 . The counties with low heat sensitivities are concentrated in Texas, Oklahoma, and Kansas. Thus, there is systematic variation in heat sensitivity across regions, as KDD sensitivity is typically greater in the cooler northern states. Adaptation is the most obvious explanation for this pattern.

The main prediction of the simple model of Section I is that KDD sensitivity should decrease as KDD increases, as farmers have more incentive to adapt. The correlation between $\hat{\beta}_{2it}$ and KDD_{it} is 0.43, supporting this prediction. Figure 4 presents a scatter plot of $\hat{\beta}_{2it}$ and KDD_{it} , as well as a regression of $\hat{\beta}_{2it}$ on $\log(KDD_{it})$, which is the best fitting curve to the approximately log-linear relationship. This is remarkably consistent with the simple parametric model of Section I, which also generates a log-linear relation.

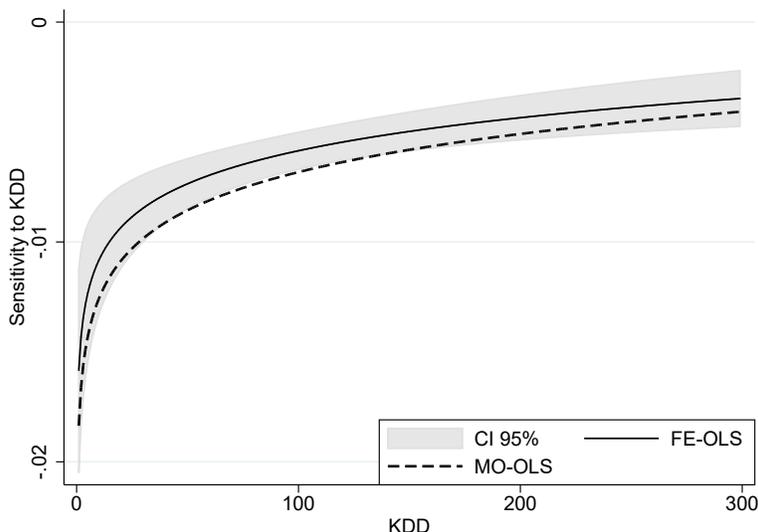
FIGURE 4. RELATIONSHIP BETWEEN $\hat{\beta}_{2it}$ AND KDD_{it} FOR U.S. CORN



Note: This graph is a scatter plot of a random 3% subsample of the coefficients on KDD_{it} from the MO-OLS model (see Table 2) against KDD_{it} itself. 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.

In Figure 5 we compare the relationship between $\hat{\beta}_{2it}$ and KDD_{it} that we obtain from the MO-OLS estimated fixed effects vs. the FE-OLS parametric model in Table 1 column (3). The MO-OLS curve implies a more negative marginal effect across all levels of KDD_{it} , but it still lies within the 95 percent confidence interval of the FE-OLS curve.

FIGURE 5. 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 the county/time specific coefficients estimated with MO-OLS (see Figure 4) and the FE-OLS regression results with adaptation, presented in Table 1 column 3.

Note that our fitted log-linear relation (based on both models) implies considerable scope for adaptation as KDD moves from 0 to 100, but at higher levels the curve flattens and adaptation is more modest. This has important implications for the potential for adaptation to mitigate future damage from climate change, if we assume the scope for future adaptation is similar to what we see historically. In particular, adaptation will become less effective at mitigating *marginal* damage under scenarios where typical KDD levels start to exceed about 100. We now turn to projecting climate change impacts.

V. Projecting the Effect of Climate Change on Crop Yield

A. Forecast Methodology

Here we project the effects of climate change on U.S. corn yields through to 2100. We present projections using FE-OLS and MO-OLS models with and without adaptation, and we assess the effectiveness of adaptation and emissions reductions in mitigating effects of climate change. We also assess the extent of disagreement in the projections that arises from differences between climate models vs. differences between econometric models.

To project future corn yields, we need predictions of temperature and precipitation for every corn growing county in the U.S. through to 2100. This requires predictions from a climate model, which further requires us to input a greenhouse gas emissions scenario. To assess sensitivity of our predictions to climate model assumptions, we use 19 different climate models and three emissions scenarios. We now describe the procedure in detail:

We utilize weather predictions from 19 general circulation models (GCMs), or simply “climate models,” from the Coupled Model Intercomparison Project v5 (CMIP5).²⁷ Climate models differ in how they represent a number of processes, such as chemical reactions, cloud formation, and vegetation growth. As we will see, they can generate rather different predictions of how the climate will respond to increased greenhouse gas emissions.

Each climate model requires, as an exogenous input, a path for the atmospheric concentration of greenhouse gases. In order to compare the output of climate models consistently, CMIP5 uses four representative concentration pathways: RCP26, RCP45, RCP60, and RCP85.²⁸ Each RCP embeds assumptions regarding the future trajectory of population growth, technological development and government policies.

We call RCP85 the “business as usual” scenario, as little is done to curb emissions, so atmospheric concentrations of greenhouse gases grow at present rates until 2100. We call RCP45 the “moderate” emissions reduction scenario. It corresponds to policies somewhat more ambitious than current government pledges. Under RCP45, greenhouse gas concentrations climb until 2100, but the growth rate significantly declines after 2060.

Finally, the RCP26 scenario represents “substantial” emissions reductions. It is consistent with the most ambitious targets under the Paris agreement. Greenhouse gas concentrations peak in 2040 and then slowly decline until 2100.²⁹ This would require significant immediate action, as emissions must start to decline substantially in the very near future.

²⁷The core of every GCM is a set of equations that describe the behavior of rotating spheres of continuous fluid that simulate the Earth’s atmosphere and oceans. The models we use 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. The CMIP protocol was introduced by The World Climate Research Program as part of the Working Group on Coupled Modelling. It ensures the outputs of GCMs are comparable, allowing scientists to analyze them systematically. The fifth version of CMIP is part of the broader effort for the IPCC Fifth Assessment Report.

²⁸The RCP scenarios correspond to different radiative forcing values in 2100. For instance, RCP26 results in a forcing value of $+2.6W/m^2$ above pre-industrial levels. We do not consider RCP60 due to its similarity to RCP45.

²⁹To be precise, atmospheric concentrations of CO₂ (and all other forcing agents converted to CO₂ equivalence) reach 1240 parts per million by 2100 under RCP85, 575 pp million under RCP45, and 435 ppm under RCP26.

The 19 climate models use greenhouse gas concentrations from the RCP scenarios to produce projections of daily min/max temperature and precipitation, across a 12x12km grid of the contiguous U.S. The grid level projections are converted to county level projections using an interpolation procedure known as “bias-correction and spatial disaggregation” (BCSD). We obtain BCSD output for our 19 climate models from the U.S. Geological Survey Geo Data Portal, who in turn rely on the Bureau of Reclamation (2013).

Finally, we convert the daily temperature and precipitation projections 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 (29).

Each climate model was run from 2006 to 2100, so the county-specific projection errors for weather in 2006-15 are observable. We center the projections so the average county-specific projection 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, and 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). Let $K\tilde{D}D_{itrm}$, $G\tilde{D}D_{itrm}$ and $P\tilde{R}E\tilde{C}_{itrm}$ denote the centered projections of the weather variables.

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

$$(30) \quad \hat{y}_{1itrm} = 2.6755 + \hat{c}_i + 1.3486 + 0.0002(G\tilde{D}D_{itrm}) - 0.0052(K\tilde{D}D_{itrm}) \\ + 0.0007(P\tilde{R}E\tilde{C}_{itrm}) - 0.00000091(P\tilde{R}E\tilde{C}_{itrm}^2)$$

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 1.3486, which is the 2006-15 mean of the time fixed effect.

Next, we use the FE-OLS specification in Table 1 column (3) that allows for adaptation:

$$(31) \quad \hat{y}_{2itrm} = 2.1977 + \hat{c}_i + 1.3574 + 0.0004(G\tilde{D}D_{itrm}) - 0.0158(K\tilde{D}D_{itrm}) \\ + 0.0022(\ln(K\tilde{D}D_{itrm}) * K\tilde{D}D_{itrm} - K\tilde{D}D_{itrm}) \\ + 0.0007(P\tilde{R}E\tilde{C}_{itrm}) - 0.00000063(P\tilde{R}E\tilde{C}_{itrm}^2)$$

All parameters are from Table 1, except \hat{c}_i and 1.3574 (the 2006-15 average time effect).

Our MO-OLS model can provide yield projections with or without adaptation. Define $\bar{\hat{\beta}}_{ki}$ as the mean of $\hat{\beta}_{kit}$ over the 2016-15 period for $k = 1, \dots, 4$, and similarly for $\bar{\hat{c}}_i$ and $\bar{\hat{\tau}}_t$. Then, a yield projection that does not allow for future adaptation can be obtained as:

$$(32) \quad \hat{y}_{3itrm} = \bar{\hat{c}}_i + \bar{\hat{\tau}}_t + \bar{\hat{\beta}}_{1i}G\tilde{D}D_{itrm} + \bar{\hat{\beta}}_{2i}K\tilde{D}D_{itrm} + \bar{\hat{\beta}}_{3i}P\tilde{R}E\tilde{C}_{itrm} + \bar{\hat{\beta}}_{4i}P\tilde{R}E\tilde{C}_{itrm}^2$$

where the marginal effect of each variable is fixed at the 2006-2015 mean for each county.³⁰

Alternatively, the MO-OLS model can allow for adaptation by setting the coefficient on $K\tilde{D}D_{itrm}$ equal to the predicted value from log-linear relationship shown in Figure 4.³¹ That is, we set $\hat{\beta}_{2it} = \alpha_1 + \alpha_2 \ln(K\tilde{D}D_{it})$ where $\hat{\alpha}_1 = -0.0183$ and $\hat{\alpha}_2 = 0.0025$:

$$(33) \quad \hat{y}_{4itrm} = \bar{\hat{c}}_i + \bar{\hat{\tau}}_t + \bar{\hat{\beta}}_{1i}G\tilde{D}D_{itrm} + (\log(K\tilde{D}D_{itrm}) * 0.0025 - 0.0183)K\tilde{D}D_{itrm} \\ + \bar{\hat{\beta}}_{3i}P\tilde{R}E\tilde{C}_{itrm} + \bar{\hat{\beta}}_{4i}P\tilde{R}E\tilde{C}_{itrm}^2$$

Finally, we aggregate our county-level yield projections to the national level using county average crop areas w_i , as in $\hat{y}_{trm} = \sum_{n=1}^N \hat{y}_{itrm}w_i / \sum_{n=1}^N w_i$.

B. Projections of Future Crop Yield

Here we present projections of corn yield from 2015 to 2100 using four econometric models, three RCP scenarios and 19 climate models. This gives $4 \cdot 19 \cdot 3 = 228$ scenarios. For each econometric model and emissions scenario r , we report the mean projection across all climate models, \hat{y}_{tr} , and 80% prediction intervals around this mean derived from the standard deviation of \hat{y}_{trm} across climate models m . This quantifies the extent of disagreement between climate models, as suggested by Auffhammer et al. (2013).

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.

³⁰Recall that our estimates capture both farmer and natural adaptation (i.e., inherent nonlinearity in the relation between yield and temperature). By shutting down both, we may exaggerate the impact of farmer adaptation.

³¹There is an asymmetry in (33) in that we allow the KDD coefficient to adapt over time, but we do not let other coefficients change over time. This is because we find no evidence that other parameters adapt to high temperatures. The correlations of $(\hat{c}_i + \hat{\tau}_t)$ and $\hat{\beta}_{1it}$ with KDD_{it} are only 0.02 and -0.04, respectively, in sharp contrast to the strong relationship between $\hat{\beta}_{2it}$ and KDD_{it} depicted in Figure 4. It is crucial to recall the distinction between adaptation *per se* (which we can project from historical data) and general technical progress. Predicting changes in τ_t and β_{1it} due to general technical progress is a more speculative exercise that we take up in Appendix G.

Appendix F describes the temperature predictions of the 19 climate models. The mean annual *KDD* level across all years/counties in the historical data is 41, while the projected means (across models) in 2050 are 69, 82 and 107 under the RCP26, 45 and 85 scenarios, respectively. Figures for 2100 are 70, 115 and 306 respectively. The models also predict increases in *GDD*, which may counteract some of the negative effects of higher *KDD*.

TABLE 4—EFFECTS OF CLIMATE CHANGE ON U.S. CORN YIELD (PCT CHANGE)

Year	Conventional FE-OLS	MO-OLS w/o future adaptation	FE-OLS with adaptation	MO-OLS with future adaptation
<u>RCP85</u>				
2030	-04 (-12, 04)	-06 (-16, 04)	-04 (-15, 06)	-06 (-14, 01)
2050	-19 (-34, -05)	-24 (-41, -05)	-21 (-36, 00)	-15 (-24, -05)
2080	-44 (-62, -26)	-51 (-72, -30)	-42 (-59, -26)	-26 (-37, -15)
2100	-62 (-81, -43)	-70 (-89, -51)	-57 (-73, -42)	-36 (-47, -25)
<u>RCP45</u>				
2030	-02 (-12, 09)	-04 (-16, 09)	-02 (-15, 11)	-04 (-13, 04)
2050	-11 (-24, 01)	-15 (-30, 00)	-12 (-26, 01)	-11 (-20, -01)
2080	-21 (-36, -06)	-27 (-44, -09)	-23 (-38, -07)	-16 (-25, -07)
2100	-21 (-38, -04)	-26 (-46, -05)	-22 (-40, 00)	-16 (-26, -06)
<u>RCP26</u>				
2030	-04 (-14, 05)	-07 (-17, 04)	-06 (-16, 05)	-07 (-13, 00)
2050	-07 (-16, 03)	-09 (-20, 01)	-08 (-19, 03)	-08 (-14, -02)
2080	-06 (-15, 04)	-09 (-20, 03)	-07 (-18, 05)	-08 (-15, -01)
2100	-07 (-19, 05)	-10 (-23, 03)	-08 (-21, 05)	-08 (-16, 00)

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 figures in brackets are the 80% (1.28 standard deviation) prediction intervals.

Table 4 summarizes our main results for corn yield. We report projections from four econometric models: the conventional FE-OLS model in (30), the FE-OLS model with adaptation in (31), and MO-OLS without and with adaptation in (32) and (33). We report 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. We present the ensemble average projection and the 80% (i.e., 1.28 standard deviation) prediction interval (in brackets).

First consider scenarios that ignore adaptation. As we see in Table 4, we predict catastrophic damage to corn yields in the pessimistic RCP85 scenario. The model ensemble average reduction in yield is 70% by 2100 according to the MO-OLS model, and 62% according to the conventional FE-OLS model. As we discussed earlier, the conventional FE-OLS model is likely to understate the impact of climate change because it ignores parameter heterogeneity. Yet we see that it still predicts a very severe impact.

With “moderate” emissions reductions (the RCP45 scenario) we observe smaller, but still significant, reductions in yield, with an ensemble average loss of 26% by 2100 (according to the MO-OLS model). And with “substantial” emissions reductions (as in RCP26, which requires significant immediate action on emissions) 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 no losses.

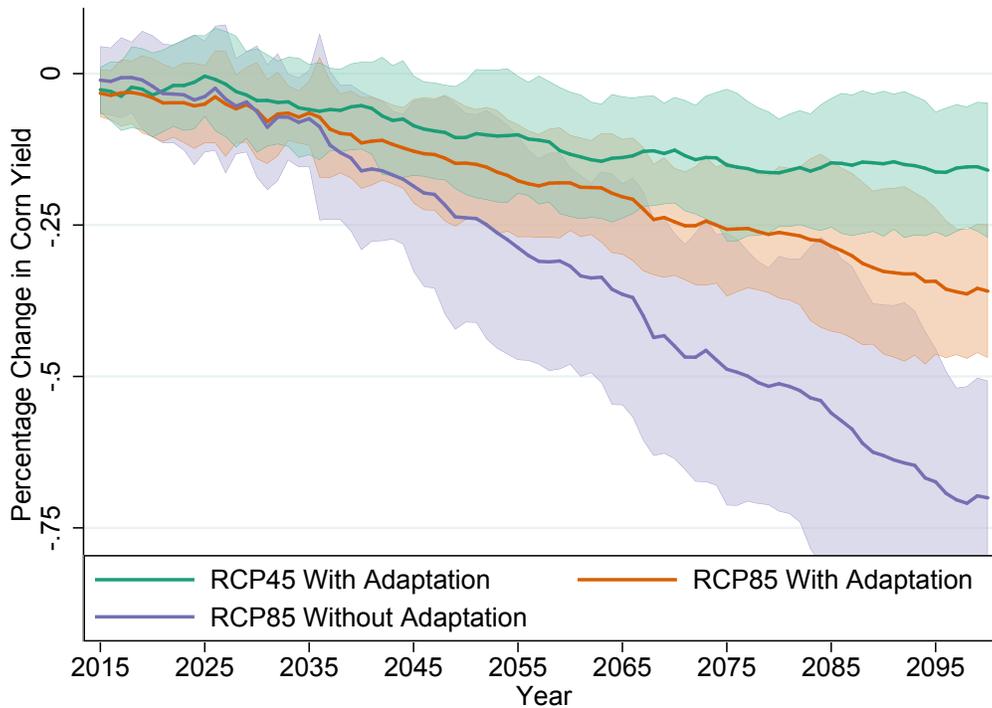
Turning to projections that incorporate adaptation to future climate change, recall that our parametric (FE-OLS) and fixed-effect (MO-OLS) models predict future adaptation based on the historical adaptation patterns shown in Figure 5. The results are reported in the right columns of Table 4. Both models predict adaptation will appreciably mitigate the damage from climate change in the RCP85 scenario. The FE-OLS model predicts adaptation will cause the decline in corn yield in 2100 to drop to 57%. The MO-OLS model implies adaptation is even more effective at mitigating damage: the reduction in crop yield in 2100 decreases from 70% to ‘only’ 36%.³²

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 26% 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

³²Thus, the MO-OLS estimates imply greater scope for adaptation to reduce yield losses than the FE-OLS estimates. This is in part because these two approaches model how the effect of KDD_{it} on y_{it} depends on KDD_{it} in different ways. In the MO-OLS approach in (33), the average effect of KDD_{it} is equal to the marginal effect, so adaptation shifts both equally. But in the FE-OLS approach in (31), the average effect of KDD_{it} on y_{it} is given by $\hat{\beta}_{20} + \hat{\beta}_{21}(\ln(KDD_{it}) - 1)$, which is always less than the marginal effect in (10). Thus, adaptation in the FE-OLS approach leads to a slower decrease in the average effect relative to the MO-OLS approach. We argue the MO-OLS specification is more intuitive, as it implies that adaptation alters the impact of all units of KDD_{it} on the crop, not just the additional units.

FIGURE 6. THE EFFECT OF CLIMATE CHANGE ON CORN YIELD BY SCENARIO (MO-OLS)



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

emission reductions are quite effective at reducing damage - e.g., for MO-OLS, compare the 36% mean drop in yield under RCP85 with adaptation with the 26% drop in yield achieved under RCP45 even without adaptation (falling to 16% when adaptation is included). Thus, we find that adaptation on a scale consistent with historical patterns cannot mitigate most of the damage from climate change - substantial emissions reductions are also required.

Figure 6 presents annual projections of corn yield using the MO-OLS model with and without adaptation. Under RCP85, most climate models predict the decline in yields by 2100 without adaptation will be catastrophic. Adaptation reigns in the most negative projections, and reduces the degree of disagreement between climate models (as is also clear in Table 4), yet the declines in yield remain severe. Under RCP45 with adaptation we predict a moderate mean decline in yield. But there is substantial disagreement between climate models, with the 80% prediction interval ranging from roughly 6% to 30%.

C. Cumulative Losses from 2020-2100

Next we examine the effectiveness of adaptation and emissions reductions at averting yield loss over the whole projection horizon. We calculate total loss of future yield as:

$$(34) \quad Loss_{r,ad} = M^{-1} \sum_{m=1}^M \sum_{t=2020}^{2100} (\hat{y}_{trm,ad} - \bar{y}_m)$$

where r is the RCP scenario, ad is the adaptation scenario, and \bar{y}_m is average crop yield from 2006 to 2015 according to the MO-OLS model in the worst-case baseline scenario (i.e., “business as usual” emissions, $r=RCP85$, and no adaptation, $ad=0$).³³ The share of damage averted under emissions and adaptation scenario ($RCPn, ad = k$) for $n = 26, 45, 85$ and $k = 0, 1$ is then defined as $1 - \frac{Loss_{RCPn,ad=k}}{Loss_{RCP85,ad=0}}$.

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

Model	RCP85+Adaptation	RCP45	RCP45+Adaptation	RCP26
MO-OLS	36 (29, 44)	55 (32, 79)	61 (49, 74)	76 (60, 91)
FE-OLS	15 (13, 17)		62 (38, 86)	

Note: Figures are the % reduction in damage relative to the RCP85 scenario with no adaptation using the MO-OLS model.

Table 5 presents the share of damage averted under four scenarios. According to the third column of Table 5, the MO-OLS model predicts a shift from the most pessimistic baseline of RCP85,0 (“business as usual,” no adaptation) to RCP45,0 (“moderate reductions,” no adaptation) will avert 55% of damage on average, with a 80% prediction interval of 32% to 79%. If we factor in adaptation (RCP45,1) the share of damage averted improves to 61%, and the 80% prediction interval narrows to 49% to 74%. Thus, our point projection is that moderate emissions reductions combined with adaptation will avert more than half the 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

³³We could of course discount future losses in (34), 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 our focus.

abatement is 76%, but the 80% prediction interval ranges from fully 91% to only 60%.

According to column (2) of Table 5, the MO-OLS model predicts adaptation alone can avert 36% of total damage to yield (on average),³⁴ while FE-OLS gives a much smaller figure of only 15%. This is a substantial difference. Yet both econometric approaches agree that adaptation alone (if consistent with historical patterns) cannot avert the majority of the severe damage to yield that is projected to occur with climate change. Significant emissions reductions are necessary to avert substantial reductions in future yields.³⁵

D. Distribution of Losses Across U.S. Counties

Here we examine the distribution of losses across counties due to future climate change. Figure 7 graphs the percentage of counties that experience losses as a function of time, RCP, and adaptation scenario. By 2100 the fraction that experience losses approaches one under RCP85 without adaptation. Allowing for adaptation reduces this fraction, but not substantially, as over 80 percent of counties still experience losses by 2100. While not definitive, these results cast doubt on the notion that corn production can be shifted to cooler corn-growing counties as a way to avoid significant damage to yields (at least not U.S. counties that are already producing corn).

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, but about 40% experience gains. This suggests that shifting production to cooler counties may be a more useful strategy if combined with substantial emissions reductions.

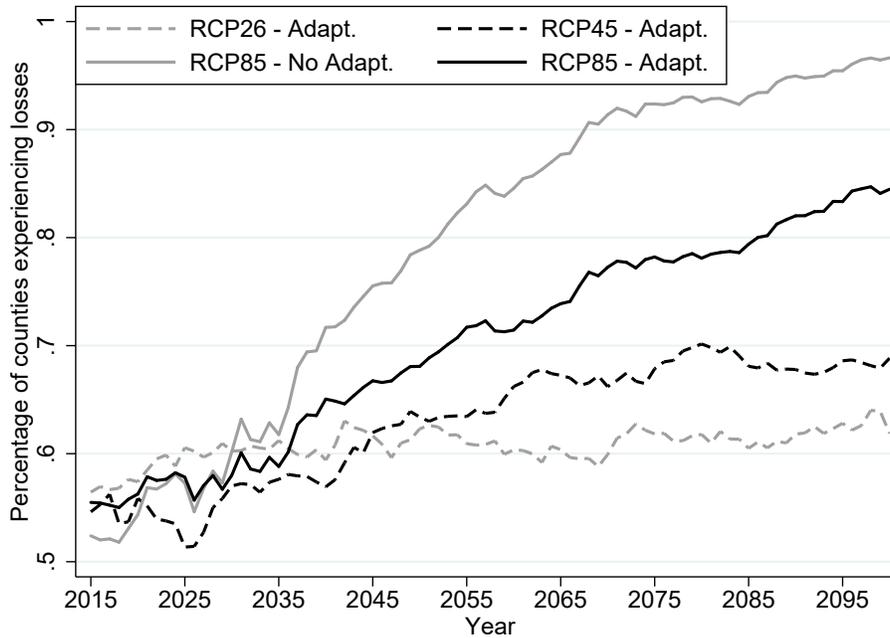
E. Extensions

So far, we have discussed projections holding agricultural technology fixed at current levels (aside from adaptation). In Appendix G we report yield projections that incorporate

³⁴As the MO-OLS model predicts a mean yield reduction of 70% in 2100 without adaptation and 36% 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.

³⁵Table 5 does not present FE-OLS results for share of losses averted with emissions reductions but no adaptation. This is because of our earlier finding that the conventional FE-OLS model (without adaptation) seriously understates damage from climate change. Thus we view share-of-loss-averted calculations based on the conventional FE-OLS model as unreliable. In contrast to FE-OLS, the MO-OLS model has the advantage that we can take the estimates from one model and compare results with and without allowing for adaptation.

FIGURE 7. 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.

projections of technical progress. Specifically, we use VARs to forecast the time effects in the MO-OLS model. We find that under the most optimistic technology forecast combined with RCP26 the mean projection of yield growth keeps pace with population growth. But there is still substantial uncertainty across climate models.

Finally, in Appendix H we present results for soybeans, the second largest U.S. crop. This is an interesting contrast, as soybeans are naturally less sensitive to heat than corn. But the production process for soybeans also exhibits less scope for adaptation (i.e., the *KDD* coefficient is less heterogeneous, and less highly correlated with *KDD*).³⁶ We predict that, without emissions reductions or adaptation, yields will drop by 34% to 74% by 2100. Adaptation is rather ineffective, reducing damage by only 0 to 22%. But emissions reductions are very effective - the moderate RCP45 scenario reduces damage by 36 to 85%.

³⁶This is consistent with a relative lack of knowledge on the root architecture and genome of the soybean through most of the sample period (e.g. see Alsajri et al. (2019) and Li et al. (2016)).

VI. Conclusion

We argue that adaptation to high temperatures generates spatial and temporal heterogeneity in the parameters of agricultural production functions. This heterogeneity enters slopes as well as intercepts, and it takes a fixed effects form, as it is correlated with temperature itself. Thus, we propose a new method, that we call “mean observation OLS” (or “MO-OLS”) that makes it feasible to estimate panel data models with unit and time fixed effects in both intercepts and slopes in very large panels.

We apply this method to estimate county-level corn yield equations for the U.S. for the 1950-2015 period. We find significant heterogeneity in the effect of temperature on yield, and, as expected, it diminishes as harsh temperatures are more common. This provides strong evidence for historical adaptation between warmer and cooler counties. We also find strong evidence of adaptation over time (i.e., declining heat sensitivity), but this has stalled since 1989. We show that models that do not allow for adaptation produce biased estimates of the effects of high temperatures, and fit the data significantly worse.

We use our econometric model to generate comprehensive county-level projections of corn yields through to 2100, based on future weather scenarios obtained from 19 climate models and three emission growth pathways. Our *econometric* approach is complimentary to work that uses weather predictions from climate models as input into *biological* models of plant growth in order to project future yields (see, e.g., Malcolm et al. 2012).

Our results imply several conclusions regarding the impact of climate change on corn yield: First, absent emissions reductions or adaptation, we predict very severe effects on yield. The average prediction across climate models is -70% by 2100. Second, we predict that almost all corn-growing counties will be adversely affected, implying there is little scope to avert damage by shifting production to cooler regions. Third, we predict that adaptation will avert 36% of damage (on average) in the no emissions reductions scenario. Thus adaptation is important, but can avert less than half of damage to yields.³⁷ Fourth, on a more optimistic note, we predict that adaptation combined with ‘moderate’ emissions reductions (i.e., similar to current government pledges) can avert 61% of damage to yields,

³⁷An important potential form of adaptation is irrigation. But results in Marshall et al. (2015) suggest that climate change itself will lead to water depletion that will inhibit irrigation in the latter half of this century.

rising to 76% under the more ambitious targets of the Paris agreement. Thus, *plausible* emissions reductions may still avert a large fraction of damage from climate change.

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

A striking feature of our results is the wide variability of projections across climate models. Indeed, we have focused 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 projections 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 projections of yield losses are moderate.

Acknowledgements: This research was supported by Australian Research Council grants FL110100247 and CE170100005. The authors acknowledge Maurice Bun, Gerdie Everaert, Seojeong Lee, Stéphane Bonhomme, Luke Hartigan, various seminar participants, the editor and three anonymous referees for constructive comments that improved the article.

REFERENCES

- Alsajri, F. A., B. Singh, C. Wijewardana, J. T. Irby, W. Gao, and K. R. Reddy (2019). Evaluating soybean cultivars for low- and high-temperature tolerance during the seedling growth stage. *Agronomy* 9(1).
- Andrews, D. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica* 61(4), 821–856.

- Annan, F. and W. Schlenker (2015, May). Federal crop insurance and the disincentive to adapt to extreme heat. *American Economic Review* 105(5), 262–66.
- Auerbach, A., Y. Gorodnichenko, and D. Murphy (2020). Local fiscal multipliers and fiscal spillovers in the USA. *IMF Economic Review* 68(1), 195–229.
- Auerbach, A. J. and Y. Gorodnichenko (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy* 4(2), 1–27.
- 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.
- Benati, L. and T. A. Lubik (2014). The time-varying beveridge curve. In *Advances in Non-linear Economic Modeling*, pp. 167–204. Springer.
- 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, January). Adaptation of US maize to temperature variations. *Nature Climate Change* 3, 68–72.
- Conley, T. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics* 92(1), 1 – 45.
- 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.
- Deschênes, 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.
- Frisch, R. and F. Waugh (1933). Partial time regressions as compared with individual trends. *Econometrica* 1(4), 387–401.
- Hsiao, C. (1975). Some estimation methods for a random coefficient model. *Econometrica* 43(2), 305–325.

- Hsiao, C. and M. H. Pesaran (2008). Random coefficient models. In *The Econometrics of Panel Data*, pp. 185–213. Springer.
- Hurwicz, L. (1950). Least squares bias in time series. In T. Koopman (Ed.), *Statistical Inference in Dynamic Economic Models*, pp. 365–383. New York: Wiley.
- Johnson, P. and C. Papageorgiou (2020, March). What remains of cross-country convergence? *Journal of Economic Literature* 58(1), 129–75.
- Keane, M. and T. Neal (2020). Comparing deep neural network and econometric approaches to predicting the impact of climate change on agricultural yield. *UNSW Research Working Paper 2020-02*. Available online at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3521260.
- Li, M.-W., D. Xin, Y. Gao, K.-P. Li, K. Fan, N. B. Muoz, W.-S. Yung, and H.-M. Lam (2016, 09). Using genomic information to improve soybean adaptability to climate change. *Journal of Experimental Botany* 68(8), 1823–1834.
- Liang, X.-Z., Y. Wu, R. G. Chambers, D. L. Schmoldt, W. Gao, C. Liu, Y.-A. Liu, C. Sun, and J. A. Kennedy (2017). Determining climate effects on US total agricultural productivity. *Proceedings of the National Academy of Sciences* 114(12), E2285–E2292.
- 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. 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. 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.

- Malcolm, S., E. Marshall, M. Aillery, P. Heisey, M. Livingston, and K. Day-Rubenstein (2012, July). Agricultural adaptation to a changing climate: Economic and environmental implications vary by U.S. region. Technical Report Economic Research Report 136, United States Department of Agriculture.
- Marshall, E., M. Aillery, S. Malcolm, and R. Williams (2015, November). Climate change, water scarcity, and adaptation in the U.S. fieldcrop sector. Technical Report Economic Research Report 201, United States Department of Agriculture.
- Ortiz-Bobea, A., E. Knippenberg, and R. G. Chambers (2018). Growing climatic sensitivity of U.S. agriculture linked to technological change and regional specialization. *Science Advances* 4(12).
- Pesaran, M. H. and R. Smith (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics* 68, 79–113.
- Porter, J., L. Xie, A. Challinor, K. Cochran, S. Howden, M. Iqbal, D. Lobell, and M. Travasso (2014). Food security and food production systems. In C. e. a. Field (Ed.), *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. pp. 485–533. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.
- 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.
- Roberts, M. and W. Schlenker (2012). *Is Agricultural Production Becoming More or Less Sensitive to Extreme Heat? Evidence from U.S. Corn and Soybean Yields*. University of Chicago Press.
- Schlenker, W. and M. Roberts (2009). Nonlinear temperature effects indicate severe dam-

ages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences* 106(37), 15594–15598.

Tannura, M. A., S. H. Irwin, and D. L. Good (2008, February). Weather, technology, and corn and soybean yields in the U.S. corn belt. Technical Report 01, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign.

United Nations Population Division (2017). 2017 revision of world population prospects. Can be found online at <https://esa.un.org/unpd/wpp/>.

Wallace, H. A. (1920). Mathematical inquiry into the effect of weather on corn yield in the eight corn belt states. *Monthly Weather Review* 48, 439–446.

Wang, J. C., S. H. Holan, B. Nandram, W. Barboza, C. Toto, and E. Anderson (2012). A bayesian approach to estimating agricultural yield based on multiple repeated surveys. *Journal of Agricultural, Biological, and Environmental Statistics* 17(1), 84–106.

Westcott, P. C. and M. Jewison (2013, July). Weather effects on expected corn and soybean yields. Technical Report FDS-13g-01, U.S. Dept. of Agriculture. Available at: https://www.ers.usda.gov/webdocs/publications/36651/39297_fds-13g-01.pdf?v=0.

Mathematical Appendix

Lemma 1

Consider a $M \times N$ square matrix \mathbf{B} and a $M \times 1$ column vector $\boldsymbol{\omega}$:

$$\mathbf{B} = \begin{pmatrix} b_{1,1} & b_{1,2} & \cdots & b_{1,n} \\ b_{2,1} & b_{2,2} & \cdots & b_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{m,1} & b_{m,2} & \cdots & b_{m,n} \end{pmatrix}, \quad \boldsymbol{\omega} = \begin{pmatrix} \omega_1 \\ \omega_2 \\ \vdots \\ \omega_m \end{pmatrix}$$

where $b_{m,n} > 0 \forall m$ and n . Then the $M \times 1$ vector sequence:

$$a_\ell = \begin{cases} \left(\frac{1}{N} \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{N} \sum_{n=1}^N b_{m,n} a_{\ell-1} & \text{if } \ell \text{ is odd} \\ \left(\frac{1}{M} \sum_{m=1}^M b_{m,n} \right)^{-1} \frac{1}{M} \sum_{m=1}^M b_{m,n} a_{\ell-1} & \text{if } \ell \text{ is even} \end{cases}$$

where $a_0 = \left(\frac{1}{M} \sum_{m=1}^M b_{m,n} \right)^{-1} \frac{1}{M} \sum_{m=1}^M b_{m,n} \omega_m$ is a convergent sequence that has the following pointwise limit in ℓ :

$$\lim_{\ell \rightarrow \infty} (a_\ell) = \bar{\omega} = \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \omega_m \text{ pointwise}$$

Proof of Lemma 1

a_0 represents an average of $\boldsymbol{\omega}$ over m for each n which is weighted by \mathbf{B} :

$$a_0 = [\bar{\omega}_0^{n=1}, \bar{\omega}_0^{n=2}, \dots, \bar{\omega}_0^{n=N}]$$

where $\bar{\omega}_0^n = \left(\frac{1}{M} \sum_{m=1}^M b_{m,n} \right)^{-1} \frac{1}{M} \sum_{m=1}^M b_{m,n} \omega_m$. Each element of a_1 , in turn, represents a weighted average of all the elements in a_0 over n for each m :

$$a_1 = [\bar{\omega}_1^{m=1}, \bar{\omega}_1^{m=2}, \dots, \bar{\omega}_1^{m=M}]$$

where $\bar{\omega}_1^m = \left(\frac{1}{N} \sum_{n=1}^N b_{m,n} \right)^{-1} [b_{m,n=1} \bar{\omega}_0^{n=1} + b_{m,n=2} \bar{\omega}_0^{n=2} + \dots + b_{m,n=N} \bar{\omega}_0^{n=N}]$.

Since $b_{m,n} > 0 \forall m$ and n , it follows that $\inf\{a_0\} \leq \inf\{a_1\}$ and $\sup\{a_0\} \geq \sup\{a_1\}$. If \exists an i, j pair such that $\bar{\omega}_0^{n=i} \neq \bar{\omega}_0^{n=j}$ and $i \neq j$, then it follows that $\inf\{a_0\} < \inf\{a_1\}$ and

$\sup\{a_0\} > \sup\{a_1\}$. Only if $\bar{\omega}_0^{n=i} = \bar{\omega}_0^{n=j} \forall i, j$ will $\inf\{a_0\} = \inf\{a_1\}$ and $\sup\{a_0\} = \sup\{a_1\}$. The same argument applies for a_2 , which is a weighted average of all elements of a_1 over n for each m , and indeed all subsequent values of ℓ in a_ℓ .

Thus, for every positive real number $\epsilon > 0$ there is a positive integer K such that for all positive integers $i, j > K$, the distance $d(a_i, a_j) < \epsilon$ (i.e. the sequence is convergent).

To demonstrate that $\lim_{\ell \rightarrow \infty} (a_\ell) = \bar{\omega}$ pointwise, first note that:

$$\sup\{a_0\} \geq \bar{\omega} \geq \inf\{a_0\}$$

as it is impossible for:

$$\left(\frac{1}{M} \sum_{m=1}^M b_{m,n} \right)^{-1} \frac{1}{M} \sum_{m=1}^M b_{m,n} \omega_m > \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \omega_m \quad \forall n$$

or:

$$\left(\frac{1}{M} \sum_{m=1}^M b_{m,n} \right)^{-1} \frac{1}{M} \sum_{m=1}^M b_{m,n} \omega_m < \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \omega_m \quad \forall n$$

when $b_{m,n} > 0 \forall m$ and n .

Then:

$$\sup\{a_1\} \geq \bar{\omega} \geq \inf\{a_1\}$$

since it is impossible for:

$$\left(\frac{1}{N} \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{N} \sum_{n=1}^N b_{m,n} \bar{\omega}_0^n > \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \bar{\omega}_0^n \quad \forall m$$

or:

$$\left(\frac{1}{N} \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{N} \sum_{n=1}^N b_{m,n} \bar{\omega}_0^n < \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \bar{\omega}_0^n \quad \forall m$$

when $b_{m,n} > 0 \forall m$ and n , and:

$$\begin{aligned} & \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \bar{\omega}_0^n = \\ & \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \left[\left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \omega_m \right] = \\ & \left(\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \right)^{-1} \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N b_{m,n} \omega_m = \bar{\omega} \end{aligned}$$

The same argument can be applied to $a_\ell \forall \ell > 0$, so that $\sup\{a_\ell\} \geq \bar{\omega} \geq \inf\{a_\ell\} \forall \ell$. Therefore, a_ℓ is a convergent sequence of vectors that always contains within it $\bar{\omega}$, which is accordingly the pointwise limit of the sequence.

Proof of Theorem 1

Given (15), (17), (18), and (23) then:

(35)

$$\begin{aligned} \hat{\beta}_{it} - \beta_{it} &= (\mathbf{Q}_{xu,N} + \mathbf{Q}_{xu,T} - \mathbf{Q}_{xu,NT}) + (-1)^L \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,L} \right. \\ &+ \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,L} - \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,L} + \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,L}) \left. \right) \\ &+ \sum_{\ell=0}^L (-1)^{\ell+1} \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell} - \right. \\ &\left. \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell}) \right) \end{aligned}$$

where $\Theta_{1,\ell} = \mathbf{Q}_{xx,T}^{-1} (\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,\ell-1})$ and $\Theta_{2,\ell} = \mathbf{Q}_{xx,N}^{-1} (\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,\ell-1})$ for $\ell > 0$, $\Lambda_{1,\ell} = \mathbf{Q}_{xx,T}^{-1} (\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell-1})$ and $\Lambda_{2,\ell} = \mathbf{Q}_{xx,N}^{-1} (\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell-1})$ for $\ell > 0$, $\Theta_{1,0} = \lambda_i$, $\Theta_{2,0} = \theta_t$, $\Lambda_{1,0} = \mathbf{Q}_{xu,N}$, and finally $\Lambda_{2,0} = \mathbf{Q}_{xu,T}$. First, using Lemma

1, (22), and (23) then:

$$(36) \quad \lim_{L \rightarrow \infty} \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,L} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,L} \right. \\ \left. - \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,L} + \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,L}) \right) = 0$$

To see this, exchange $b_{m,n}$ and ω_m in Lemma 1 for $\mathbf{x}_{it} \mathbf{x}'_{it}$ and either $\boldsymbol{\lambda}_i$ or $\boldsymbol{\theta}_t$. Since Lemma 1 showed the sequence a_ℓ converges pointwise to $\bar{\omega}$ in ℓ , then also the vector sequence:

$$q_\ell = \mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,\ell} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,\ell}$$

must converge to $\mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{1,L} + \mathbf{x}_{it} \mathbf{x}'_{it} \Theta_{2,L})$ in ℓ which gives us the result in (36).

Furthermore, given the Weak Law of Large Numbers, the Continuous Mapping Theorem, and assumptions **A.1-A.4**, as $N \rightarrow \infty$:

$$(37) \quad \mathbf{Q}_{xx,N}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} u_{it} \right) \xrightarrow{p} E(\mathbf{Q}_{xx,N}^{-1}) E(\mathbf{x}_{it} u_{it}) = E(\mathbf{Q}_{xx,N}^{-1}) 0 = 0$$

Furthermore, as $T \rightarrow \infty$:

$$(38) \quad \mathbf{Q}_{xx,T}^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} u_{it} \right) \xrightarrow{p} E(\mathbf{Q}_{xx,T}^{-1}) E(\mathbf{x}_{it} u_{it}) = E(\mathbf{Q}_{xx,T}^{-1}) 0 = 0$$

and lastly as $(N, T) \xrightarrow{j} \infty$:

$$(39) \quad \mathbf{Q}_{xx,NT}^{-1} \left(\frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N \mathbf{x}_{it} u_{it} \right) \xrightarrow{p} E(\mathbf{Q}_{xx,NT}^{-1}) E(\mathbf{x}_{it} u_{it}) = E(\mathbf{Q}_{xx,NT}^{-1}) 0 = 0$$

Given (37) - (39) and the Continuous Mapping Theorem then:

$$(40) \quad \sum_{\ell=0}^L (-1)^{\ell+1} \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell} - \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell}) \right) \xrightarrow{p} 0$$

Therefore, as required for Theorem 1:

$$(41) \quad \hat{\boldsymbol{\beta}}_{it} - \boldsymbol{\beta}_{it} \xrightarrow{p} 0$$

Proof of Theorem 2

Since $\bar{\boldsymbol{\beta}} = \boldsymbol{\beta} + E(\boldsymbol{\lambda}_i) + E(\boldsymbol{\theta}_t) = E(\boldsymbol{\beta}_{it})$, $\hat{\boldsymbol{\beta}}_{MO} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\boldsymbol{\beta}}_{it}$, and the result from Theorem 1 that $\hat{\boldsymbol{\beta}}_{it} - \boldsymbol{\beta}_{it} \xrightarrow{p} 0$ when $L \rightarrow \infty$ and then $(N, T) \xrightarrow{j} \infty$, the Weak Law of Large Numbers shows that:

$$(42) \quad \hat{\boldsymbol{\beta}}_{MO} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\boldsymbol{\beta}}_{it} \xrightarrow{p} E(\boldsymbol{\beta}_{it})$$

which implies Theorem 2.

Proof of Theorem 3

Given (35), (36), (25), and $\bar{\boldsymbol{\beta}} = \boldsymbol{\beta} + E(\boldsymbol{\lambda}_i) + E(\boldsymbol{\theta}_t)$ when $L \rightarrow \infty$ then:

$$(43) \quad \sqrt{NT}(\hat{\boldsymbol{\beta}}_{MO} - \bar{\boldsymbol{\beta}}) = \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T ((\boldsymbol{\lambda}_i - E(\boldsymbol{\lambda}_i)) + (\boldsymbol{\theta}_t - E(\boldsymbol{\theta}_t))) + \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T (\boldsymbol{\Psi}_{it} + \boldsymbol{\Xi}_{it})$$

where $\boldsymbol{\Psi}_{it} = (\mathbf{Q}_{xu,N} + \mathbf{Q}_{xu,T} - \mathbf{Q}_{xu,NT})$ and furthermore $\boldsymbol{\Xi}_{it} = \sum_{\ell=0}^L (-1)^{\ell+1} \left(\mathbf{Q}_{xx,N}^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{Q}_{xx,T}^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell} - \mathbf{Q}_{xx,NT}^{-1} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{1,\ell} + \mathbf{x}_{it} \mathbf{x}'_{it} \Lambda_{2,\ell}) \right)$.

Consider now the asymptotics where $(N, T) \xrightarrow{j} \infty$, assumptions **A.1-A.4** and the Weak Law of Large Numbers implies that both $\boldsymbol{\Psi}_{it} \xrightarrow{p} 0$ and $\boldsymbol{\Xi}_{it} \xrightarrow{p} 0$ (as shown in Theorem 1).

Accordingly:

$$(44) \quad \sqrt{NT}(\hat{\boldsymbol{\beta}}_{MO} - \bar{\boldsymbol{\beta}}) \xrightarrow{d} N(0, \Sigma_{MO})$$

where

$$(45) \quad \Sigma_{MO} = \text{Var}(\boldsymbol{\lambda}_i) + \text{Var}(\boldsymbol{\theta}_t)$$

since $\boldsymbol{\lambda}_i$ and $\boldsymbol{\theta}_t$ are independent by definition.

Now consider the nonparameteric estimate of Σ_{MO} that was proposed in (28):

$$\hat{\Sigma}_{MO} = \frac{1}{NT-1} \sum_{i=1}^N \sum_{t=1}^T ((\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})' + (\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})')$$

Given (35) and (36) then:

$$(46) \quad \hat{\boldsymbol{\beta}}_{it} = \boldsymbol{\beta} + \boldsymbol{\lambda}_i + \boldsymbol{\theta}_t + \boldsymbol{\Psi}_{it} + \boldsymbol{\Xi}_{it},$$

and

$$(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}}) = \left(\boldsymbol{\lambda}_i - \frac{1}{N} \sum_{i=1}^N \boldsymbol{\lambda}_i \right) + \left(\boldsymbol{\Psi}_{it} - \frac{1}{T} \sum_{t=1}^T \boldsymbol{\Psi}_{it} \right) + \left(\boldsymbol{\Xi}_{it} - \frac{1}{T} \sum_{t=1}^T \boldsymbol{\Xi}_{it} \right) \xrightarrow{p} (\boldsymbol{\lambda}_i - E(\boldsymbol{\lambda}_i)),$$

and using a symmetric argument

$$(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}}) \xrightarrow{p} (\boldsymbol{\theta}_t - E(\boldsymbol{\theta}_t))$$

where $\hat{\boldsymbol{\beta}}_{\bar{i}} = \frac{1}{T} \sum_{t=1}^T \hat{\boldsymbol{\beta}}_{it}$ and $\hat{\boldsymbol{\beta}}_{\bar{i}} = \frac{1}{N} \sum_{i=1}^N \hat{\boldsymbol{\beta}}_{it}$. Therefore:

$$(47) \quad \frac{1}{NT-1} \sum_{i=1}^N \sum_{t=1}^T ((\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})' + (\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})(\hat{\boldsymbol{\beta}}_{it} - \hat{\boldsymbol{\beta}}_{\bar{i}})') \xrightarrow{p} \text{Var}(\boldsymbol{\lambda}_i) + \text{Var}(\boldsymbol{\theta}_t)$$

and $\hat{\Sigma}_{MO} \xrightarrow{p} \Sigma_{MO}$ as required.