



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Estimating Non-linear Weather Impacts on Corn Yield— A Bayesian Approach

Tian Yu and Bruce A. Babcock

Working Paper 11-WP 522
April 2011

**Center for Agricultural and Rural Development
Iowa State University
Ames, Iowa 50011-1070
www.card.iastate.edu**

Tian Yu is a research associate at Morgan Stanley. She was a PhD candidate and graduate research assistant in the Department of Economics at Iowa State University at the time this research was conducted. Bruce Babcock is a professor of economics and director of the Center for Agricultural and Rural Development at Iowa State University.

This paper is available online on the CARD Web site: www.card.iastate.edu. Permission is granted to excerpt or quote this information with appropriate attribution to the authors.

Questions or comments about the contents of this paper should be directed to Bruce Babcock, 578 Heady Hall, Iowa State University, Ames, IA 50011-1070; Ph: (515) 294-6785; Fax: (515) 294-6336; E-mail: babcock@iastate.edu.

<p>Iowa State University does not discriminate on the basis of race, color, age, religion, national origin, sexual orientation, gender identity, genetic information, sex, marital status, disability, or status as a U.S. veteran. Inquiries can be directed to the Director of Equal Opportunity and Compliance, 3280 Beardshear Hall, (515) 294-7612.</p>
--

Estimating Non-linear Weather Impacts on Corn Yield – a Bayesian Approach

Tian Yu and Bruce A. Babcock

Abstract

We estimate impacts of rainfall and temperature on corn yields by fitting a linear spline model with endogenous thresholds. Using Gibbs sampling and the Metropolis - Hastings algorithm, we simultaneously estimate the thresholds and other model parameters. A hierarchical structure is applied to capture county-specific factors determining corn yields. Results indicate that impacts of both rainfall and temperature are nonlinear and asymmetric in most states. Yield is concave in both weather variables. Corn yield decreases significantly when temperature increases beyond a certain threshold, and when the amount of rainfall decreases below a certain threshold. Flooding is another source of yield loss in some states. A moderate amount of heat is beneficial to corn yield in northern states, but not in other states. Both the levels of the thresholds and the magnitudes of the weather effects are estimated to be different across states in the Corn Belt.

Key words: Bayesian estimation, Gibbs sampler, hierarchical structure, Metropolis - Hastings algorithm, non-linear

JEL codes: C11, C13, Q10, Q54.

In rain-fed agricultural regions, weather conditions have substantial impacts on crop productivity. Favorable weather conditions for dryland crop production, including a proper amount of heat and rainfall during the growing season, are critical factors determining

yield outcomes (O'Brien 1993). Most previous studies found that corn yield decreases in temperature and increases in the amount of rainfall (Lobell and Asner 2003; Deschenes and Greenstone 2007). Understanding how weather variables affect crop yield is essential in measuring yield risk and rating crop insurance plans (Yu and Babcock 2010).

In quantifying the impact of weather on corn yield, one widely used approach is to estimate a reduced form statistical relationship between corn yield and weather variables. Two commonly used reduced forms are linear and quadratic functions. The linear specification seems convenient but it could be restrictive as well. There is evidence that the response of corn yield to temperature and rainfall may not be constant over the whole range of possible weather outcomes. In fact, an increase of temperature within the range between 8° Celsius and 32° Celsius is found to be beneficial to corn yield, but increasing the temperature further beyond 34° Celsius leads to yield losses (Schlenker, Hanemann, and Fisher 2006). If weather factors are beneficial to corn yield in some ranges but not in others, then the linear specification would generate a misleading result. One way to model the non-linearity is to include a quadratic term. However, the quadratic functional form restricts the yield response to be symmetric. There is evidence that bad growing conditions typically cause more yield losses than good growing conditions cause yield gains. In fact, based on county-level panel data of corn yield, temperature, and precipitation from 1950 to 2004 in 2000 counties in the U.S., Schlenker and Roberts (2006) found that an increase in temperature above 25° Celsius decreases corn yield growth rates at an increasing rate. Nonparametric estimation is a viable approach to estimate the asymmetric nonlinear impacts of weather variables (Schlenker and Roberts 2006). One limitation of the nonparametric approach, however, is that in some ranges of the weather variable, especially at the two tails of the distribution, the number of observations might be too small so that the standard errors tend to be very large.

In this article, we propose to estimate a two-knot linear spline model with endogenous thresholds. We simultaneously estimate the threshold parameters and other model param-

eters using a sampling-based Bayesian approach. Our specification captures the nonlinear and asymmetric feature of weather effects. The endogenous-knot linear-spline specification is relatively simple and yet flexible. Estimation in the Bayesian framework brings the advantages of high computational efficiency and quick convergence. In light of the findings that a modest increase in temperature benefits corn yield in the northern regions of the U.S. but not in other areas (Adams et al. 1990), we also examine possible geographical differences in how the weather factors influence corn yield.

The rest of this article is organized as follows: in section two we specify the two-knot linear spline model. In section three, we describe the Bayesian approach, which is applied to estimate the model. In section four, we present the estimation results. In the last section, we present a brief conclusion.

The Yield Model

We specify corn yield to be composed of a linear trend plus a function of weather variables, specifically:

$$\begin{aligned}
 Y_{i,t} = & \alpha_i + \sum_{r=1}^R \beta_{1,r}(CRD_r \times Time) \\
 & + \beta_2 \min(0, (Temp_{i,t} - \theta_l)) + \beta_3 Temp_{i,t} + \beta_4 \max(0, (Temp_{i,t} - \theta_u)) \\
 (1) \quad & + \beta_5 \min(0, (Rain_{i,t} - \lambda_l)) + \beta_6 Rain_{i,t} + \beta_7 \max(0, (Rain_{i,t} - \lambda_u)) + \varepsilon_{i,t}.
 \end{aligned}$$

Subscripts i , r , and t denote county, Crop Reporting District (CRD), and time, respectively. We denote the total number of counties as N , the total number of years as T , and the total number of CRDs as R . Y denotes corn yield. $Time$ is a time trend variable, which takes values 0 to 28 for years 1980 to 2008. CRD_r , $r = 1, 2, \dots, R$, denotes the regional dummy variable. $CRD_r = 1$, if the yield observation is from crop reporting district r , and $CRD_r = 0$ otherwise. $Temp$ and $Rain$ denote mean monthly temperature and mean monthly rainfall in the growing season, respectively.¹ Without loss of generality, we recenter the temperature

and rainfall variables at zero by subtracting the historical mean from each temperature or rainfall observation. The recentering process does not affect our estimation of other parameters except α_i . After recentering, α_i measures the average corn yield in the base year (1980) in county i when temperature and rainfall are at their historical mean levels. We assume that the error term $\varepsilon_{i,t}$ is i.i.d. normal with mean zero and variance σ_ε^2 .

In our specification, corn yield is expressed in an additive form of a linear trend and two-knot linear spline functions of rainfall and temperature. For the linear trend, we permit the intercept term α_i to vary across counties. Thus, α_i captures time-invariant county-specific factors that influence corn yield. Since fixed-effect estimation in a Bayesian setting is inefficient, we specify a hierarchical structure for α_i . Specifically, we assume that α_i , for $i = 1, 2, \dots, N$, is independently and identically distributed from a normal distribution $N(\alpha, \sigma_\alpha^2)$. α and σ_α^2 are hierarchical parameters, which are estimated simultaneously with other model parameters. We allow the trend slope to vary across CRDs by including an interaction term between the regional dummy variable and the time trend variable. $\beta_{1,r}$ is fixed for any given CRD but differs across CRDs. Coefficients β_2 to β_7 are restricted to be constant across all counties. Equation (1) is essentially a mixed model, with a random effect α_i and a fixed effect β . In a Bayesian framework, we simultaneously estimate model parameters α , β , θ , and λ .

The two-knot linear spline functions of rainfall and temperature capture the potentially asymmetric and nonlinear feature of weather effects. We allow weather effects to be different when temperature (or the amount of rainfall) is lower than usual, within the middle range, or higher than usual. This is achieved by introducing endogenous threshold parameters. λ_l (or θ_l) and λ_u (or θ_u) denote the lower and upper thresholds that divide the domain of the weather variable into three ranges. The min and max operators serve as switches that turn on the specific β that we want to estimate when the weather variable falls into the specific range. The yield model allows a changing yield response as the weather regime

changes. To better illustrate this feature, we rewrite model (1) as:

$$Y_{i,t} = \tilde{\alpha}_i + \sum_{r=1}^R \beta_{1,r}(CRD_r \times Time) + \beta_{temp,s}Temp_{i,t} + \beta_{rain,s}Rain_{i,t} + \epsilon_{i,t}$$

$$(2) \quad \beta_{temp,s} \equiv \begin{cases} \beta_{temp,cl} = \beta_2 + \beta_3 & \text{if } Temp \leq \lambda_l \\ \beta_{temp,nt} = \beta_3 & \text{if } \lambda_l < Temp < \lambda_u \\ \beta_{temp,ht} = \beta_3 + \beta_4 & \text{if } Temp \geq \lambda_u \end{cases}$$

$$\beta_{rain,s} \equiv \begin{cases} \beta_{rain,dr} = \beta_5 + \beta_6 & \text{if } Rain \leq \lambda_l \\ \beta_{rain,nr} = \beta_6 & \text{if } \lambda_l < Rain < \lambda_u \\ \beta_{rain,fl} = \beta_6 + \beta_7 & \text{if } Rain \geq \lambda_u \end{cases}$$

Subscript s denotes the state of nature, which is defined by the weather variable falling into one of the threshold-divided ranges. $\beta_{temp,s}$ (or $\beta_{rain,s}$) measures the marginal effect of temperature (or rainfall) in each state of nature. The marginal effect can potentially change depending on the weather condition. For example, one inch of rainfall in drought years could result in a different amount of change in corn yield than in normal years.

The Bayesian Approach

We now turn to the estimation methodology. Note that (1) is non-linear in threshold parameters λ and θ . Non-linear least square estimation (NLS) and maximum likelihood estimation (MLE) could potentially be used to estimate (1). In this article, we take the sampling-based Bayesian approach. The advantages of Bayesian estimation include easy implementation, fast convergence, and computational efficiency. We apply the Markov chain Monte Carlo (MCMC) in estimating (1). The implementation of the MCMC is essentially sequentially taking draws from (sequentially updated) conditional posterior distributions of model parameters. The ease of applying the MCMC to our model results from the observation that conditional on threshold parameters, (1) becomes a well known linear regression model with a hierarchical structure. Under conjugate priors, the conditional

posterior distributions of parameters except λ and θ are readily derived. And it is easy to apply the Gibbs sampling to simulate draws from the conditional posterior distributions (Lindley and Smith 1972; Gelfand et al. 1990; Chib and Carlin 1999). The conditional posterior distributions of λ and θ require some derivations and are not of any recognizable distributional forms. Thus, we apply the Metropolis - Hastings algorithm (Chib and Greenberg 1995; Gelman, Roberst, and Gilks 1995) in the step of drawing λ and θ from their conditional posterior distributions. The estimation process would become clear once we have the prior distributions, the likelihood function, and the posterior distributions, which will be specified and derived in the following sections.

Priors

Following Chib and Carlin (1999), we assume normal priors for α and β , and inverse gamma priors for the variance parameters. These are conjugate priors in the sense that the posterior distributions belong to the same family as the prior probability distributions. We assume conjugate priors due to the computational ease associated with this specification. Results are not sensitive to the specific form of the prior distribution. For notation simplicity, we stack all the fixed-effect parameters into one vector and define as: $\beta \equiv [\beta_{1,1}, \beta_{1,2}, \dots, \beta_{1,R}, \beta_2, \beta_3, \dots, \beta_7]'$. We assume the prior distribution of β to be $N(\mu_\beta, V_\beta)$, where N denotes the normal distribution. As specified above, α_i is of normal distribution $N(\alpha, \sigma_\alpha^2)$. We assume the prior distributions of the hierarchical mean and variance parameters α and σ_α^2 to be $N(\mu_\alpha, V_\alpha)$ and $IG(a_1, a_2)$, respectively. IG denotes the inverse gamma distribution. We assume the prior distribution of the variance of the error term σ_ϵ^2 to be $IG(e_1, e_2)$. We assume that μ_β and μ_α equal to the corresponding OLS estimates of a linear yield model, $Y_{i,t} = \alpha_i + \sum_{r=1}^R \beta_{1,r}(CRD_r \times Time) + \beta_{temp}Temp_{i,t} + \beta_{rain}Rain_{i,t} + \epsilon_{i,t}$, with $\mu = \beta_{temp}^{OLS} \times [1, 1, 1, 0, 0, 0] + \beta_{rain}^{OLS} \times [0, 0, 0, 1, 1, 1]'$. We assume $V_\alpha = 1600$ and $V_\beta = 100$ so that the normal priors are reasonably diffused. For the inverse gamma distributions, we set $e_1 = a_1 = 3$ and $e_2 = a_2 = \frac{1}{2 \times 100}$ so that the standard deviation of

the inverse gamma distribution is 10.² Finally, we assume the prior distribution of the lower and upper thresholds to be the joint uniform distribution over the domain of the weather variable, with the restriction that the upper threshold is larger than the lower threshold. Specifically, $p(\theta_l, \theta_u) = \frac{2}{(Temp_{max} - Temp_{min})^2} I(Temp_{min} < \theta_l < \theta_u < Temp_{max})$ and $p(\lambda_l, \lambda_u) = \frac{2}{(Rain_{max} - Rain_{min})^2} I(Rain_{min} < \lambda_l < \lambda_u < Rain_{max})$, where $I(\cdot)$ denotes the index function. We set $Temp_{max}$ (or $Rain_{max}$) at the recentered 95% percentile of all temperature (or rainfall) observations and $Temp_{min}$ (or $Rain_{min}$) at the recentered 5% percentile of all temperature (or rainfall) observations.

The Likelihood Function

The other piece of information we need in order to derive the posterior distributions is the likelihood function. For notation simplicity, we stack the explanatory variables as:

$$\begin{aligned}
 X(\theta, \lambda)_{i,t} \equiv & [(CRD_1 \times T), (CRD_2 \times T), \dots, (CRD_R \times T), \\
 & (\min(0, (Temp_{i,t} - \theta_l))), Temp_{i,t}, (\max(0, (Temp_{i,t} - \theta_u))), \\
 (3) \quad & (\min(0, (Rain_{i,t} - \lambda_l))), Rain_{i,t}, (\max(0, (Rain_{i,t} - \lambda_u)))]].
 \end{aligned}$$

Then the likelihood function is:

$$\begin{aligned}
 L(\cdot) &= \prod_{i=1}^{N \times T} p\{Y_{i,t} | \alpha_i, \beta, \sigma_\varepsilon^2, X_{i,t}(\theta, \lambda)\} \\
 (4) \quad &= (2\pi\sigma_\varepsilon^2)^{-\frac{N \times T}{2}} \exp\left[-\frac{1}{2\sigma_\varepsilon^2} \sum_{i=1}^N \sum_{t=1}^T (Y_{i,t} - \alpha_i - X_{i,t}(\theta, \lambda)\beta)^2\right].
 \end{aligned}$$

Notation $X_{i,t}(\theta, \lambda)$ emphasizes the fact that $X_{i,t}$ depends on thresholds θ and λ . To be succinct, we drop θ and λ and simply note $X_{i,t}$ when referring to $X_{i,t}(\theta, \lambda)$ in the following text.

Conditional Posteriors

Based on the priors and the likelihood function specified above, we derive the conditional posterior distributions. As mentioned above, standard results for the linear regression model apply to our model once we condition on the threshold parameters. Following Chib and Carlin (1999), the posterior distributions are derived by applying conclusions by Lindley and Smith (1972). Specifically, $p(\alpha_i | \beta, \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y) \sim N(D_1 d_1, D_1)$, where $D_1 = (\frac{T}{\sigma_\varepsilon^2} + \frac{1}{\sigma_\alpha^2})^{-1}$, $d_1 = \frac{\iota_T'(Y_i - X_i \beta)}{\sigma_\varepsilon^2} + \frac{\alpha}{\sigma_\alpha^2}$.³ Similarly, the posterior distribution of β is also normal. $p(\beta | \alpha_i, \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y) \sim N(D_2 d_2, D_2)$, where $D_2 = (\frac{X'X}{\sigma_\varepsilon^2} + V_\beta^{-1})^{-1}$, $d_2 = \frac{X'(y - \bar{\alpha})}{\sigma_\varepsilon^2} + V_\beta^{-1} \mu_\beta$, $\bar{\alpha} \equiv [(\iota_T \alpha_1)', (\iota_T \alpha_2)', \dots, (\iota_T \alpha_N)']'$. Again, the posterior distribution of α is normal. $p(\alpha | \alpha_i, \beta, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y) \sim N(D_3 d_3, D_3)$, where $D_3 = (\frac{N}{\sigma_\alpha^2} + V_\alpha^{-1})^{-1}$, $d_3 = \frac{\iota_N'[\alpha_1, \alpha_2, \dots, \alpha_N]'}{\sigma_\alpha^2} + V_\alpha^{-1} \mu_\alpha$.

Conditional posterior distributions of the variance parameters are inverse gamma. $p(\sigma_\alpha^2 | \alpha_i, \beta, \sigma_\varepsilon^2, \theta, \lambda, Y) \sim IG(\frac{N}{2} + a_1, (\frac{1}{a_2} + 0.5 \sum_{i=1}^N (\alpha_i - \alpha)^2)^{-1})$. $p(\sigma_\varepsilon^2 | \alpha_i, \beta, \sigma_\alpha^2, \theta, \lambda, Y) \sim IG(\frac{N \times T}{2} + e_1, (\frac{1}{e_2} + 0.5(Y - [\iota_T \alpha_1, \iota_T \alpha_2, \dots, \iota_T \alpha_N]' - X\beta)')(Y - [\iota_T \alpha_1, \iota_T \alpha_2, \dots, \iota_T \alpha_N]' - X\beta))^{-1})$. The posteriors for θ and λ are not derived in the same fashion, but are straightforward under uniform priors. Since the conditional posterior is proportional to the likelihood function multiplied by the prior for θ and λ , the kernel of the conditional posterior of θ and λ is simply the likelihood function as specified in (4).

Implementing the Gibbs Sampling and the Metropolis - Hastings Algorithm

We take a sampling-based approach to estimate the hierarchical model (1). We implement the Gibbs sampler to draw β , $\{\alpha_i\}$, α , σ_ε^2 , and σ_α^2 . Since the conditional posterior distribution of θ and λ are not of any recognizable distributional form, we employ the random-walk Metropolis-Hastings algorithm to draw θ and λ . The procedure of implementing the Gibbs Sampler and the Metropolis - Hastings algorithm is as follows. (1) Start with initial values of model parameters: β , $\{\alpha_i\}$, α , σ_ε^2 , σ_α^2 , θ , and λ . Calculate $X(\theta, \lambda)$ with θ and λ evalu-

ated at initial values. (2) Draw $\tilde{\beta}$ from its conditional posterior distribution, conditional on the initial values of other model parameters. (3) Draw $\{\alpha_i\}$ from their conditional posterior distributions, conditional on the most recent draw of $\tilde{\beta}$ and other model parameters at their initial values. (4)-(6) Draw α , σ_ε^2 , and σ_α^2 one by one from their conditional posterior distributions, conditional on other parameters, which are updated to their most recent draws. (7) Sample $\theta^* = \theta^{[-1]} + \varepsilon_\theta$, where $\theta^{[-1]}$ denotes the last draw of θ in the chain and ε_θ is a draw from normal distribution $N(0, \sigma_\theta)$. Draw $u \sim U(0, 1)$. Denote $p(\cdot)$ as the conditional posterior distribution of θ , conditional on the most recent draws of other model parameters. Calculate the index function $I(Temp_{min} < \theta_l^* < \theta_u^* < Temp_{max})$. If $u \leq \frac{p(\theta^*)}{p(\theta^{[-1]})}$ and the index function is equal to one, then keep θ^* as a draw from the conditional posterior, that is, $\tilde{\theta} = \theta^*$. Otherwise, use the last draw from the chain $\tilde{\theta} = \theta^{[-1]}$. Update $X(\tilde{\theta}, \lambda)$. (8) Draw $\tilde{\lambda}$ in a similar fashion as with $\tilde{\theta}$ in step (7), and update $X(\theta, \tilde{\lambda})$ accordingly. (9) Repeat steps (2)-(8) 20,000 iterations, updating the posterior conditionals at each iteration using the most recently simulated values in the chain. (10) Discard an early set of parameter simulations (the first 5,000 iterations) as the burn-in period. Use the subsequent draws to make Bayesian posterior inferences.

Following Chib and Carlin (1999), we improve the above process by drawing the random effect parameters $\{\alpha_i\}$ and fixed effect parameter β in a single block. We replace steps (2)-(3) by drawing $(\{\alpha_i\}, \beta)$ from $p(\{\alpha_i\}, \beta | \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, Y)$. Specifically, we draw the group via the method of composition. We first draw $\tilde{\beta}$ from $p(\beta | \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y)$ and then draw each α_i independently from its complete posterior distribution evaluated at $\tilde{\beta}$: $p(\alpha_i | \tilde{\beta}, \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y)$. Denote $\Sigma = \sigma_\varepsilon^2 I_T + \sigma_\alpha^2 \iota_T \iota_T'$. Then $p(\beta | \alpha, \sigma_\varepsilon^2, \sigma_\alpha^2, \theta, \lambda, Y) \sim N(D_4 d_4, D_4)$, where $D_4 = (X'(I_N \otimes \Sigma^{-1})X + V_\beta^{-1})^{-1}$, $d_4 = X'(I_N \otimes \Sigma^{-1})(Y - \iota_{NT}\alpha) + V_\beta^{-1}\mu_\beta$. The strategy of grouping together correlated parameters will generally facilitate the mixing of the chain and thereby reduce numerical standard errors associated with Gibbs sampling estimates.

Data

We collected weather and yield data in major non-irrigated corn-production states from 1980 to 2008. Those states are Illinois, Indiana, Iowa, Michigan, Minnesota, Missouri, Ohio, and Wisconsin. County-level production and planted acreage data were collected from the National Agricultural Statistics Service (NASS) to calculate yield per planted acre. Observations with zero production or missing acreage data were deleted. To focus our attention on major production areas, only counties with yield data in all years from 1980 to 2008 were kept. Weather data were collected from the National Oceanic and Atmospheric Administration (NOAA). We obtained data of monthly mean temperature (MNTM) and total monthly precipitation (TPCP) from all weather stations located in the eight states. For most of the weather stations, NOAA identified the county where each weather station was located. For stations not identified to any county, we found the nearest station with a county name (by calculating the distance between weather stations using latitude and longitude information) and assigned the unidentified station to the county of the nearest station. Most of the counties are matched with at least one weather station. For counties with multiple weather stations, we took the simple average of weather records from all weather stations located in the county. For each year, county level corn yields were matched with county-level weather data. We substituted missing values of rainfall or temperature with the average value of the CRD that the county belongs to.⁴

Estimation Results

Using a panel of county-level corn yield with matched rainfall and temperature data, we estimate model (1) state by state. For each state, we apply the MCMC as described above. We run 20,000 iterations with the first 5000 iterations as the burn-in period. Parameters simulated by the Gibbs sampler converges very fast. The posterior distributions of threshold parameters are insensitive to changes in initial values, which indicates that simulations

using the M-H algorithm also converge. In applying the M-H algorithm, we tune the variance parameter of the random walk chain so that the acceptance rate is maintained at around 0.3.⁵ Based on the 15,000 kept draws for each parameter, we calculate the mean and standard deviation of the posterior distributions. Table 1 and table 2 present these results.

In table 1, $\beta_{1,r}$, for $r = 1, 2, \dots, R$, measures the CRD-specific slope of the linear trend. There are nine CRDs in each state except Michigan and Minnesota, where we only have data for seven CRDs for each state. Trend estimates range from low of nearly zero in CRD 6 in Ohio to as high as about 2.5 bushels per acre per year in Illinois, Iowa, and Minnesota. Trend estimates are positive and statistically significant in most CRDs as expected. α_0 measures the state average of the county-specific intercept term α_i . Since we re-centered weather variables at zero, α_0 indicates the state average yield in 1980 if temperature and the amount of rainfall were at the historical mean levels. α_0 is estimated to be around 70-120 bushels per acre. σ_α^2 is the variance parameter of the distribution of α_i . Variance of the county-specific intercept ranges from 77 in Iowa to 477 in Minnesota, averaging around 200. The variance parameter of the residual term is around 200.

Table 2 presents the posterior mean and standard deviation of parameters that capture the impacts of weather. β_3 is the estimate of the marginal effect of temperature on corn yield when temperature is between the lower and the upper thresholds. β_3 is estimated to be negative in all states and is statistically significant in all states except Minnesota and Michigan. As temperature increases one degrees Fahrenheit within the range, corn yield is estimated to decrease about 8 bushels per acre in Ohio and around 6 bushels per acre in Illinois, Iowa, Missouri and Wisconsin. β_2 measures the difference between the marginal effect of temperature when temperature is below the lower threshold and the marginal effect of temperature when temperature is between the two thresholds. β_2 is estimated to be positive and statistically significantly in all states except Missouri. This indicates that, on the margin, heat is less harmful in cooler weather conditions. The difference in marginal effect is about 3 to 8 bushels per acre. β_4 measures the difference between the marginal

effect of temperature when temperature is above the upper threshold and the marginal effect of temperature when temperature is between the two thresholds. β_4 is estimated to be insignificant in Illinois, Michigan, Missouri, Ohio, and Wisconsin, indicating that the marginal damage of excess heat remains about the same once temperature reaches beyond the lower threshold. In other words, the upper temperature threshold is redundant for these states.⁶ However, β_4 is useful in identifying the extra damage due to excessive heat when temperature reaches beyond the upper threshold in Iowa and Minnesota. β_4 is negative and statistically significant in these two states. As temperature increases every degree Fahrenheit beyond the upper threshold, corn yield losses an additional 9.6 bushels per acre in Minnesota and an additional 6 bushels per acre in Iowa as compared with the marginal effect of temperature when temperature is between the two thresholds. Note that the estimated values of β_2 , β_3 , and β_4 indicate that in Indiana the marginal loss from heat is largest when temperature is within the middle range, which is against common sense. We suspect that the estimated values of the lower and upper temperature thresholds in Indiana are too close so that estimation of β_3 might be dominated by noise. As we will present later, in cases when the two thresholds are close, the one-knot specification offers a better estimation.

The marginal effect of rainfall is captured by parameters β_5 , β_6 , and β_7 . β_6 measures the marginal effect of rainfall when rainfall is between the lower and upper thresholds. β_6 is positive and statistically significant in Illinois and Wisconsin but is negative in Minnesota. Within this middle range of rainfall, one inch increase in rainfall increases corn yield by 7 bushels per acre in Illinois and 4 bushels per acre in Wisconsin but decreases corn yield by 2.5 bushels per acre in Minnesota. β_6 is insignificant in other states. β_5 is positive and statistically significant in all states except Michigan and Missouri. This indicates that the marginal effect of rainfall is very different in dry weather conditions as compared with the ‘normal weather’ (the amount of rainfall between the two thresholds). The difference is that one inch increase in rainfall when the amount of rainfall is below the lower threshold

brings in an additional marginal benefit of 10 to 12 bushels per acre to corn yield in Illinois, Indiana, Iowa, Minnesota, Ohio, and Wisconsin. β_7 measures the difference in marginal effects of rainfall between flooding and the middle range rainfall. It is negative and significant in all states except Michigan. β_7 ranges from 5 bushels per acre to 17 bushels per acre.

The bottom part of table 2 presents estimation results for threshold parameters. Since we recentered rainfall and temperature variables to zero, θ and λ indicate the distance between the threshold and the historical mean of the weather variable. Results show that the lower threshold of temperature is about 2.5 degrees Fahrenheit below the mean temperature in Illinois and Michigan, and about 1 to 2 degrees Fahrenheit above the mean temperature in Ohio and Wisconsin. The lower threshold is around the historical mean temperature in Iowa and Minnesota. The upper threshold of temperature also varies across states, ranging from the mean temperature to 3.6 degrees Fahrenheit above the mean temperature. As pointed out above, β_4 is insignificant in Illinois, Michigan, Missouri, Ohio, and Wisconsin. The upper threshold of temperature in these states are ineffective. In other words, there is a statistically significant change in temperature effect around the lower threshold but not around the upper threshold. In most states, the lower threshold of rainfall is less than an inch below the historical mean rainfall. The upper threshold of rainfall ranges from 0.5 inches to 2 inches above the mean. To compare across states the absolute values of the thresholds instead of the relative distances, we add back the historical mean of rainfall and temperature to the posterior mean of θ and λ and present the results in table 3. Temperature thresholds vary significantly across states, with lower values in northern states such as Michigan, Minnesota and Wisconsin. Rainfall thresholds exhibits less variation. The lower threshold of rainfall is around 3.5 inches. And the upper threshold of rainfall is around 6 inches in Iowa, Minnesota, and Missouri, and is about 4.5 to 5.5 inches in other states.

The Marginal Effect

To better evaluate the marginal effect of weather variables in each weather regime, we simulate $\beta_{temp,s}$ and $\beta_{rain,s}$. For each of the 15,000 kept draws of β_2 to β_7 , we simulate $\beta_{temp,s}$ and $\beta_{rain,s}$ according to equation (2). We present the mean and standard deviation of the simulated $\beta_{temp,s}$ and $\beta_{rain,s}$ in table 4. The marginal effect of temperature when temperature is below the lower threshold is measured by $\beta_{temp,cl}$. $\beta_{temp,cl}$ is positive and significant in northern states such as Michigan, Minnesota, and Wisconsin, but negative and significant in Illinois, Iowa, Missouri, and Ohio. One degree Fahrenheit increase in temperature in this weather regime increases corn yield by 1.6 bushels per acre in Wisconsin, by 2.5 bushels per acre in Minnesota, and by 3 bushels per acre in Michigan. On the other hand, the marginal effect is to decrease corn yield by 1.7 bushels per acre in Illinois, by 2 bushels per acre in Iowa, by 2.7 bushels per acre in Ohio, and by 5.7 bushels per acre in Missouri. In Illinois, Michigan, Ohio, and Wisconsin, the marginal effect of temperature stays roughly the same once temperature is above the lower threshold. Once temperature is above the lower threshold, the marginal damage of increasing temperature by one degree Fahrenheit is about 6-8 bushels per acre in Illinois, Ohio, and Wisconsin and around 1.5 bushels per acre in Michigan. On the other hand, in Iowa and Minnesota, the adverse marginal effect of heat is much higher when temperature increases beyond the upper threshold compared with when temperature is between the two thresholds. The marginal effect of temperature is about 11 bushels per acre in Iowa and Minnesota in the hot weather condition. In Missouri, the marginal effect of temperature remains at about -6 bushels per acre over the whole range of temperature.

The marginal effect of rainfall when the amount of rainfall is below the lower threshold is universally positive and significant. The marginal benefit of one inch of rainfall in this rainfall regime ranges from 7 bushels per acre in Missouri to 18.7 bushels per acre in Illinois. The marginal effect of rainfall when the amount of rainfall is above the upper

threshold is negative and significant in Indiana, Iowa, Minnesota, Missouri, and Wisconsin. The marginal damage of rainfall in this rainfall regime is around 5 bushels per acre in Indiana, Missouri, and Wisconsin and reaches 13-14 bushel per acre in Iowa and Minnesota. In Illinois, Michigan, and Ohio, the marginal effect of rainfall when rainfall reaches above the upper threshold is insignificant.

Features of Weather Effects

To intuitively see the impact of weather on corn yield, we plot the posterior mean of corn yield against temperature and rainfall, state by state, in figure 1 and figure 2. For the plot of yield against temperature, we generate 500 values of temperature evenly distributed over the observed range of temperature. We then evaluate model (1) with the time variable being 28 (year 2008), the rainfall variable being the state average, and the temperature variable being each one of the 500 generated values. For each temperature value, we simulate 15,000 corn yield draws using the simulated coefficients from the 15,000 MCMC iterations. To represent corn yield at the state average level, we use α_0 as the intercept term of the linear trend and use the average of CRD-specific trend slopes to calculate the trend yield. For each of the 500 temperature values, we calculate the posterior mean of corn yield from the 15,000 simulations. We then plot the 500 posterior means of corn yield against the 500 values of temperature. We also plot the posterior mean of corn yield against rainfall in a similar way except that we instead use the state average temperature and 500 values of rainfall.

The plots reveal several features of the impacts of weather on corn yield. First, the impacts of both weather variables are nonlinear in most states. The exceptions are that corn yield is almost linear in temperature in Indiana and Missouri. If a linear specification were applied to estimate the weather impacts, the estimated slope would be the weighted average of slopes of the three linear splines in our specification, with frequencies of the occurrences in each weather regime as weights. Given that the distribution of temperature

is almost symmetric and that the distribution of rainfall is positively skewed, the average impact is likely to be negative for temperature and positive for rainfall, which is consistent with findings in literature. As the plots indicate, the underlining weather impacts actually vary across weather regimes. Thus, our specification is a more comprehensive and precise way to estimate weather impacts. Secondly, the impact is asymmetric in most cases. Particularly, an excessive amount of heat causes more damage to corn yield than does a moderate amount of heat benefits corn yield in the northern states. Similarly, corn yield is more responsive to a lack of rainfall than to excessive rainfall in Ohio, Illinois, Indiana, and Michigan. And there are several plots that indicate yield could be monotone in the weather variable. Thus, a quadratic specification would be too restrictive. With endogenous thresholds, the linear spline specification is reasonably flexible and simple. Thirdly, corn yield is concave in both rainfall and temperature. Consider lack of heat as a favorable input to corn growth, then as this input increases (as temperature decreases), the marginal benefit from the favorable weather input decreases. Similarly, the marginal benefit of rainfall decreases as rainfall increases. The notion of decreasing marginal benefit, which is valid for other agricultural inputs, is true for the weather input, arguably the most important production input for crops. Finally, there is not a universal pattern of how weather variables affect corn yield. Rather, there exist several patterns, depending on the geographical location of the state. There are three patterns of the temperature effect. In Illinois, Iowa, and Ohio, the temperature effect is always negative, but with a flatter slope below the threshold at around 73 degrees Fahrenheit and a steeper slope above the threshold. In Indiana and Missouri, the temperature effect is negative and linear. In northern states, including Michigan, Minnesota, and Wisconsin, corn yield increases as temperature increases until a threshold, and then decreases. Patterns of rainfall fall into two general categories. In both categories, corn yield increases sharply as the amount of rainfall increases until it reaches the lower threshold. The difference between the two categories is how yield responds to rainfall above the threshold. In Illinois, Indiana, Michigan, and Ohio, corn yield becomes insensitive

to rainfall above the threshold. In Minnesota, Missouri, Iowa, and Wisconsin, corn yield decreases sharply when the amount of rainfall reaches beyond the upper threshold.

The plots also reflect three main sources of yield losses. One is excessive heat, that is, when temperature reaches beyond a threshold. This threshold is about 73, 75.5, 66.5, 72, 72.5, 70.5 degrees Fahrenheit in Illinois, Iowa, Michigan, Minnesota, Ohio, and Wisconsin respectively. The second cause is a lack of rainfall, that is, when the amount of rainfall is less than the lower threshold, which is between 2.5 to 4 inches. Flooding could also induce an equally large yield loss in Iowa, Minnesota, Missouri and Wisconsin.

The One-knot Specification

The plot of corn yield against temperature in Indiana raises the issue that the two thresholds might be too close to facilitate valid estimation. We instead estimate a yield model with a one-knot linear spline function of temperature and a two-knot linear spline function of rainfall using data in Indiana. Estimation results are presented in table 5.

Trend parameters and variance parameters are similar to the two-knot specification results. The rainfall thresholds and marginal effects of rainfall are also similar to the two-knot specification results. The lower threshold of rainfall is about 0.5 inches below mean (about 3.6 inches in absolute value). The upper threshold of rainfall is about 1.3 inches above mean (about 5.4 inches in absolute value). When the amount of rainfall is within the thresholds, the marginal effect of rainfall is insignificant. The differences of the marginal effect of rainfall in both dry condition and flooding condition are statistically different from the marginal effect of rainfall in normal weather condition. In particular, when the amount of rainfall is below the lower threshold, one inch of rainfall increases corn yield by 14 bushels per acre. On the contrary, about 4 bushels per acre corn yield is lost as the amount of rainfall increases every inch beyond the upper threshold. The threshold of temperature is around the mean, which is about 73.6 degrees Fahrenheit. The marginal effect of temperature is negative and significant both below and above the threshold. However, there is a statistically

significant difference in marginal effects in the two temperature regimes. The marginal damage is about 3 bushels per acre when temperature is below the threshold and is about 5 bushels per acre when temperature is above the threshold. The overall weather impact on corn yield in Indiana is revealed in the middle two subplots in figure 3.

As pointed out above, the upper threshold of temperature is redundant for Illinois, Michigan, Missouri, Ohio, and Wisconsin. Will inclusion of the redundant knot affect our estimation? The answer is no. We estimated the one-knot specification model and compared results to the two-knot specification. The estimated weather effects are very similar. For example, we plot the posterior mean of corn yield against weather variables based on the one-knot specification for Illinois and Michigan in figure 3. There is no recognizable difference between the corresponding subplots for these two states in figure 1 and figure 3.

Conclusions

Using a sampling-based Bayesian approach, we estimate a yield model with a hierarchical structure. The linear trend has a county-specific random effect and a CRD-specific slope. Weather impacts are captured by flexible linear-spline functions with endogenous thresholds. We find the impacts of rainfall and temperature on corn yields to be non-linear and asymmetric in most states in the Corn Belt. The temperature effect is linear in only two out of eight states. Rainfall effect is non-linear in all states. Weather impacts are also asymmetric. An excessive amount of heat causes more damage to corn yields than a moderate amount of heat benefits corn yields in the northern states. Corn yield is more responsive to droughts than to flooding in Ohio, Illinois, Indiana, and Michigan.

In general, corn yield decreases sharply as temperature increases above a threshold, although estimated thresholds vary across states. A moderate amount of heat is beneficial to corn yield in northern states. Drought is a big threat to corn yield in all states. Below a threshold of about 2.5 to 4 inches, the marginal benefit of increasing the amount of rainfall is large. In some states, corn yield stays insensitive to rainfall once the amount of rainfall

reaches a certain level. But in other states, an excess amount of rainfall (more than 5-6 inches) causes severe losses. Patterns of weather impacts tend to cluster in geographically neighboring areas. Temperature effects differ between the north and the south, while rainfall effects differ between the east and the west. Universally, corn yield is concave in both weather variables, which is consistent with the notion of decreasing marginal benefit of good weather.

Table 1. Posterior Mean and Standard Deviation of Trend and Variance Parameters

	IL	IN	IA	MI	MN	MO	OH	WI
$\beta_{1,1}$	2.43 (0.09)	1.93 (0.11)	2.42 (0.10)		2.37 (0.22)	1.74 (0.12)	1.44 (0.10)	0.74 (0.12)
$\beta_{1,2}$	1.87 (0.12)	1.85 (0.11)	1.97 (0.10)	1.05 (0.21)		1.69 (0.14)	1.39 (0.11)	0.98 (0.14)
$\beta_{1,3}$	2.47 (0.11)	1.59 (0.11)	2.08 (0.10)	1.71 (0.21)		1.61 (0.15)	0.88 (0.12)	0.63 (0.16)
$\beta_{1,4}$	2.31 (0.10)	1.65 (0.11)	2.29 (0.09)		2.80 (0.12)	1.83 (0.16)	1.62 (0.11)	1.25 (0.11)
$\beta_{1,5}$	2.15 (0.12)	1.61 (0.08)	2.40 (0.09)	1.28 (0.15)	2.09 (0.11)	1.53 (0.12)	1.57 (0.10)	1.03 (0.13)
$\beta_{1,6}$	1.84 (0.11)	1.58 (0.14)	2.14 (0.10)	1.68 (0.12)	1.25 (0.14)	1.63 (0.16)	0.16 (0.19)	0.74 (0.11)
$\beta_{1,7}$	1.82 (0.09)	1.78 (0.09)	1.99 (0.11)	1.51 (0.12)	2.75 (0.14)	2.00 (0.20)	1.45 (0.11)	1.39 (0.12)
$\beta_{1,8}$	1.77 (0.12)	1.39 (0.12)	1.89 (0.10)	1.45 (0.09)	2.14 (0.13)	2.37 (0.40)	0.84 (0.32)	1.63 (0.13)
$\beta_{1,9}$	1.63 (0.12)	1.24 (0.21)	1.92 (0.10)	1.79 (0.10)	2.24 (0.13)	2.19 (0.18)	0.73 (0.12)	1.09 (0.14)
α_0	104.76 (3.12)	119.33 (6.28)	109.90 (3.10)	79.12 (4.91)	88.71 (4.04)	77.11 (6.20)	105.99 (5.23)	95.50 (8.81)
σ_α^2	180.69 (30.64)	111.02 (21.36)	77.31 (13.43)	243.85 (59.21)	476.99 (95.85)	273.82 (54.95)	202.47 (39.01)	249.87 (50.75)
σ_ε^2	231.81 (6.92)	206.65 (6.79)	253.67 (6.90)	179.96 (7.87)	295.39 (10.81)	355.59 (12.44)	222.87 (7.52)	220.81 (7.88)

Note: Standard errors are in parenthesis.

Table 2. Posterior Mean and Standard Deviation of Parameters of Weather Impacts

	IL	IN	IA	MI	MN	MO	OH	WI
β_2	4.19 (1.54)	8.16 (2.85)	3.32 (0.94)	4.49 (2.20)	4.05 (1.27)	0.07 (3.51)	5.17 (2.92)	8.11 (2.89)
β_3	-5.88 (1.40)	-11.59 (2.89)	-5.39 (1.04)	-1.52 (2.18)	-1.58 (1.28)	-5.80 (2.81)	-7.82 (3.00)	-6.53 (2.97)
β_4	-1.04 (1.64)	7.62 (2.97)	-5.92 (1.73)	-0.05 (2.48)	-9.59 (1.63)	-1.10 (3.44)	-0.03 (3.37)	-1.70 (3.39)
β_5	11.60 (2.99)	11.33 (2.48)	12.60 (1.69)	9.62 (6.98)	10.97 (1.64)	5.35 (4.19)	11.13 (3.28)	10.00 (2.32)
β_6	7.04 (2.88)	2.28 (2.40)	1.89 (1.26)	3.13 (5.57)	-2.57 (1.39)	1.72 (3.89)	5.11 (3.41)	3.67 (1.78)
β_7	-8.42 (2.34)	-6.37 (2.25)	-16.67 (1.12)	-2.93 (5.79)	-10.85 (2.49)	-7.31 (3.34)	-5.50 (3.06)	-7.85 (1.94)
θ_l	-2.43 (0.63)	1.65 (0.29)	0.89 (0.77)	-2.30 (0.72)	-0.15 (0.53)	-1.11 (2.18)	1.08 (0.55)	2.18 (0.53)
θ_u	1.46 (1.56)	2.33 (0.38)	3.59 (0.35)	0.69 (2.03)	2.49 (0.36)	0.80 (2.29)	2.21 (0.72)	3.28 (0.51)
λ_l	-1.33 (0.45)	-0.48 (0.22)	-0.95 (0.28)	-0.20 (0.35)	-0.19 (0.28)	-0.18 (0.73)	-0.70 (0.27)	-0.73 (0.23)
λ_u	0.42 (0.55)	1.28 (0.65)	1.67 (0.21)	0.72 (0.58)	2.06 (0.23)	1.52 (0.87)	0.84 (0.83)	0.70 (0.29)

Note: Standard errors are in parenthesis.

Table 3. Comparing Thresholds Across States

	IL	IN	IA	MI	MN	MO	OH	WI
lower threshold of temperature	71.51	74.55	72.90	66.41	69.24	74.79	72.43	70.50
upper threshold of temperature	75.40	75.22	75.60	69.40	71.89	76.70	73.57	71.60
lower threshold of rainfall	2.55	3.62	3.49	3.18	3.87	4.04	3.22	3.40
upper threshold of rainfall	4.30	5.37	6.11	4.10	6.12	5.74	4.76	4.83

Table 4. Marginal Effects of Temperature and Rainfall

	IL	IN	IA	MI	MN	MO	OH	WI
$\beta_{temp,cl}$	-1.69 (0.87)	-3.42 (0.26)	-2.07 (0.31)	2.97 (1.06)	2.47 (0.48)	-5.73 (1.25)	-2.66 (0.41)	1.58 (0.29)
$\beta_{temp,nt}$	-5.88 (1.40)	-11.59 (2.89)	-5.39 (1.04)	-1.52 (2.18)	-1.58 (1.28)	-5.80 (2.81)	-7.82 (3.00)	-6.53 (2.97)
$\beta_{temp,ht}$	-6.92 (0.73)	-3.97 (0.90)	-11.32 (1.22)	-1.57 (1.01)	-11.17 (1.33)	-6.90 (1.16)	-7.86 (1.18)	-8.23 (1.16)
$\beta_{rain,dr}$	18.65 (3.92)	13.61 (1.34)	14.49 (2.22)	12.75 (2.21)	8.40 (1.73)	7.08 (1.32)	16.24 (2.00)	13.67 (1.95)
$\beta_{rain,nr}$	7.04 (2.88)	2.28 (2.40)	1.89 (1.26)	3.13 (5.57)	-2.57 (1.39)	1.72 (3.89)	5.11 (3.41)	3.67 (1.78)
$\beta_{rain,fl}$	-1.38 (1.00)	-4.09 (1.88)	-14.77 (0.83)	0.20 (1.75)	-13.42 (1.85)	-5.59 (1.63)	-0.39 (1.96)	-4.18 (1.00)

Note: Standard errors are in parenthesis.

Table 5. Estimation Results for the One-Knot Specification, Indiana

$\beta_{1,1}$	$\beta_{1,2}$	$\beta_{1,3}$	$\beta_{1,4}$	$\beta_{1,5}$	$\beta_{1,6}$	$\beta_{1,7}$	$\beta_{1,8}$	$\beta_{1,9}$
1.93 (0.11)	1.85 (0.11)	1.58 (0.11)	1.66 (0.11)	1.60 (0.08)	1.59 (0.13)	1.78 (0.09)	1.39 (0.12)	1.25 (0.21)
α_0	σ_α^2	σ_ε^2	β_2	β_3	β_5	β_6	β_7	
106.86 (1.89)	111.76 (21.81)	207.27 (6.76)	1.99 (0.53)	-5.31 (0.44)	11.40 (2.32)	2.28 (2.19)	-6.47 (2.06)	
θ	λ_l	λ_u	$\beta_{temp,cl}$	$\beta_{temp,nt}$	$\beta_{rain,dr}$	$\beta_{rain,nr}$	$\beta_{rain,fl}$	
0.75 (0.60)	-0.50 (0.22)	1.31 (0.57)	-3.32 (0.32)	-5.31 (0.44)	13.68 (1.25)	2.28 (2.19)	-4.19 (1.76)	

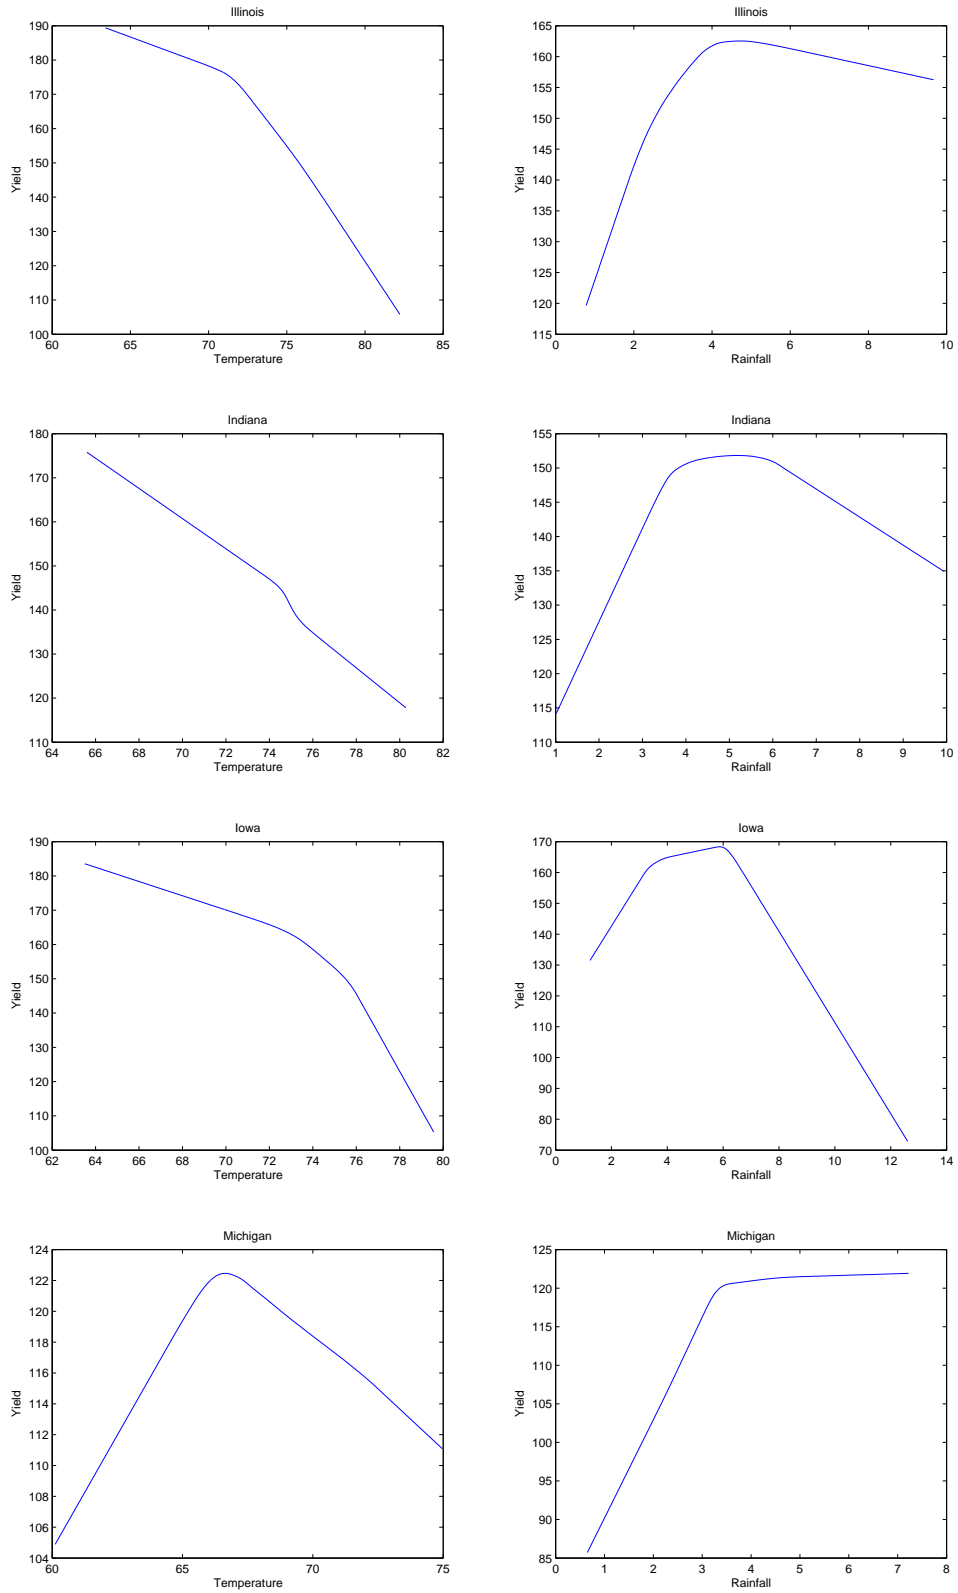


Figure 1. Plots of weather impacts on corn yield




Figure 2. Plots of weather impacts on corn yield - continued

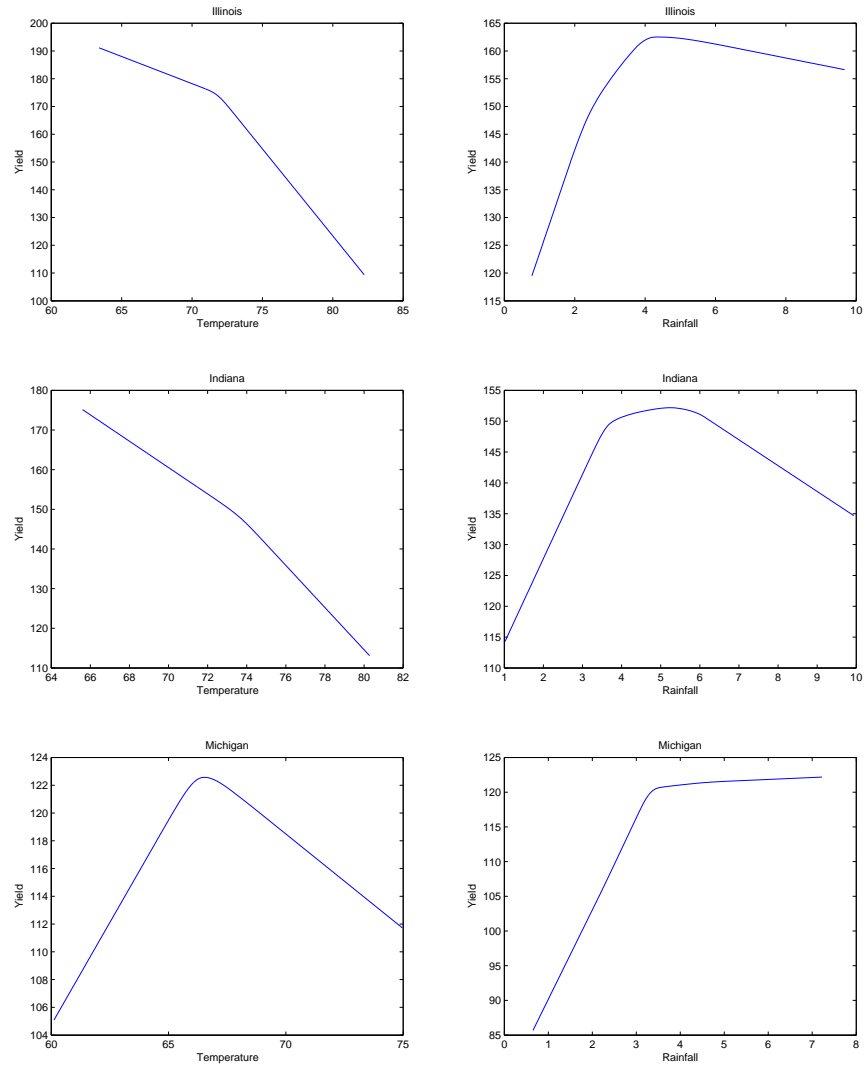


Figure 3. Plots for the one-knot specification

Endnotes:

¹We specify the growing season as from June to August.

²Results are insensitive to the specific values of the prior distribution parameters.

³ $\mathbf{1}$ is a column vector of ones. Its dimension is specified by its subscript.

⁴Less than 10% of the observations have missing values.

⁵A 0.3 acceptance rate is a common rule of thumb to achieve optimal mixing.

⁶We also estimated the model with one temperature threshold for these state and obtained similar results.

References

- Adams, R.M., C. Rosenzweig, R.M. Peart, J.T. Ritchie, B.A. McCarl, J.D. Glycer, R.B. Curry, J.W. Jones, K.J. Boote, and L.H. Allen, Jr. 1990. "Global Climate Change and U.S. Agriculture." *Nature* 345:219 – 224.
- Chib, S., and B. Carlin. 1999. "On MCMC sampling in hierarchical longitudinal models." *Journal Statistics and Computing* 9:17–26.
- Chib, S., and E. Greenberg. 1995. "Understanding the Metropolis-Hastings Algorithm." *The American Statistician* 49:327–335.
- Deschenes, O., and M. Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review* 97:354 – 385.
- Gelfand, A.E., S.E. Hills, A. Racine-Poon, and A.F.M. Smith. 1990. "Illustration of Bayesian Inference in Normal Data Models Using Gibbs Sampling." *Journal of the American Statistical Association* 85:972–985.
- Gelman, A., G.O. Roberst, and W.R. Gilks. 1995. "Efficient Metropolis Jumping Rules." *Bayesian Statistics* 5:599–607.
- Lindley, D.V., and A.F.M. Smith. 1972. "Bayes Estimates for the Linear Model." *Journal of the Royal Statistical Society. Series B (Methodological)* 34:1–41.
- Lobell, D.B., and G.P. Asner. 2003. "Climate and Management Contributions to Recent Trends in U.S. Agricultural Yields." *Science* 299:1032.
- O'Brien, D.M. 1993. "Forecasting the Probability Distribution of United States Harvest Time Average Corn Prices." PhD dissertation, Iowa State University.
- Schlenker, W., W.M. Hanemann, and A.C. Fisher. 2006. "The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions." *Review of Economics and Statistics* 88:113 – 125.
- Schlenker, W., and M.J. Roberts. 2006. "Nonlinear Effects of Weather on Corn Yields."

Review of Agricultural Economics 28:391–398.

Yu, T., and B.A. Babcock. 2010. “Are Corn and Soybean Becoming More Drought Tolerant?” *American Journal of Agricultural Economics* 92:1310–1323.