Heterogenous Patterns of Crop Yield Growth Stagnation across U.S. Counties in the Next Decade

Noé J Nava,[1,*]

[1]Economic Research Service, U.S. Department of Agriculture, Kansas City, MO 64105, USA
[*] Corresponding author: Noé J Nava, Noe.Nava@usda.gov

November 2022

Abstract

Crop yield outlooks indicate that climate change can put an end to the U.S. agricultural productivity gains experienced over the last couple of decades. The dominant explanation for this trend is raising temperatures and the increased frequency of extreme weather events. Previous research has primarily provided agricultural outlooks for the whole country with forecasts for the middle of the century while short-term projections and the heterogeneity of climate impacts remains understudied. I develop a probabilistic model of the interaction of weather variations and crop yield growth to provide short-term corn and soybean projections for U.S. counties. Crop yield simulations reveal that (1) warming temperatures can potentially stagnate U.S. crop yield growth in 2032 with a probability equal or higher than 75%, and (2) the largest productivity losses are associated with counties in the Corn Belt. My findings stress the urgency of adopting short-term adaptation and mitigation strategies to cope with climate change in the U.S.

JEL codes: C11, C15, Q13, Q54
Keywords: Bayesian data analysis, climate change, crop yield growth, U.S. agriculture.

1. Introduction

The most recent Intergovernmental Panel of Climate Change (IPCC) report revamps the urgency of its message to adapt to the dire conditions of a warming planet. Climate scientists expect the global temperature to exceed 1.5 °C above pre-industrial levels within the next few decades, forcing governments to evaluate short- and long-run adaptation and mitigation measures. Agriculture is the most sensitive sector to changes in weather patterns and extreme weather events (Lobell *et al.*, 2013). To provide an outlook on agriculture, practitioners employ several statistical techniques that rely on climate forecasts such as the hedonic model of farmland prices introduced by Mendelsohn, Nordhaus, and Shaw (1994). As researchers' interest expand, variations of the canonical methodology have been proposed to include other agricultural outputs such as agricultural profits as a measure of producers' welfare (Deschênes and Greenstone, 2007) or crop yields that inform on food availability (Schlenker and Robers, 2009; Roberts, Schlenker, and Eyer, 2013). Most of these studies have primarily provided agricultural outlooks for the whole country with forecasts for the middle of the century while short-term projections and the heterogeneity of climate impacts remains understudied (a couple exceptions are the work of Ortiz-Bobea and Tack (2018) that combine short- and long-term projections of crop yields, and the work of Dall'Erba, Chen, and Nava (2021) that investigate the effect of drought conditions on agricultural profit for specific U.S. states). In this manuscript, I focus attention on short-term agricultural outlooks and provide U.S. county-level corn and soybeans yield projections for 2032 to determine the heterogeneity associated with the decline of agricultural productivity found in other studies and the extent to which crop yield growth will stagnate in the next decade.

The estimation of damages from global warming in the agricultural sector has its origins in the work of Mendelsohn, Nordhaus, and Shaw (1994). Practitioners first derive relationships that relate climate to the variable of interest and then estimate the parameters associated with the climate variables. Using their model and estimated parameters along with climate scenarios, practitioners perform post-estimation simulations for the value of their dependent variable conditional on future climate. Due to its simplicity, several studies report employing variations of Mendelsohn, Nordhaus, and Shaw (1994)'s empirical framework to estimate the impact of climate change in different countries, including Canada (Reinsborough, 2003), China (Wang *et al.*, 2009), some European countries (van Passel, Masseti, and Mendelsohn, 2016), and South Africa (Gbetibouo and Hassan, 2005). Despite its success to provide insights on the climate threats to the food supply, a Bayesian reformulation of the interaction between weather variations and crop yield growth can allow for flexible forecasting, such as short-term projections of U.S. counties, by incorporating model structure (Gelman, 2008; Gelman, Lee, and Guo, 2015; van de Schoot, *et al.,* 2021).

This manuscript develops a hierarchical Bayesian framework that relates weather variation to U.S. county-level crop yield growth. A probabilistic model of precipitation, growing degree days, and extreme degree days is parametrized to affect the location and spread parameters of a crop yield distribution. Prior and hyper-priors are imposed to account for the combined effect of weather variables on crop yield growth and the known direction of their effect on agricultural productivity. Then, a Beta-type likelihood distribution of crop yields binds together model parameters with yield and weather data. Using corn and soybean yield data in the U.S. along with historical weather observations, a Hamiltonian Monte Carlo Markov-Chains (H-MCMC) algorithm samples the posterior distributions of the effect of weather variables on crop yields. Posterior distributions are then employed with weather forecasts used in the most recent IPCC report to study crop yield trends in the U.S.

Corn and soybean yield simulation results for the next decade indicate that increasing temperatures in the U.S. have the potential to put an end to the upward trend in crop productivity documented in previous studies within the next decade (see, for instance, Huffman, and Evenson, 1992; Alston *et al.* 2010; Ball, Schimmelpfenning, and Wang, 2013; Liang *et al.,* 2017). I find that (1) warming temperatures can potentially stagnate U.S. crop yield growth in 2032 with a probability equal or higher than 75%, and (2) the largest productivity losses are associated with counties in the Corn Belt. To test the robustness of my conclusions, I adopt two crop-growth regimes that account for non-weather influences and serve as an upper- and lower-limit of my projections. Conclusions hold regardless of whether an upward crop growth trend is assumed, and they are manifested in changes in the location and spread of the crop yield potential distributions formed by individual county results. An advantage of my Bayesian formulation is that I can study county-level results to determine what drives the decline in crop yields at the national level. Individual results from Midwestern and Eastern U.S. counties for the next decade indicate that yield potential in Midwestern states will decrease by up to 11.2 bushels per acre for corn and by up to 1.9 bushels per acre for soybeans within the next decade. Non-midwestern counties, however, will experience mild increases in crop yield potential no greater than 5.6 bushels per acre for corn and 1.5 bushels per acre for soybeans.

My study provides two primary contributions and the first one being methodological. Other studies, particularly the work of Zhu, Goodwin, and Ghosh (2011), employ a Beta-type distribution whose parameters are not unit-specific since their interest is on analyzing temporal changes in risk or the work of Nelson and Preckel (1989) that develop estimators for a Beta-type log-likelihood with an upper bound to study crop yields. Instead, I develop a Beta-type log-likelihood that has the advantage of having a location and spread parameter (similarly to commonly used normal distributions) that can be further parametrized to study unit-level production inputs. Closer to my work but using a multivariate normal distribution, Ramsey, Tack, and Balota (2022) are interested in estimating projections that account for climate scenarios, taking advantage of both temporal and cross-sectional variation. My approach goes beyond Ramsey, Tack, and Balota (2022)'s parametric estimation in that it allows estimable parameters to vary by location. A final contribution of my study is providing evidence that climate change will stagnate crop growth in the next decade, putting additional emphasis on the need to have short-term adaptation and mitigation policies. Previous studies have stressed the need to shift policy efforts to the short run when dealing with climate change (Loë, Kreutzwiser, Moraru, 2001; Tubiello and Ronsenzweig, 2008; Kistner *et al.,* 2017; Vimic *et al.,* 2022).

The next section describes my empirical strategy to study crop yield potential. Section 3 explains how the dataset is constructed. Section 4 employs my formulation to data on corn and soybeans in Midwestern and Eastern U.S. counties and discusses the main results. Finally, Section 5 offers concluding remarks.

2. Empirical Strategy

When it comes to farmers' responses to weather conditions, I assume that farmers' output is a function of their responses to variations in weather. That is, $y_{it} \equiv \Psi(x_{it}, \gamma^*(x_{it}))$, where x_{it} is farmers' observed weather each year t in location i, and $\gamma^*(x_{it})$ is input employment as a response to observed weather. To study farmers production output, I propose a Beta-type probability distribution that builds upon the work of Nelson and Preckel (1989), and Cribari-Neto and Zeileis (2010):

$$B(\mu_{it}, \phi) = \frac{\Gamma(\phi)}{\Gamma(\mu_{it}\phi)\Gamma((1-\mu_{it})\phi)} \times \frac{y_{it}^{\mu_{it}\phi-1}(y_u - y_{it})^{(1-\mu_{it})\phi-1}}{y_u^{(\phi-1)}} \tag{1}$$

where μ_{it} is a parameter that controls the location of the distribution and ϕ controls the dispersion.[1] The $\Gamma(.)$ notation denotes a gamma probability distribution, and y_u is an upper limit for crop yield realizations.[2] Equation (1) is advantageous for several reasons. First, (1)'s range is $(0, y_u)$ such that simulation draws, and parametric interpretations are restricted to be non-negative and less than y_u. Because of its flat tails (e.g., no probability density) outside $(0, y_u)$, tail tests from (1) are, in general, more accurate than those from a normal distribution whose tails extend indefinitely. Also, μ_{it} can be further parametrized to study the role of agricultural inputs on crop yields. For instance, the simplest model of the crop-response relationship is linear: $\mu_{it} = \beta^0 + \beta x_{it}$. The parameters β^0 and β are interpreted as the constant and the effect of the input x_{it} on crop yields, respectively. Unfortunately, a linear relationship has the flaw that μ_{it} can fall outside (1)'s domain, rendering (1) undefined. This can be addressed with a link function, $g(\mu_{it})$: $\mathbb{R}^1 \mapsto (0,1)$, that keeps (1) well-defined. A final consideration concerns ϕ, the parameter of dispersion in (1). I opt for a model that draws realizations for the level of dispersion as $\phi \sim \chi^2_{25}$, where the degrees of freedom as chosen to let the H-MCMC algorithm sample from an ample selection of draws.

To keep (1) stable and parsimonious, I choose to model μ_{it} first as being drawn from a normal distribution,

$$\mu_{it} = N(\boldsymbol{X_{it}\beta_i} + \Psi(t), \sigma), \tag{2}$$

where the $\boldsymbol{\beta_i}$ is a vector of estimable parameters associated with the weather variables in $\boldsymbol{X_{it}}$. Notice that the estimation is defined such that location-specific parameters are calibrated with data that vary by year. The choice of (2) as a varying-slopes model is based on my interest on estimating crop-yield projections at the county level. The term $\Psi(t)$ governs a crop-growth temporal trend

[1] Equation (1) is well-defined if $1 > \mu_t > 0$, and $\phi > 0$.
[2] The upper limit of equation (1) y_u, can be estimated within the estimation procedure discussed later as proposed by Nelson and Preckel (1989). I tested a linear parametrization of y_u formed by weather variables as discussed later. I find that doing this prevents my H-MCMC algorithm from reaching convergence. Instead, I opt for the informative prior $y_u = 300$. I test increasing the upper limit by multiples of 50 up to $y_u = 500$. I find that doing this does not affect my conclusions, suggesting that my upper limit assumption is not restrictive.

that accounts for non-weather influences, allowing me to impose crop-growth regimes in my simulations[3] (Ortiz-Bobea and Tack, 2018). I parametrize $\Psi(t)$ such that my projections follow the same trend as my historical data (upward trend) with additional simulations in which crop growth is assumed to be stagnated at 2017 levels (horizontal trend). The choice of these two crop-growth regimes corresponds to my interest of showing that my conclusions hold under opposite crop-growth regime assumptions, effectively creating an upper- and lower-limit for my results. The dispersion parameter in (2) is defined as $\sigma \sim \mathrm{N}(0,1)$ to allow the H-MCMC algorithm to sample an ample selection of draws.

Lastly, samples from (2) can render (1) undefined if they fall outside (1)'s domain, so I propose a simple link function,

$$\mu_{it}^{*} = \frac{\exp(\mu_{it})}{1+\exp(\mu_{it})}, \tag{3}$$

where μ_{it}^{*} is the transformation inputted into (1). Notice that the range of μ_{it}^{*} is not a limitation of the employment of equations (1) through (3) since $\mathrm{E}[y_{it}] = \mu_{it}^{*} y_u$, and $\mathrm{Var}[y_{it}] = \frac{\mu_{it}^{*}(1-\mu_{it}^{*})}{1+\phi} y_u^2$. Therefore, it is sufficient to scale predictions by y_u to obtain an interpretable effect on average crop yields and their variance by y_u^2.

Having established the relationships that connect my priors defined in (2) and (3) with my data (e.g., Equation (1)), I now define a set of hyper-priors that govern the rest of my model. Particularly, I am interested in a Hierarchical Bayesian estimation in which the vectors $\boldsymbol{\beta_i}$ for each location i are calibrated with temporal variation (unit-level variation across time) and cross-sectional variation. That is, I am interested in my H-MCMC algorithm sampling from the following proportion,

$$p(\boldsymbol{\beta_i}, \boldsymbol{\tau} | y_{it}, \boldsymbol{X_{it}}) \propto p(y_{it}, \boldsymbol{X_{it}} | \boldsymbol{\beta_i}) p(\boldsymbol{\beta_i}, \boldsymbol{\tau}), \tag{4}$$

where the additional vector, $\boldsymbol{\tau}$**,** contains additional parameters with their own distributions and includes the previously defined σ^2 and ϕ, $p(y_{it}, \boldsymbol{X_{it}} | \boldsymbol{\beta_i})$ is the likelihood function in (1) and

[3] I experimented with several temporal and spatial structures including autoregressive models, but complex autoregressive structures prevent proper convergence, limiting model reliance.

$p(\boldsymbol{\beta_i}, \tau) = p(\boldsymbol{\tau})p(\boldsymbol{\beta_i}|\boldsymbol{\tau})$ denotes my priors and hyper-priors model structure, being jointly determined.

I am interested in the role of precipitation (PPT), growing degree days (GDD), and extreme degree days (EDD). Thus, the vector $\boldsymbol{\beta'_i} = (\beta_i^{GDD}, \beta_i^{EDD}, \beta_i^{PPT})$. Then, I model farmers' responses to weather variations jointly with the following specification,

$$\begin{pmatrix} \beta_i^{GDD} \\ \beta_i^{EDD} \\ \beta_i^{PPT} \end{pmatrix} \sim \mathrm{N}(\mathbf{M}, \boldsymbol{\Sigma}) \,, \tag{5}$$

where $\mathbf{M}' = (0^+, 0^-, 0)$ is a vector of means, and the super-scripts indicate the H-MCMC algorithm to only sample positive values +, or negative values –.[4] In addition, the $\boldsymbol{\Sigma}$ is the variance-covariance matrix that regulates the variation of the distribution as well as the correlation across the parameters. Permitting a correlation between the weather impacts incorporates the combined effect of weather on crop yields. With no proper evidence on the combined effect of weather on crop yields, I allow my H-MCMC algorithm to select the estimates from a random structure. I decompose $\boldsymbol{\Sigma}$ into a correlation matrix, $\boldsymbol{\Omega}$, and a scaling factor, $\boldsymbol{\lambda}'$, such that $\boldsymbol{\Sigma} = \mathrm{diag}(\boldsymbol{\lambda})\boldsymbol{\Omega}\mathrm{diag}(\boldsymbol{\lambda})$ (Gelman, and Hill, 2007). The latter formulation allows the H-MCMC algorithm to sample on the distributions that govern the effect of weather:

$$\begin{aligned} &\boldsymbol{\lambda} \sim \mathrm{Cauchy}(0,1), \\ &\boldsymbol{\Omega} \sim \mathrm{LKJ}(\eta), \text{ where } \eta \in \{.5,1,2\}, \end{aligned} \tag{6}$$

where the half-Cauchy distribution is denoted as $\mathrm{Cauchy}(0,1)$, and its location and dispersion parameters are chosen to allow the H-MCMC algorithm to sample an ample selection of draws. The $\mathrm{LKJ}(\eta)$ refers to the Lewandowski-Kurowicka-Joe distribution, and its parameters are

[4] I test a non-informative vector of means, $\mathbf{M}' = (0,0,0)$, to allow the posterior distributions in $\boldsymbol{\beta_i}$ to have any value in the real line. I find that the assumption for GDD in (5) is restrictive since part of the posterior distribution of β_i^{GDD} shifts to the negative side of the real line. Nevertheless, I choose to keep an informative prior as illustrated in (5) since all models with such informative priors have MAPE measures between 31-35% when being tested against post-2017 data. MAPE measures from models with uninformative priors jump to 76% for corn and 55% for soybeans, indicating that estimations with uninformative priors are driven by unexplained variation.

calibrated such that they improve posterior predictions and control for the combine effect of the weather variables on crop yields[5] (Gabry *et al.*, 2019).

3. Data

The National Agricultural Statistics Service (NASS) collects and reports corn and soybeans data for the years between 1992-2020. Nevertheless, agricultural census years (ending in 2 and 7) report the most observations. To extend the scope of my study and allow for enough temporal and cross-sectional variation, I focus only on the agricultural census years.[6] I describe the complete crop yield data set in Figure 1 where the top panel describes corn, and the bottom panel describes soybeans. A salient insight is the upward trend of yields for both crops that has been studied extensively (Huffman, and Evenson, 1992; Alston *et al.,* 2010; Ball, Schimmelpfenning, and Wang, 2013; Liang *et al.,* 2017). Except for a few years, median yields fall closely to their upward trend lines. Econometric analyses interested in how farmers respond to weather variations would separately account for the unexpected deviations from the upward trend by either isolating the yearly variation with a dummy indicator or conditioning on additional observed variables (e.g., Midwestern farmers experienced exceptional drought conditions in 2012), but as I described in the previous section, my hierarchical Bayesian formulation of crop yields accounts for the *between* and *within* variation in the yearly data. A final insight from Figure 1 is the largest observed value each year. Corn yields have the largest values with three observations falling above 250. In (1)'s estimation, I chose $y_u = 300$ as the largest possible observed value in my analysis. While such restriction has no implications for estimation, it does limit the range of crop yield simulations

[5] My HMCMC algorithm is implemented in R on a 2.3 GHz Quad-Core Intel Core i7 MacBook Pro. Estimation time is above 18 hours for 8,000 iterations at a 0.5 warm up rate. Number of iterations and warm up rate are chosen to achieve convergence.
[6] My analysis includes 1,272 corn county producers and 1,113 soybean county producers for each year.

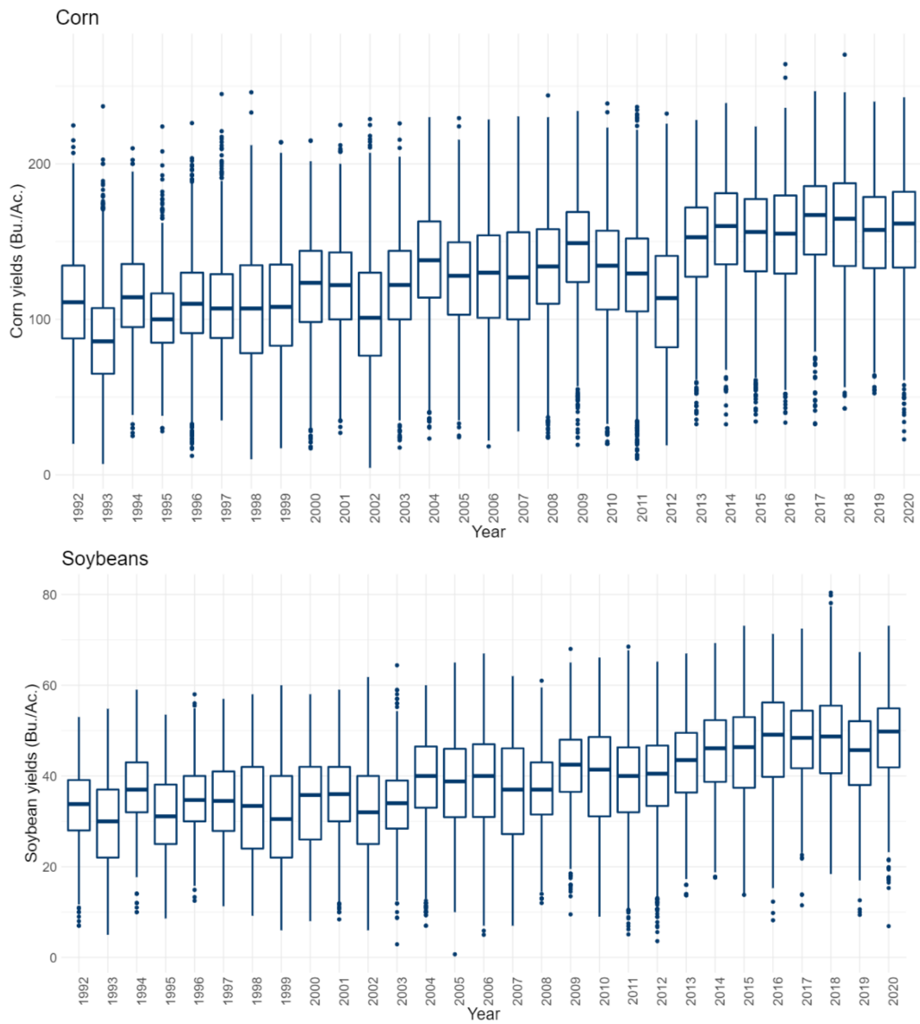

Figure 1: Corn and soybean yield trends in the U.S.: 1992-2020
Note: (a) Corn yields and (b) soybean yields. Blue and green lines show historical trends.

discussed in section 4 to fall below 300. I relaxed the latter assumption, re-ran my simulation experiments, and find no difference in results likely because $y_u = 300$ is large enough as shown in Figure 1.

Historical daily weather data come from the University of Oregon's Parameter-elevation Regression on Independent Slopes Model (PRISM) database. I process PRISM daily temperature (°C) into crop-specific measures of heat: GDD and EDD, where GDD and EDD have different upper-limit thresholds: 29°C for corn, 30°C for soybeans (Schlenker and Roberts, 2008). Most measures of heat exposure approximate the temperature distribution within a day by the midpoint between minimum and maximum daily temperature. To improve the predictive power of my

Bayesian application, a crop's heat exposure distribution within a day is approximated by a synodal function that only requires maximum and minimum daily temperature, following the work of Snyder (1985), Schlenker and Roberts (2009), and Tack, Barkley, and Nalley (2015).[7] PRISM precipitation is taken as is from the PRISM database where daily cumulative precipitation is measured in millimeters. Following standard practices in the literature, I focus on the corn and soybeans growing-season months, April-October, for Midwestern and Eastern U.S. counties (Burke and Emerick, 2016). A tradition of studies that follow closely the canonical work of Mendelsohn, Nordhaus, and Shaw (1994) include in their analyses all weather seasons of the year and observations from all U.S. counties. In contrast to these hedonic approaches that employ farmland prices as they account for current and future valuation of climate, my interest is to narrow down the effect of weather on corn and soybeans only. Thus, having weather variables from all seasons can add irrelevant variation to my estimations. Similarly, my focus on Midwestern and Eastern U.S. counties is intended to isolate the effect of growing degree days, extreme degree days, and precipitation from irrigation capacity (Schlenker, Hanemann, and Fisher, 2006).

My analysis uses future weather forecasts employed in the most recent IPCC report that assumes increases in global temperature of 1.5 – 3.0 ℃ above pre-industrial levels. I focus on the climate scenario categorized as "middle of the road" Shared Socioeconomic Pathway (SSP2-4.5) that accounts for medium challenges to mitigation and adoption of environmentally sustainable practices.[8] The National Aeronautics and Space Administration (NASA) collects and makes available all weather forecast simulations employed in the most recent IPCC report through its NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP) database (Thrasher *et al.,* 2012). NEX-GDDP is 18 TB, so its processing and manipulation requires extensive computational requirements. To manage computational requirements, I use Amazon's AWS storage and computing services (Polzehl, Papafitsors, and Tabelow, 2020).

In Figure 2, I plot historical and forecasted growing degree days (top figure), extreme degree days (middle figure), and precipitation (bottom figure) data. Each plotted density curve represents

[7] An ideal measure of crop heat exposure relies on mapping temperature variations within a day to the crops' heat exposure. This method is unpractical, and most climate models only report minimum and maximum daily temperature. My chosen methodology approximates daily temperature variation with $T(t) = \frac{T_{max}+T_{min}}{2} - \frac{T_{max}-T_{min}}{2}\cos(t)$ and $t \in [0, 2\pi]$, where T_{max} and T_{min} are maximum and minimum daily temperature, respectively. Then, heat exposure is approximated by evaluating the integral bounded by $T(t)$ and the crops' heat resistance levels. I thank Jesse Tack from the Kansas State University for introducing me to this heat exposure approximation and sharing his code.

[8] I obtain the average of my chosen experiments: CanESM5, TaiESM2, MIRO6, and ACCESS-CM2.

a year. I select gaps of 10 years for the sake of exposition and to account for the stationarity of climate over short periods of time. Values are logged (base 2) to ease comparison. A closer look to the tails and the peaks of the density can provide insights into the mechanism that govern forecasted results in the next section. Historical and forecasted weather data shows a pattern of raising temperatures towards the middle of the century whose effect has been considered in the literature (Lobell *et al.,* 2013). Raising temperatures affect crop growth through an increase of growing degree days (positive effect) and extreme degree days (negative effect). This pattern is observed in both panels. Raising temperatures increase the density of growing degree days on the right-hand side of the distribution with thicker right-tails and flatter left-tails. In turn, temperatures above the crop's heat resistance level increases the number of extreme degree days. The middle panel shows that the years of 2022 and 2032 have a lower peak on the right-hand side, suggesting that the temperature during the growing season in the U.S. is surpassing the crop's heat resistance level. While 2012 was an exceptionally dry year for the U.S. as shown by a density that curves towards the left with a low peak, precipitation patterns suggest that total rain will stay relatively constant. Tails do not vary in thickness across the decades, and the location of the peak is constant with mild variations of less than 10%.

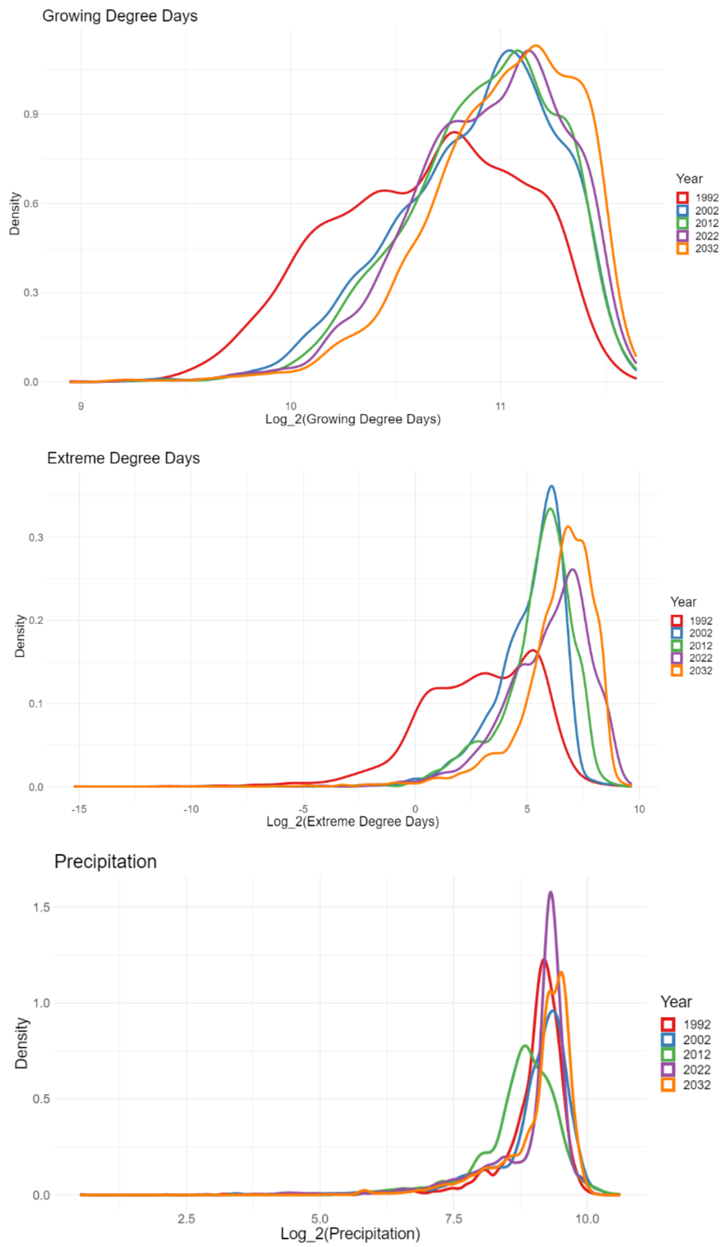

Figure 2: Historical and forecasted weather trends in Midwestern and Eastern U.S. counties: 1992-2032.

Note: Top figure is growing degree days. Middle figure is extreme degree days. Bottom figure is precipitation. Growing degree days and extreme degree days are calculated using 30°C.

4. Crop yield outlooks

Figures 3 and 4 describe simulation results for corn and soybeans, respectively. For each figure, the top panels present the crop yield probability distributions, aggregated across all county-observations, for the years of 2022, 2027, and 2032 compared to crop yield realizations for the years of 2012 and 2017 across the same counties. The bottom panels present the slope of the trends towards 2032 using 2017 values as the reference. The left figures illustrate the no growth regime while the right figures illustrate the growth regime as described in my empirical strategy section. To illustrate, comparing corn (in Figure 3) yield realizations in 2012 (in red) with corn yield distributions for the year of 2027 (in purple) in the no growth regime (left), my results suggest that the U.S. in 2027 can potentially reduce its corn production to 2012 levels. This can be observed in a bunching of counties on the left-hand side of the 2027 corn yield potential distribution, suggesting that a significant number of counties will experience low productivity because of higher

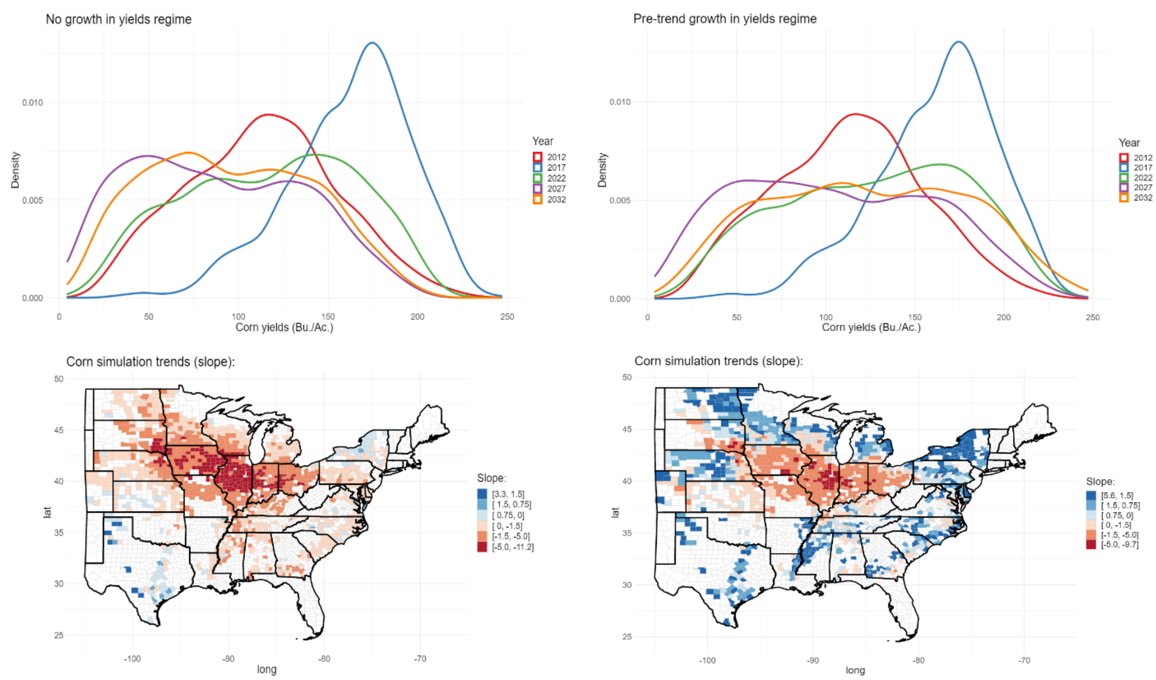

Figure 3: U.S. counties corn yield simulation results.

Note: Top panels describe corn yield probability distributions for each year including historical observations. Bottom panels describe corn yield potential trends (slope) for the projected years.

temperatures. This is an interesting insight, since my results indicate that raising temperatures will have a similar effect on corn yield distributions as the 2012 drought had on corn producers. In a similar manner, we can study the results for soybeans as shown in Figure 4. All soybeans yield probability distributions (2022-2032) peak around the same location as in 2012, but there is a large probability mass closer to zero, suggesting that soybeans productivity in future years will also be similar to the results in 2012. None of my results suggest that future crop production in the U.S. will be similar to crop production in 2017, which seems to be a good year for corn producers, but some U.S. counties can have crop yield realizations similar to observations in 2017.[9]

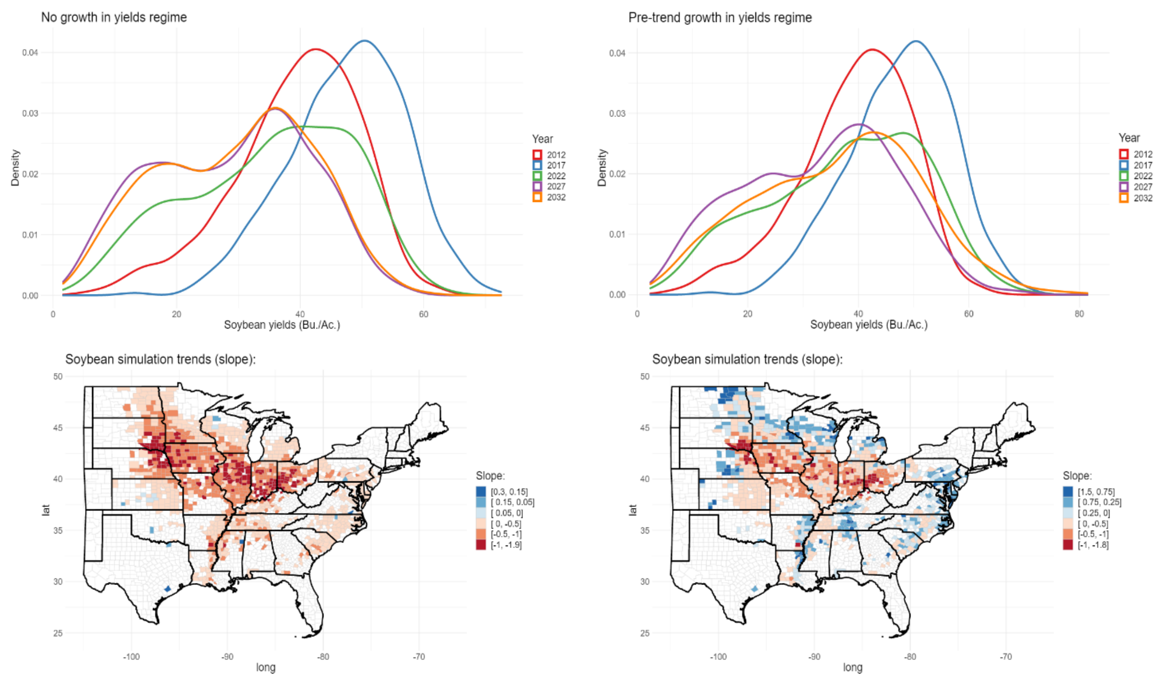

Figure 4: U.S. counties soybean yield simulation results.

Note: Top panels describe soybean yield probability distributions for each year including historical observations. Bottom panels describe soybean yield potential trends (slope) for the projected years.

Corn and soybean yield simulation results indicate that increasing temperatures in the U.S. have the potential to put an end to the upward trend in crop productivity documented in previous studies within the next decade (see, Figure 1). In my study, I find that regardless of the growth

[9] To arrive to these results, I consider the heterogeneity of posterior distributions, so my post-estimation simulations rely on a two-step approach. First, I obtain the expected value of the posterior distributions, $\widehat{\boldsymbol{\theta}}_{\boldsymbol{\iota}}$, via sampling (e.g., the impact of weather on crop yields) with $\mathrm{E}\left[\widehat{\boldsymbol{\theta}}_{\boldsymbol{\iota}} | y_{it}, x_{it}\right] = \int \gamma_i p(\gamma_i, \tau | y_{it}, x_{it}) d\gamma$, where the posterior distribution is equation (4). This first step allows me to obtain a weighted average for the impact of weather on crop yields that depends on the probability of the different realizations of the posterior distribution, effectively accounting for multimodal probabilistic distributions. In the second step, I focus on the most probable crop yield realization for each county estimated with the maximum a posterior (MAP) of (1): $\hat{y}_{it}^{MAP} = \arg\max_{y_i} p\left(y_{it}, \hat{x}_{it} \middle| \widehat{\boldsymbol{\theta}}_{\boldsymbol{\iota}}\right)$, where $\hat{y}_{it}^{MAP}$ denotes my crop yield projection, and $\hat{x}_{it}$ denotes forecasted weather data. Therefore, results described in Figures 3 and 4, and the leftmost panels in Figures 5 and 6 represent the most probable crop yield realization for each county given my data and posterior estimations.

regime being used, simulation results reveal that crop yield realizations observed in 2017 are less likely within the next decade. My conclusions are based on the location and spread of the crop yield potential distributions shown in 2022, 2027, and 2032, and the implied slopes (Figures 3 and 4). Both corn and soybean simulation results indicate that a significant number of counties have yield probability distributions that are substantially lower than the 2017 peak. Previous studies find that crop productivity in the U.S. may decrease by an average of 20-30% for the average between 2025 to 2075 due to high temperatures (see, for instance, Schlenker, Hanemann, and Fisher, (2006), and Lobell *et al.*, (2013)).

Individual results from Midwestern and Eastern U.S. counties for 2022, 2027, and 2032 indicate that the most likely yield outputs in Midwestern states will decrease by up to 11.2 bushels per acre for corn and by up to 1.9 bushels per acre for soybeans within the next decade, regardless of whether an upward crop growth trend is assumed. Non-midwestern counties, however, will experience mild increases in crop yields no greater than 5.6 bushels per acre for corn and 1.5 bushels per acre for soybeans. Having small losses and gains indicates stagnation of the agricultural productivity experienced over the last decades (see Figure 1). There are a number of counties whose yields may be closer to zero; my results, however, cannot provide evidence of whether a county with relatively low levels of yields (e.g., close to zero) will decide to produce each year or not since I cannot impose an ad-hoc cut-off to exclude simulation results that can be considered as *too low*. For instance, corn yield data in 2012, a year with extreme drought conditions, presents cases in which counties produced less than 50 bushels, a substantially low value for that year, per acre while data in 2017 presents cases in which only a handful of counties produced less than 50 bushels per acre.

Besides producing county-level crop yield projections, an additional advantage of my Bayesian formulation described in section 2 is that I can assign probability figures for specific intervals of crop yield realizations and calculating the most probable crop yield realizations for a given year. For this exercise, I decide to focus on Corn Belt states only since Figures 3 and 4 indicate that the largest losses are associated with the largest corn and soybeans producers in this region. In Figures 5 and 6, I describe probabilistic crop yield outlooks for corn and soybeans, respectively, across the pre-2017 growth regime (top figures) and 2017 growth regime (bottom figures). The leftmost column presents the most likely crop yield scenario in 2032; the middle column presents the probability of each county having a crop yield realization lower than the 2017 average in 2032; and the rightmost column presents the probability of each county having a crop yield realization lower than the 2032 average in the same year.

A first insight from the results described in Figure 5 is that the assumption of a crop growth regime determines if county-level projections surpass the 150 bushels per acre threshold (as a reference, average corn yields was 162.07 in 2017). When examining the probability of each county surpassing the 2017 average, results indicate that such crop yield realizations are extremely unlikely, regardless of the crop-growth regime assumption; most of the counties in the Corn Belt states will experience crop yield realizations below 162.07 with a probability equal or higher than 75%. The rightmost maps in Figure 5 divide counties by those that likely will experience higher than average corn yields from those that will not in 2032. Most of the southern Corn Belt counties will experience lower than average crop yield realizations with a probability equal or higher than 75%. Results are similar for soybeans (Figure 6). When examining the probability of each county surpassing the 2017 average, results indicate that such soybeans yield realizations are extremely unlikely, regardless of the crop-growth regime assumption since most of the counties in the Corn Belt states will experience soybeans yield realizations below the 2017 average of 47.57 with a probability equal or higher than 75%. The rightmost maps in Figure 6 divide counties by those that likely will experience higher than average soybeans yields from those that will not. A few Corn Belt counties in the Southmost part of our reference will experience lower than average soybean yield realizations with a probability equal or higher 75%.

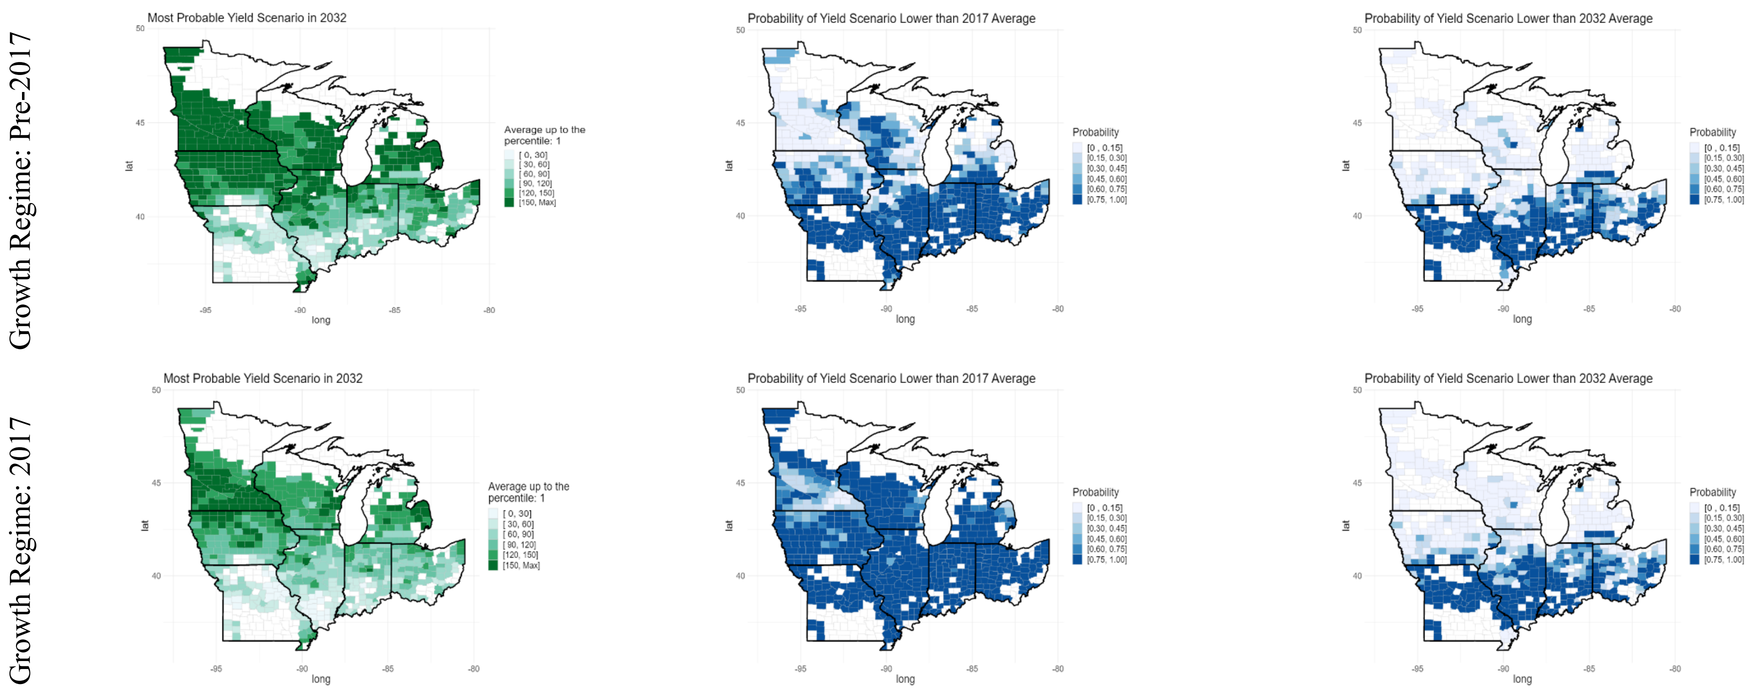

Figure 5: Probabilistic corn yield outlooks across U.S. counties in the Corn Belt states.

Note: Top panels describe corn yield probabilistic outlooks under the assumption of pre-2017 crop growth regime, while bottom panels describe corn yield probabilistic outlooks under the assumption of 2017 crop growth regime. Average corn yields were 162.07 in 2017. Expected corn yields are 123.36 for the pre-2017 growth regime and 95.57 for the 2017 growth regime in 2032.

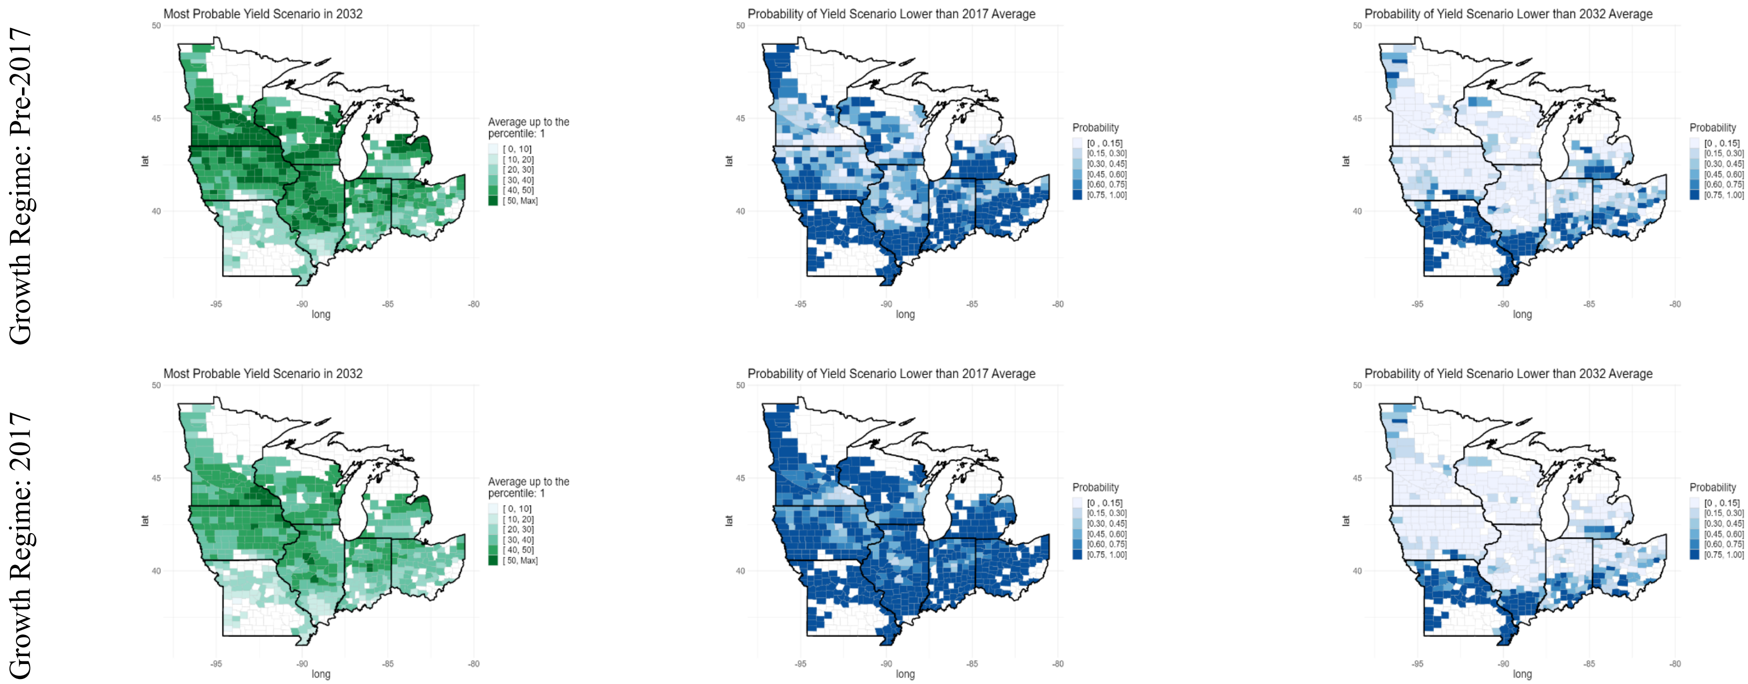

Figure 6: Probabilistic soybeans yield outlooks across U.S. counties in the Corn Belt states.

Note: Top panels describe soybeans yield probabilistic outlooks under the assumption of pre-2017 crop growth regime, while bottom panels describe soybeans yield probabilistic outlooks under the assumption of 2017 crop growth regime. Average soybeans yields were 47.57 in 2017. Expected soybeans yields are 36.38 for the pre-2017 growth regime and 29.61 in 2032.

5. Concluding remarks

In this manuscript, I develop a hierarchical Bayesian framework of the interaction between crop yield growth and weather variations to determine the heterogeneity associated with the decline of agricultural productivity found in other studies and the extent to which crop yield growth will stagnate in the next decade. A probabilistic model of precipitation, growing degree days, and extreme degree days is parametrized to affect the location and spread parameters of a crop yield distribution. Prior and hyper-priors are imposed to account for the combined effect of weather variables on crop yield growth and the known direction of their effect on agricultural productivity. Then, a Beta-type likelihood distribution of crop yields binds together model parameters with yield and weather data. Using corn and soybean yield data in the U.S. along with historical weather observations, a H-MCMC algorithm samples the posterior distributions of the effect of weather variables on crop yields. Posterior distributions are then employed with weather forecasts used in the most recent IPCC report to study crop yield trends in the U.S.

Corn and soybean yield simulation results for the next decade indicate that increasing temperatures in the U.S. have the potential to put an end to the upward trend in crop productivity documented in previous studies within the next decade. I find that (1) warming temperatures can potentially stagnate U.S. crop yield growth in 2032 with a probability equal or higher than 75%, and (2) the largest productivity losses are associated with counties in the Corn Belt. A contribution of my study is providing evidence that climate change will stagnate crop growth in the next decade, putting additional emphasis on the need to have short-term adaptation and mitigation policies. My study, thus, contributes to a growing literature that stresses the need to shift policy efforts to the short-run when dealing with climate change (Loë, Kreutzwiser, Moraru, 2001; Tubiello and Ronsenzweig, 2008; Kistner *et al.,* 2017; Vimic *et al.,* 2022).

References

Alston, J. M., Andersen, M. A., James, J. S., & Pardey, P. G. (2010). Persistence Pays: US Agricultural Productivity Growth and the Benefits from Public R & D Spending.

Ball, E., Schimmelpfennig, D., & Wang, S. L. (2013). Is US agricultural productivity growth slowing? *Applied Economic Perspectives and Policy, 35*, 435–450.

Burke, M., & Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy, 8*, 106–40.

Cribari-Neto, F., & Zeileis, A. (2010). Beta regression in R. *Journal of statistical software, 34*, 1–24.

Dall'Erba, S., Chen, Z., & Nava, N. J. (2021). US interstate trade will mitigate the negative impact of climate change on crop profit. *American Journal of Agricultural Economics*, *103*(5), 1720-1741.

Deschênes, O., & Greenstone, M. (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review, 97*, 354–385.

De Loë, R., Kreutzwiser, R., & Moraru, L. (2001). Adaptation options for the near term: climate change and the Canadian water sector. *Global Environmental Change*, *11*(3), 231-245.

Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2019). Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society), 182*, 389–402.

Gbetibouo, G. A., & Hassan, R. M. (2005). Measuring the economic impact of climate change on major South African field crops: A Ricardian approach. *Global and Planetary change, 47*, 143–152.

Gelman, A. (2008). Objections to Bayesian statistics. *Bayesian Analysis, 3*, 445–449.

Gelman, A., & Hill, J. (2007). Data analysis using regression and hierarchical/multilevel models. *New York, NY: Cambridge*.

Gelman, A., Lee, D., & Guo, J. (2015). Stan: A probabilistic programming language for Bayesian inference and optimization. *Journal of Educational and Behavioral Statistics, 40*, 530–543.

Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., . . . Modrák, M. (2020). Bayesian workflow. *arXiv preprint arXiv:2011.01808*.

Group, P. C., & others. (2011). PRISM climate data. *Oregon State University*.

Huffman, W. E., & Evenson, R. E. (1992). Contributions of public and private science and technology to US agricultural productivity. *American Journal of Agricultural Economics, 74*, 751–756.

Kistner, E., Kellner, O., Andresen, J., Todey, D., & Morton, L. W. (2018). Vulnerability of specialty crops to short-term climatic variability and adaptation strategies in the Midwestern USA. *Climatic change*, *146*(1), 145-158.

Liang, X.-Z., Wu, Y., Chambers, R. G., Schmoldt, D. L., Gao, W., Liu, C., . . . Kennedy, J. A. (2017). Determining climate effects on US total agricultural productivity. *Proceedings of the National Academy of Sciences, 114*, E2285–E2292.

Lobell, D. B., Hammer, G. L., McLean, G., Messina, C., Roberts, M. J., & Schlenker, W. (2013). The critical role of extreme heat for maize production in the United States. *Nature Climate Change, 3*, 497–501.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1994). The impact of global warming on agriculture: a Ricardian Analysis. *American Economic Review*, 753–771.

Nelson, C. H., & Preckel, P. V. (1989). The conditional beta distribution as a stochastic production function. *American Journal of Agricultural Economics, 71*, 370–378.
Ortiz-Bobea, A. (2020). The role of nonfarm influences in Ricardian estimates of climate change impacts on US agriculture. *American Journal of Agricultural Economics, 102*, 934–959.
Ortiz-Bobea, A., & Tack, J. (2018). Is another genetic revolution needed to offset climate change impacts for US maize yields? *Environmental Research Letters, 13*, 124009.
Polzehl, J., Papafitsoros, K., & Tabelow, K. (2020). Patch-Wise Adaptive Weights Smoothing in R. *Journal of Statistical Software, 95*, 1–27.
Reinsborough, M. J. (2003). A Ricardian model of climate change in Canada. *Canadian Journal of Economics/Revue Canadienne d'économique, 36*, 21–40.
Roberts, M. J., Schlenker, W., & Eyer, J. (2013). Agronomic weather measures in econometric models of crop yield with implications for climate change. *American Journal of Agricultural Economics, 95*, 236–243.
Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. *Proceedings of the National Academy of Sciences, 106*, 15594–15598.
Schlenker, W., Hanemann, W. M., & Fisher, A. C. (2006). The impact of global warming on US agriculture: An econometric analysis of optimal growing conditions. *Review of Economics and Statistics, 88*, 113–125.
Snyder, R. L. (1985). Hand calculating degree days. *Agricultural and Forest Meteorology, 35*, 353–358.
Tack, J., Barkley, A., & Nalley, L. L. (2015). Effect of warming temperatures on US wheat yields. *Proceedings of the National Academy of Sciences, 112*, 6931–6936.
Tack, J., Coble, K., & Barnett, B. (2018). Warming temperatures will likely induce higher premium rates and government outlays for the US crop insurance program. *Agricultural economics, 49*, 635–647.
Thrasher, B., Maurer, E. P., McKellar, C., & Duffy, P. B. (2012). Bias correcting climate model simulated daily temperature extremes with quantile mapping. *Hydrology and Earth System Sciences, 16*, 3309–3314.
Tubiello, F. N., & Rosenzweig, C. (2008) Developing climate change impact metrics for agriculture. *Integrated Assessment Journal*, 8(1).
van de Schoot, R., Depaoli, S., King, R., Kramer, B., Märtens, K., Tadesse, M. G., . . . others. (2021). Bayesian statistics and modelling. *Nature Reviews Methods Primers, 1*, 1–26.
van Passel, S., Massetti, E., & Mendelsohn, R. (2017). A Ricardian analysis of the impact of climate change on European agriculture. *Environmental and Resource Economics, 67*, 725–760.
Vuković Vimić, A., Djurdjević, V., Ranković-Vasić, Z., Nikolić, D., Ćosić, M., Lipovac, A., ... & Vujadinović Mandić, M. (2022). Enhancing Capacity for Short-Term Climate Change Adaptations in Agriculture in Serbia: Development of Integrated Agrometeorological Prediction System. *Atmosphere*, *13*(8), 1337.
Wang, J., Mendelsohn, R., Dinar, A., Huang, J., Rozelle, S., & Zhang, L. (2009). The Impact of Climate Change on China's Agriculture. *Agricultural Economics, 40*, 323–337.
Zhu Y., Goodwin B. K., & Ghosh, S. K. (2011). Modeling Yield Risk Under Technological Change: Dynamic Yield Distributions and the U.S. Crop Insurance Program. *Journal of Agricultural and Resource Economics Association,* 36(1): 192-210.