

Corn Production Shocks in 2012 and Beyond

Implications for Harvest Volatility

Steven T. Berry, Michael J. Roberts, and
Wolfram Schlenker

Historically, 25 percent of an average year's global corn production is held in inventories to buffer weather shocks and allow for a smooth consumption between years. As inventory levels are drawn down, prices increase, thereby giving farmers an incentive to increase production in the following years to refill depleted inventory levels.

While individual countries might face significant production shocks, these idiosyncratic shocks average out over the globe. Global corn production shocks (deviations from a trend) ranged from -13 percent to +7 percent in 1961 to 2010, with a standard deviation of 4 percent (Roberts and Schlenker 2013). International trade smoothes production shocks between countries unless these countries institute export bans.

There are, however, certain exceptions to this rule. The production of some crops is highly spatially correlated and subject to the same common weather shocks. A prime example is corn production in the United States, which is grown in the Midwest. Since the US produces roughly 40 percent of the world's corn, any impact to US production has the potential to significantly affect global production and global price levels.

Current and future corn price volatility depends directly on production shocks. One of the main drivers of production shocks are weather fluctua-

Steven T. Berry is the James Burrows Moffatt Professor of Economics at Yale University and a research associate of the National Bureau of Economic Research. Michael J. Roberts is associate professor of economics and an affiliate with Sea Grant at the University of Hawaii at Manoa. Wolfram Schlenker is associate professor at the School of International Public Affairs at Columbia University and a research associate of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c12806.ack>.

tions. An accurate model that translates weather fluctuations into production fluctuations is hence a crucial first step in examining food price volatility.

In this chapter we extend earlier work on the effects of weather on corn production (Schlenker and Roberts 2009). We previously allowed for a highly nonlinear effect of weather on corn yields, but assumed the effect of various temperatures to be *constant* throughout the growing season that we fixed to March through August. The main innovations of this chapter are: First, we allow the effect of various weather measures to evolve over the growing season. Second, we no longer keep the growing season fixed to March through August, but rather use annual state-level data on planting and harvest dates to capture weather measures over the actual growing season. Third, we predict yields for 2012 using the traditional as well as the new model. Since the 2012 heat wave happened during the part of the growing season when it is most harmful, the new model predicts larger production shortfalls. Fourth, we contrast 2012 to what is expected under climate change.

2.1 Model

We start by estimating a baseline model of yields that assumes a fixed growing season (March through August) and a constant effect of weather variables over the growing season. This baseline replicates a specification from earlier research (Schlenker and Roberts 2009). In a second step, following Ortiz-Bobea and Just (2013), we consider models that account for planting date and temperature effects that vary over the growing season.

2.1.1 Baseline Model 1

The baseline model relates log yield y_{it} in county i and year t to four weather variables:

$$(1) \quad y_{it} = \beta_1 m_{it} + \beta_2 h_{it} + \beta_3 p_{it} + \beta_4 p_{it}^2 + c_i + f_s(t) + \epsilon_{it},$$

where m_{it} is growing degree days between 10°C and 29°C, accounts the beneficial effects of moderate temperatures, h_{it} are degree days above 29°C that capture the damaging effect of extreme heat, and p_{it} and p_{it}^2 are season-total precipitation and its square.¹ County fixed effects c_i account for baseline differences between counties and state-specific time trends f_s account for

1. *Growing degree days* are based on cumulative heat exposure above a threshold temperature, which is sometimes also truncated by an upper bound. Degree days 10°C–29°C count all temperatures below 10°C as zero, temperatures between 10°C and 29°C as the difference between the observed temperature and 10°C, and temperatures at or above 29°C as 19. For example, twenty-four hours of exposure to a temperature of 11°C counts as one growing degree day while twenty-four hours of exposure to a temperature of 12°C counts as two degree days, and so on. In our weather data, we incorporate the entire distribution of temperatures between the daily minimum and maximum, thereby counting fractions of a day (see the data in section 2.2). Degree days above 29°C put the lower bound at 29°C and have no upper bound.

technological progress as average yields have been trending upward over time. Errors are clustered at the state level to adjust for spatial correlation.

The data underlying these regressions is constructed using daily fine-scaled weather measures on a 2.5 × 2.5 mile grid for the contiguous United States. We follow the same algorithm as Schlenker and Roberts (2009), but update the data through 2012. We use only counties east of the 100-degree meridian (excluding Florida) in the regression because the response function might be different for highly irrigated areas.² The data set spans the years 1950 to 2012.

2.1.2 Model 2: Time-Varying Parameters

Model 2 allows the effect of weather variables to vary over the growing season. Ortiz-Bobea and Just (2013) extend our earlier work by separating the growing season into three subintervals, and then estimate separate (constant) coefficients for each of the subintervals. This chapter allows the effect of weather variables to vary continuously over time. To make locations comparable, we use yearly data on planting and harvesting dates and normalize the season to have length 1. A value of 0.5 stands for the day that occurred in the middle of the growing season.

In a first step we only allow the coefficient β_2 that measures the effect of extreme heat to vary over the growing season. The reason is that extreme heat has consistently been found to have the largest influence on year-to-year variability of crop yields. There is agronomic evidence that heat matters especially during the flowering period, and the effect of weather measures might hence evolve over time. Model 2 is defined as:

$$(2) \quad y_{it} = \beta_1 m_{it} + g_2(h_{0it}, \dots, h_{D_{it}}) + \beta_3 p_{it} + \beta_4 p_{it}^2 + c_i + f_s(t) + \epsilon_{it}.$$

In the baseline model we summed daily degree days above 29°C over all days of the fixed growing season $\sum_{d=\text{March } 1}^{\text{August } 31} h_{dit}$, while $g_2()$ now allows the effect of h_{dit} to vary over the growing season. Note that we also no longer fix the growing season to March 1st through August 31st, but allow it to vary year to year. Different places might have different growing season lengths, and there is year-to-year variation in planting and harvesting dates at a given location. We define a growing season to last from planting (time 0) to harvest (time 1).

We construct a restricted cubic spline with k knots over the growing season, which will result in $k - 1$ spline variables $s_j()$. We consider models with between 3 and 7 knots, with the knots placed at standard fractions of the growing season.³ We normalize the growing season to length one, so the

2. Table 2.1 shows that these counties account for 91 percent of US production.

3. Spline knots locations are as follows: $k = 3$ indicates 3 knots set at 0.1, 0.5, and 0.9 fractions of the total growing season; $k = 4$ indicates knots set at 0.05, 0.35, 0.65, and 0.95; $k = 5$ spline knots set at 0.02, 0.26, 0.5, 0.74, and 0.98; $k = 6$ knots are set at 0.02, 0.212, 0.404, 0.596, 0.788, and 0.98; $k = 7$ knots are set at 0.02, 0.18, 0.34, 0.5, 0.66, 0.82, and 0.98.

“weighted” sum of daily degree days above 29°C (h_{dit}) over all days d of the growing season $d = 0, 1, 2 \dots D_{it}$ in county i in year t depends on the phase of the growing season $x_{dit} = (d - 1)/(D_{it} - 1)$.

$$g_2(h_{it}) = \sum_{d=0}^{D_{it}} h_{dit} \underbrace{\sum_{j=1}^k s_j(x_{dit})}_{\text{weight(time)}} = \sum_{j=1}^{k-1} \underbrace{\sum_{d=1}^{D_{it}} s_j(x_{dit}) h_{dit}}_{H_{jit}} = \sum_{j=1}^{k-1} \beta_{2j} H_{jit}$$

The second equality simply exchanges the order of summation. We are ultimately left with $j = 1 \dots k - 1$ variables H_{jit} , which are the sum of daily degree days above 29°C (h_{dit}) weighted by the value of the spline function $s_j(x_{dit})$ for each day (phase) of the growing season.

We also estimate an extended model that allows the effect of other weather variables to vary over the growing season. It includes a fifth variable, which is the interaction of daily degree days above 29°C and daily precipitation

$$(3) \quad y_{it} = g_1(m_{0it}, \dots, m_{D_{it}}) + g_2(h_{0it}, \dots, h_{D_{it}}) + g_3(p_{0it}, \dots, p_{D_{it}}) \\ + g_4(p_{0it}^2, \dots, p_{D_{it}}^2) + g_5(h_{0it} \times p_{0it}, \dots, h_{D_{it}} \times p_{D_{it}}) \\ + c_i + f_s(t) + \varepsilon_{it}.$$

Besides the time-varying effect of additional weather variables, the extended model differs in another aspect: earlier models use season-total precipitation and season-total precipitation squared. The extended model uses *daily* precipitation as well as *daily* precipitation squared, and allows the effects of these variables to vary over the growing season.

2.2 Data

We pair data on annual county-level corn yields with fine scaled-weather measures that were constructed on a 2.5 × 2.5 mile grid for the entire United States. We follow the same algorithm of Schlenker and Roberts (2009), but update the data through 2012. These data give daily minimum and maximum temperature, as well as precipitation for each grid cell. Degree days above a threshold b are calculated by fitting a sine-curve between the daily minimum and maximum temperature in each cell and integrating over the difference between the temperature curve and the threshold (Snyder 1985). Daily weather measures for all grids in a county are weighted averages, where the weights are the cropland area in each grid cell that were obtained from a satellite scan. This gives daily weather measures for each county.

In the baseline model, we sum degree days over all days of the growing season, which was fixed to March 1st through August 31st for all counties and years. These variables were calculated using all counties east of the 100-degree meridian (excluding Florida). The second row of table 2.1 displays the fraction of the US growing area and production that falls in these countries for the three most recent years before the heat wave occurred; that

Table 2.1 Descriptive statistics of county samples

	Production		Area harvested		Yield	
	Billion bushels	Percent of US total (%)	Million acres	Percent of US total (%)	Bushel per acre	Percent of US total (%)
US Total	12.63		81.64		154.73	
Eastern counties	11.44	90.53	73.28	89.76	156.07	100.87
Planting dates	10.95	86.72	69.49	85.12	157.63	101.87

Notes: Table summarizes the subsets of counties used in this study. Data are given for the three years before the 2012 heat wave; that is, 2009 to 2011. Eastern counties are all counties east of the 100-degree meridian except Florida. Counties with planting dates are all eastern counties where state-level planting dates are available.

is, 2009 to 2011. With approximately 90 percent of the growing area and total production, these counties account for the largest share of US corn production. Given this large coverage, average yields from these counties closely match overall US yields as shown in the last two columns of the table.

2.2.1 Weather Anomalies for a Fixed Growing Season (Model 1)

Weather measures for counties east of the 100-degree meridian (excluding Florida) that grow corn are displayed on the maps in figure 2.1. The top graph shows 2012 anomalies of season-total degree days above 29°C for a fixed growing season of March through August; that is, the difference between 2012 and the average from 1950 to 2011. The bottom graph shows the 2012 anomalies for season-total precipitation. There is a lot of heterogeneity across counties, with some counties experiencing above normal conditions while others experience below-normal conditions for both weather variables. The Corn Belt was hotter and drier than usual, while southern counties had a cooler and wetter than average year. Note the variation in extreme heat: some highly productive counties in the Corn Belt experienced up to 100 extra degree days above 29°C. As we show below, each degree day above 29°C reduces log yields by 0.006, so the effect of an extra 100 degree day above 29°C is a decrease of 60 log points.

For comparison, the production-weighted average exposure to degree days above 29°C is 33 among all eastern counties in 1950 to 2011. Since bad weather in highly productive areas can cause a loss that is not compensated by better-than-average weather in less productive areas, we summarize weather outcomes by constructing the production-weighted average of all eastern counties. Production weights are the product of actual area (which is known at the beginning of the season) and predicted yields according to a trend.⁴

4. We fit a restricted cubic spline with 3 knots to the yield history of each county.

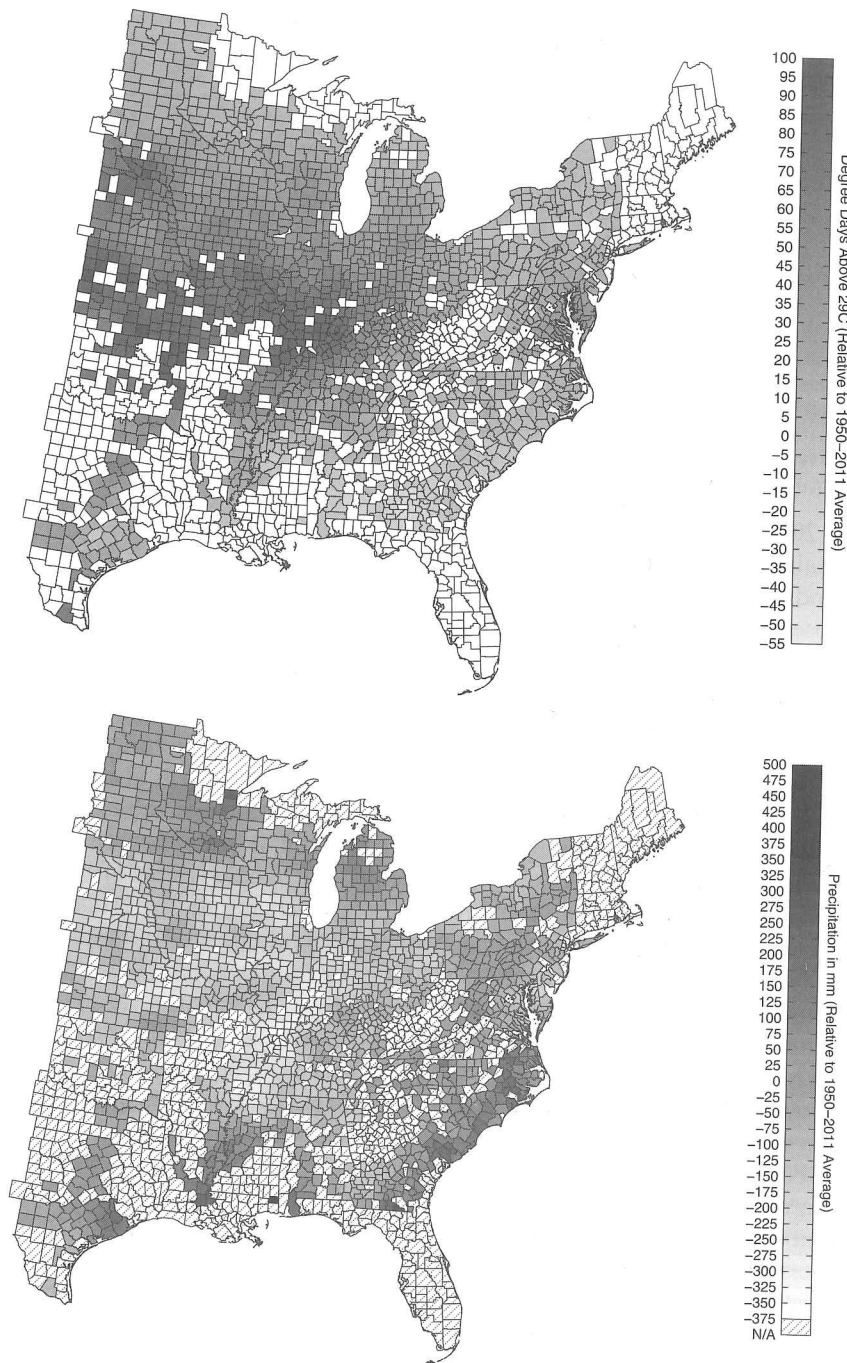


Fig.2.1 Spatial distribution of degree days above 29°C and precipitation in 2012
Notes: Spatial distribution of weather anomalies over the fixed 2012 growing season (March–August). Top panel shows degree days above 29°C, while the bottom panel shows precipitation totals.

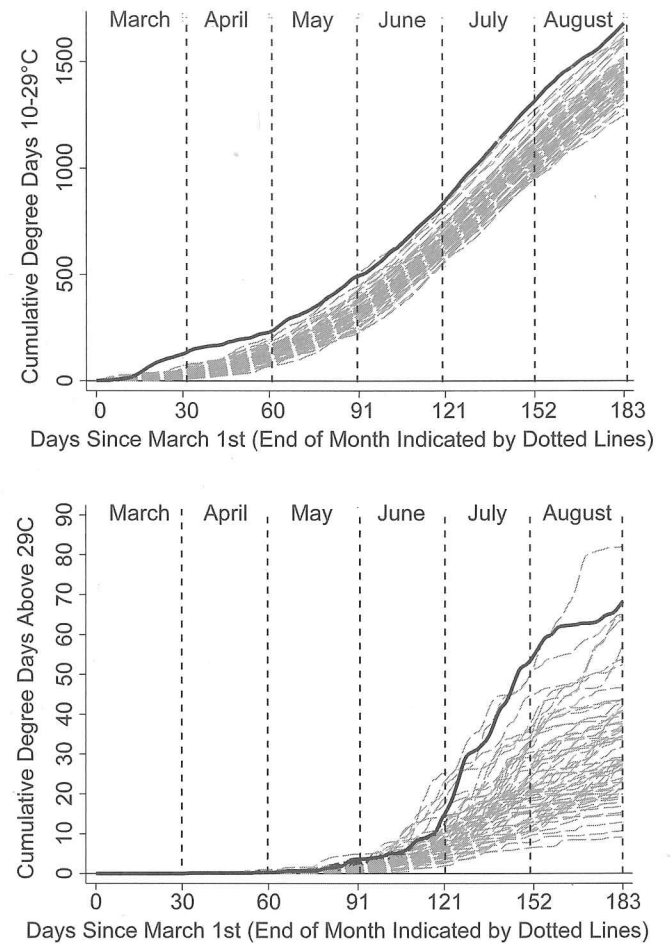


Fig. 2.2 Degree days 10°C–29°C and degree days above 29°C in 2012 relative to 1950–2011

Notes: Panels show cumulative total of degree days 10°C–29°C and cumulative total of degree days above 29°C for the eastern United States except Florida. Weather measures are the weighted average of all counties east of the 100-degree meridian excluding Florida, where the weights are predicted yields along a trend line (restricted cubic spline with 3 knots) times the actual growing area. Cumulative totals for the years 1950 to 2011 are added as thin dashed lines, while 2012 is shown as a thick solid line.

Figure 2.2 shows the evolution of the cumulative season total degree days measures over the 184 days of the growing season, ranging from March 1st (day 0) to August 31st (day 183). We average cumulative season totals up to a given day of the growing season. Historic exposures for the years 1950 to 2011 are shown as gray dashed lines, while the outcome for 2012 is shown as a thick solid line.

The top panel of figure 2.2 shows degree days 10°C – 29°C . Degree days above 10°C – 29°C start to increase earlier than usual in 2012, since the United States had a warm spring. The beneficial side effect of a warmer spring is that it allows for earlier planting. The bottom graph of figure 2.2 shows degree days above 29°C . July and August are traditionally the months where temperatures climb above 29°C most frequently and degree days above 29°C increase most rapidly. July 2012 was exceptionally hot by historic standards. At the beginning of July, the measures were slightly above normal, but by the end of July, it had superseded the hottest year among the 1950 to 2011 historic baseline, which was 1988. Note that 1988 had a hotter August than 2012, and as a result the season total degree days above 29°C was highest in 1988, followed by 2012.

The top graph of figure 2.3 displays the cumulative season-total precipitation. Precipitation was below normal in 2012, and the only year with drier conditions in the 1950 to 2011 historic baseline is again 1988. Note, however, that the relative deviation from the mean is much lower for precipitation than for degree days above 29°C . Finally, the bottom graph of figure 2.3 shows cumulative vapor pressure deficit, which is the difference between how much water the air can hold when it is saturated and how much water is currently in the air. This measure is used in agronomic crop models and has also been shown to predict yields in a statistical model (Roberts, Schlenker, and Eyer 2013). Similar to precipitation, this measure indicates that crops were adversely affected (a higher than usual deficit is bad for crops), yet the relative deviation from the mean was less than for degree days above 29°C .

2.2.2 Planting and Harvest Dates (Model 2)

The second model relaxes two assumptions: first, we no longer fix the growing season to March through August, but instead used data from the National Agricultural Statistics Service (NASS) on planting and harvesting dates. The NASS reports on a weekly level what fraction of the corn area in major corn-producing states was planted and harvested. We define the beginning of the growing season as the Monday of the week by the end of which at least 50 percent of the corn area in a state had been planted. Similarly, the end of the growing season is the last day of a week when at least 50 percent of the growing area had been harvested in a state.

The average planting date for each county is shown in the top graph of figure 2.4. Southern places tend to plant earlier, as they are not limited by the probability of late freezes. Northern places also have a larger intrayear cycle in solar radiation, which is an important component of crop growth that limits farmers from shifting the planting date too far forward. We do not fix the growing season in each place but allow it to vary between years according to annual NASS reports. In case only the planting date is available for a state, but not the harvest date, we approximate the harvest date by adding the average growing season length to the reported planting date. By the

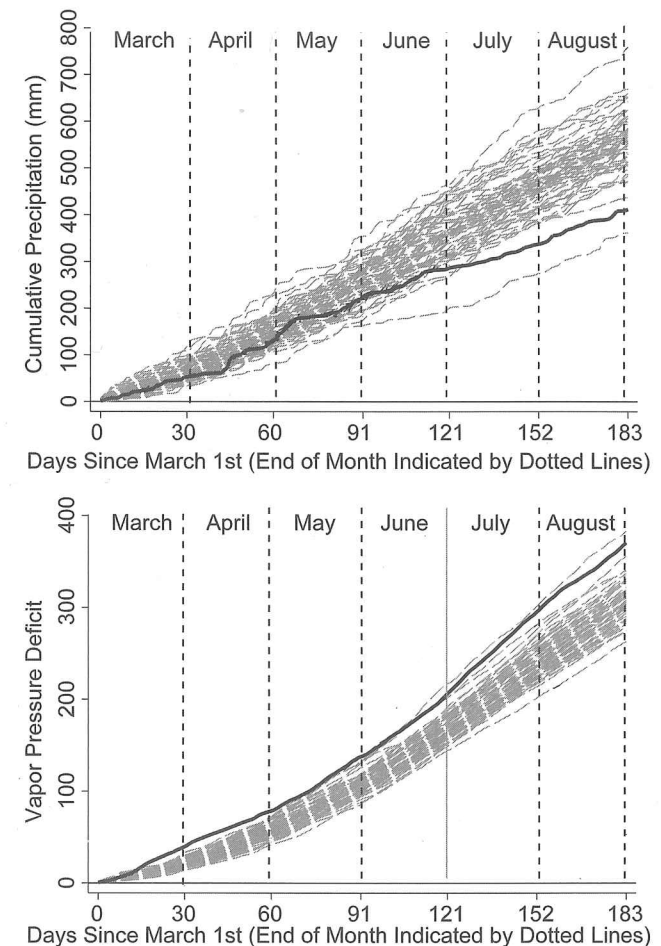


Fig. 2.3 Precipitation and vapor pressure deficit in 2012 relative to 1950–2011

Notes: Panels show precipitation and vapor pressure deficit for eastern United States except Florida. Weather measures are the weighted average of all counties east of the 100-degree meridian excluding Florida, where the weights are predicted yields along a trend line (restricted cubic spline with 3 knots) times the actual growing area. Cumulative totals for the years 1950 to 2011 are added as thin dashed lines, while 2012 is shown as a thick solid line.

same token, if the harvest date is reported but the planting date is missing, we approximate the latter by subtracting the average growing season length from the harvest date.

Southern places have a longer growing season as shown in the bottom graph of figure 2.4. As mentioned above, we make the different growing seasons comparable by rescaling them such that the first day is 0, while the last day is 1. After fitting spline polynomials over the season, we aggregate the variables to an annual level.

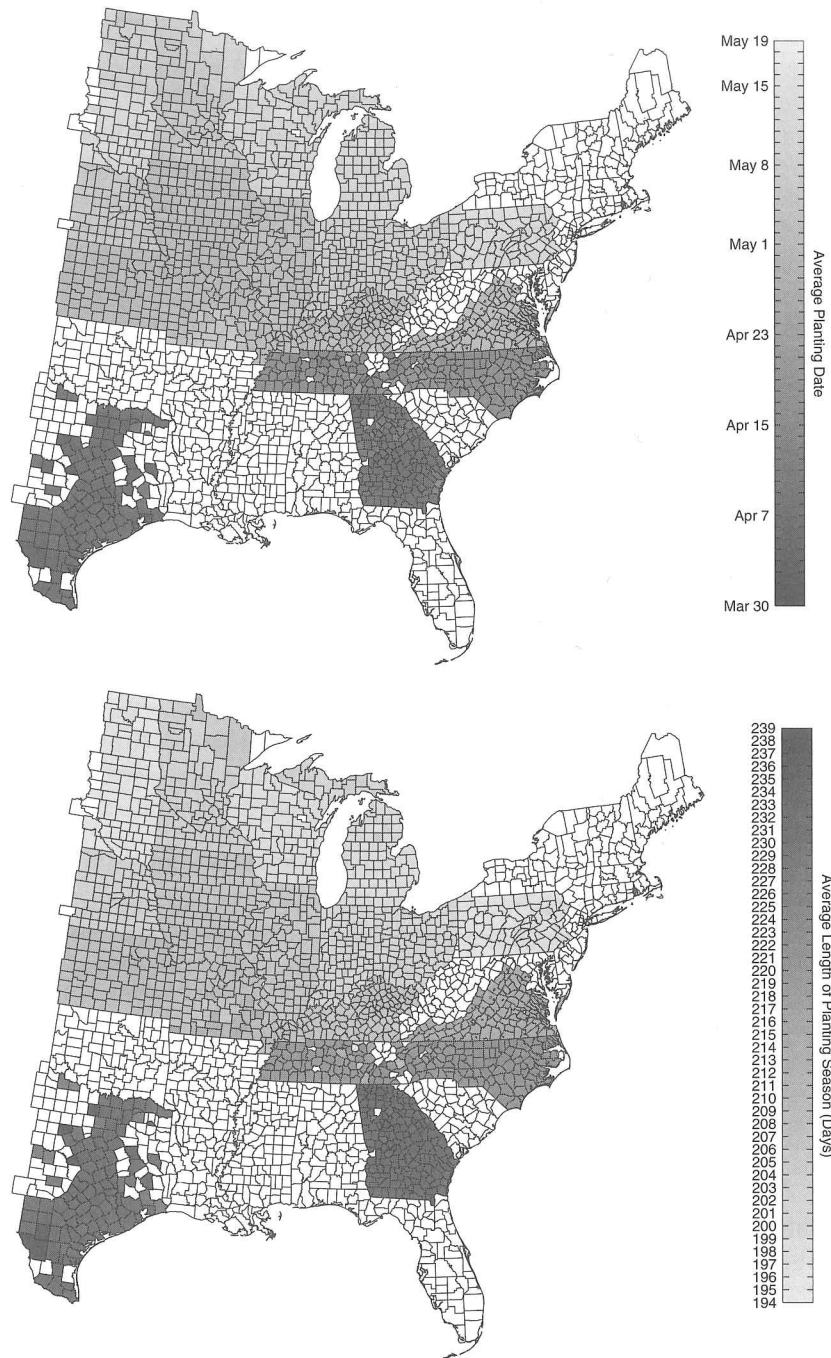


Fig. 2.4 Average planting date and growing season length (1979–2011)

Notes: Top graph shows average planting date in 1979 to 2011, while bottom graph shows average growing season length. Both planting dates and growing season length are reported annually for each state. Counties within each state might have different values because they grew corn in different years.




Fig. 2.5 Counties with yearly state-level data on planting dates (1979–2011)

Notes: Figure displays counties for which annual planting and/or harvesting dates as well as yields were reported. Counties are shaded by the number of yearly observations that are available for 1979 to 2011. Planting and harvest dates are reported on a state level, while yields are reported for each county. The number of observations can differ within a state because yields are not reported for all counties in a state.

The number of yearly observations for which we have yield and planting data in a county is shown in figure 2.5. The first year in which data on planting and harvesting dates is available is 1979, but many states started to report planting dates at a later time. Counties in the eastern United States (excluding Florida) that report planting and/or harvest dates are summarized in the third row of table 2.1. States that report planting dates account for 85 percent of the corn growing area and 87 percent of the US corn production in the most recent three years before the 2012 heat wave (2009 to 2011).

The second innovation of model 2 is to relax the assumption that the effect of some, and eventually all, weather variables are constant over the growing season. As outlined in section 2.1, we interact daily measures of the weather variables with spline polynomials. This allows the effect to differ over the growing season in a flexible way.

Figure 2.6 displays the average daily exposure over the growing season for four weather variables: degree days 10°C–29°C, degree days above 29°C, precipitation, and vapor pressure deficit. We use either restricted cubic splines

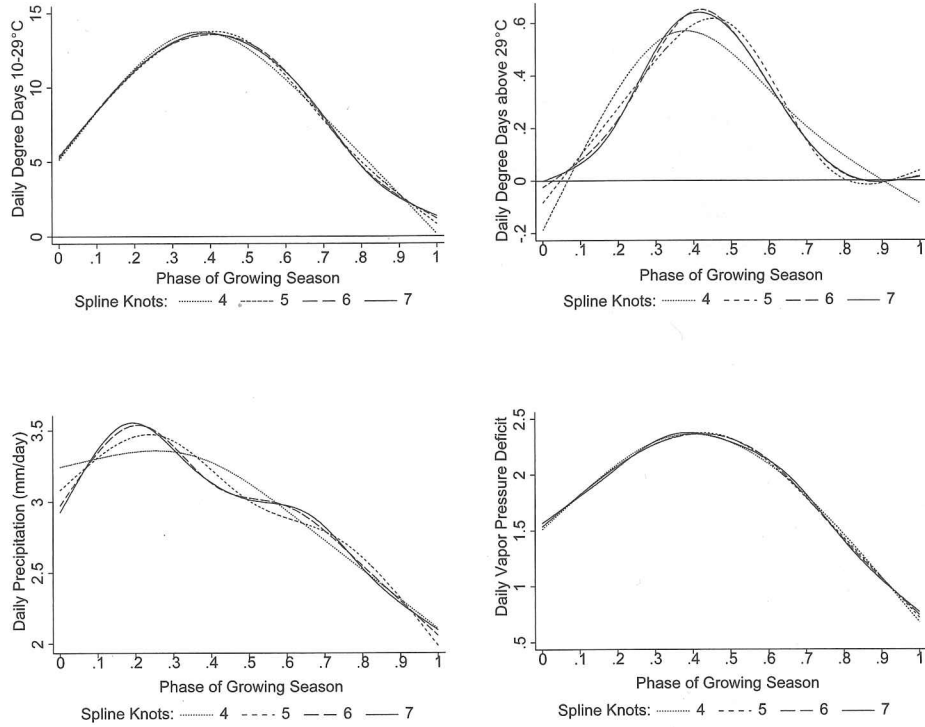


Fig. 2.6 Exposure to various weather variables over the growing season

Notes: Panels show the average exposure to various weather measures over the growing season. We use year- and state-specific estimates of the National Agricultural Statistical Service (NASS) to define the growing season: the week in which the planted area exceeds 50 percent is the start (x -value of 0) and the week that the harvested area exceeds 50 percent is the end (x -value of 1). Daily values are smoothed using restricted cubic splines with 4, 5, 6, or 7 knots.

with 4, 5, 6, or 7 knots. The results seem fairly stable as long as we include at least 5 knots.

2.3 Empirical Results

We start by replicating the results for a fixed growing season (March through August) that assume constant marginal effects of the weather variables before relaxing both assumptions.

2.3.1 Baseline Model 1

Results for a panel analysis for eastern counties (excluding Florida) for the years 1950 to 2011 is given in table 2.2. All columns use the same set of observations, but vary the set of time controls that are used to capture overall trends in yields. Columns (a), (b), and (c) use state-specific restricted cubic

Table 2.2 The effect of weather on maize yields using a fixed growing season, March through August

	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
Thousand degree days 10–29°C	0.314*** (0.068)	0.301*** (0.066)	0.303*** (0.063)	0.346*** (0.082)	0.343*** (0.079)	0.361*** (0.074)
Hundred degree days above 29°C	-0.622*** (0.068)	-0.616*** (0.066)	-0.625*** (0.065)	-0.580*** (0.069)	-0.583*** (0.068)	-0.584*** (0.067)
Precipitation (m)	1.028*** (0.212)	1.016*** (0.208)	1.029*** (0.198)	1.092*** (0.217)	1.061*** (0.216)	1.095*** (0.209)
Precipitation (m) squared	-0.806*** (0.165)	-0.800*** (0.160)	-0.818*** (0.153)	-0.807*** (0.163)	-0.787*** (0.161)	-0.814*** (0.156)
Total production impact (%)	-14.43	-14.65	-14.74	-13.29	-13.46	-13.12
RMSE—2012 county prediction	0.3250	0.3442	0.3437	0.3514	0.3476	0.4732
Pred. error total prod. 2012 (%)	3.36	8.82	10.22	10.86	9.96	36.89
R ²	0.7734	0.7784	0.7810	0.7920	0.7955	0.7972
Observations	115,205	115,205	115,205	115,205	115,205	115,205
Counties	2,276	2,276	2,276	2,276	2,276	2,276
Spline knots	3	4	5	3	4	5
Year fixed effects	No	No	No	Yes	Yes	Yes

Notes: Table regresses log maize yields for counties east of the 100-degree meridian (except Florida) in the years 1950 to 2011 on four weather variables as well as time controls. Columns (a), (b), and (c) include state-specific restricted cubic splines in time with 3, 4, and 5 knots, respectively, as time controls. The last three columns additionally include year fixed effects. Panel B gives the predicted production shortfall below trend from the 2012 weather outcomes in percentage points. Panel C compares prediction for 2012 to actual observed yields. The first row shows the root mean squared prediction error of all county-level log yields, while the second row gives the prediction error of total production for all counties combined. Errors are clustered at the state level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

splines with 3, 4, and 5 knots, respectively. On top of that, columns (2a), (2b), and (2c) also include year-fixed effects to capture overall shocks, like changes in global food prices or technological breakthroughs. The results are very stable across specifications. Moderate heat (degree days 10°C – 29°C) is beneficial, while extreme heat (degree days above 29°C) is highly damaging. Note that 2,000 degree days 10°C – 29°C increase expected yields by as much as 100 degree days above 29°C decrease them. Moreover, a coefficient of -0.6 on degree days above 29°C implies that the 100 additional degree days are lowering expected yields by 60 log points. Recall that several counties in the Corn Belt experienced heat anomalies of that magnitude in 2012 (see figure 2.1). Finally, precipitation and precipitation squared suggest that the relationship is hill-shaped (both too little and too much rain are harmful). The optimum is around 0.63m, or 25 inches, which matches closely the estimate of optimal rainfall from agronomic studies.⁵

The effect of the 2012 weather outcomes on expected yields are shown in figure 2.7. The top graph depicts predicted deviations from the time trend in log points (using the specification from column [1a] in table 2.2). There is significant heterogeneity: some counties are predicted to be as much as 56 percent below normal, while others experience yields up to 32 percent above normal. Unfortunately, yield declines are concentrated in the more productive areas. The bottom graph of the figure does not show relative impacts, but predicted *total impacts*. We multiply the observed harvest area in 2012 by the predicted production shortfall per area.⁶ While northern and southern areas experience small absolute increases, counties of the Corn Belt are predicted to experience large declines. The overall impact for our sample is a 14.4 percent production shortfall below trend as shown in column (1a), panel B, in table 2.2.

The observed yields in 2012 have been published after an earlier version of this chapter gave our predicted production shortfalls. Panel C therefore compares the how well our prediction compares to the actual observed yields in 2012. The first row gives the root mean squared error, which is the square root of the sum of the squared difference between the prediction in each county and the observed outcomes in 2013. The second row derives the percent error when predicted total production for all counties in the sample is compared against the observed outcome for 2012. All numbers are positive, suggesting that our model overpredicted yields, or underpredicted the damaging effects of extreme heat. Note that the error on predicted production is

5. Ozone pollution is correlated with high temperatures and one might wonder whether the coefficient on extreme temperatures captures the reduced form effect of both temperature and ozone. Boone, Schlenker, and Siikamäki (2013) estimate a model that includes both degree days above 29°C as well as various ozone measures. While ozone is very damaging for maize yields, its inclusion only slightly changes the coefficient on degree days above 29°C as the latter is a highly nonlinear transformation of temperature and hence not directly related to ozone exposure.

6. We obtain similar results if we instead use the average harvest area for the previous three years 2009 to 2011.

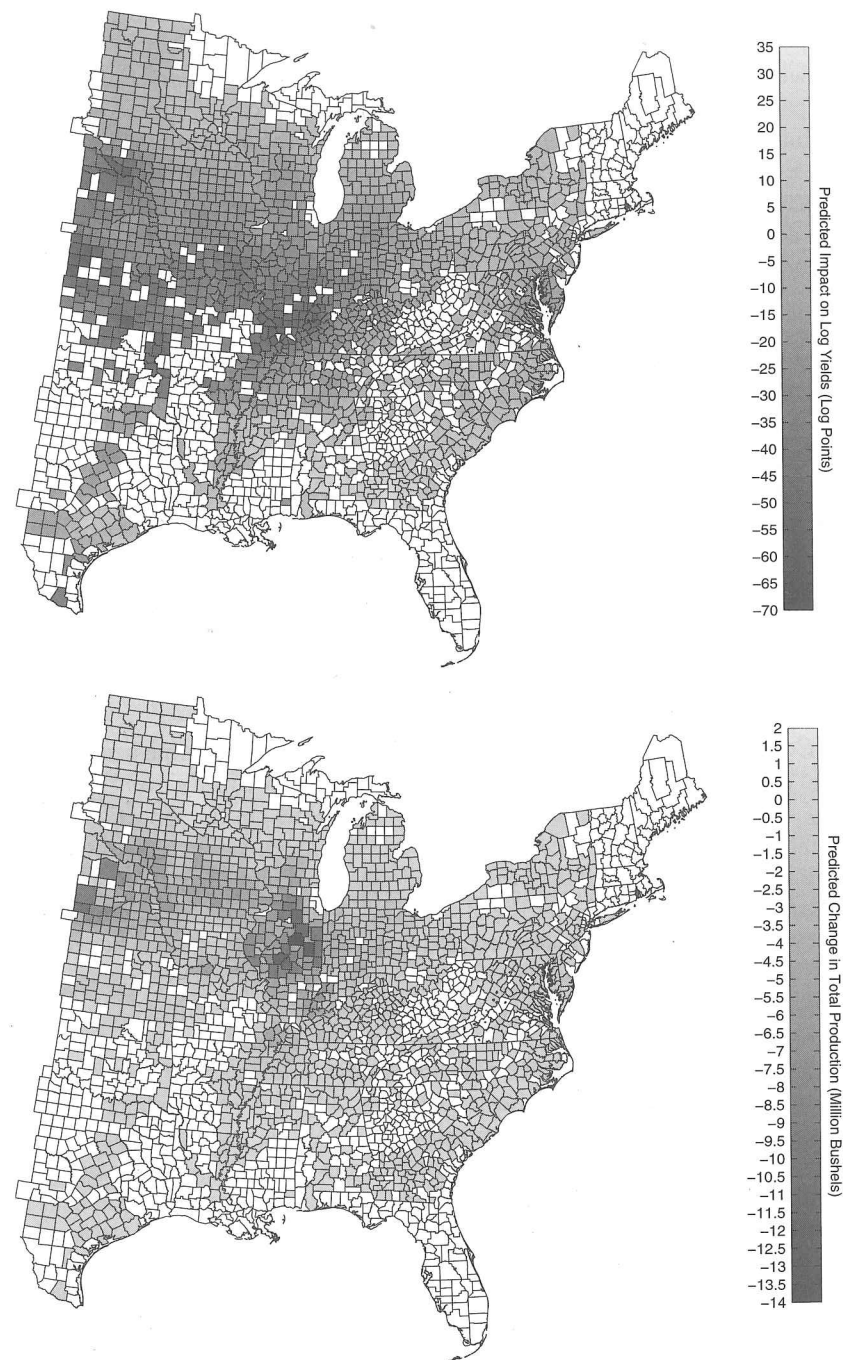


Fig. 2.7 Predicted yields and production in 2012

Notes: Predicted yield and production impacts in 2012 by county using the regression specification in column (1a) of table 2.2. The top panel shows changes in predicted yields in log points, while the bottom shows predicted changes in total production (using the average area of 2009 to 2011 as growing area). Total predicted production in the shown counties was 11.4 billion bushels, and the production shortfall was 1.7 billion bushels, or 15 percent.

quiet large if we use a fifth-order time polynomial, but not for the weather impacts. While the predicted production impact is comparable among all columns in panel B, the error on predicted total production is large when we use more flexible time trends. The reason is that the predicted trend is badly estimated for years outside the range observed in the data in a flexible model, which uses the last few years of observed data to interpolate the trend out of sample.

2.3.2 Model 2: Time-Varying Growing Season and Parameters

When we allow the effect of weather variables to vary over the growing season, we have to restrict the data set to a smaller set of counties for which annual planting and harvest dates are available. Column (1) of table 2.3, therefore, still forces the effect of each weather variable to be constant over the growing season, but runs the regression on the subset of counties for which planting dates are available and uses the weather measures when they are averaged over the actual growing season (instead of March through August). The coefficient on the two degree days variables remain rather unchanged. Panel C summarizes the predicted decrease in total production from the observed 2012 weather outcomes, which is 18.5 percent in column (1), that is, larger in magnitude than what we had observed for the bigger sample in column (1a) of table 2.2.

Columns (2a), (2b), (2c), and (2d) allow the effect of extreme heat, which had the largest effect on year-to-year yield variability to vary over the growing season. The columns use $k = 4, 5, 6,$ or 7 spline knots, respectively. The coefficient estimates on the $k - 1$ spline polynomials are difficult to interpret, and hence we plot them over the growing season in figure 2.8. There is considerable heterogeneity over the growing season: the most damaging effects occur during phase 0.3 to 0.4 of the growing season irrespective of how many spline knots we use. The behavior at the boundaries (close to 0 and 1) should be interpreted with caution, as there is little mass at these endpoints as shown in figure 2.6.⁷

Panel B of table 2.3 tests whether the time-varying portion (not the constant effect of degree days above 29°C) are statistically significant, which is always the case. Predicted damages of the 2012 heat wave increase to 21 percent in panel C, which is not surprising as most of the excessive heat happened in July, which is in the 0.3 to 0.4 window when extreme heat is most damaging. The spatial distribution of the predicted impacts for the specification in column (2b) is given in the top graph of figure 2.9.⁸ Note that

7. Recall that the largest exposure to degree days above 29°C happens around 0.4 to 0.5 of the growing season; that is, the effect is not simply largest when exposure is highest.

8. Since the state-specific planting dates are only available for some years starting in 1979, weather anomalies are calculated as the difference to the observed weather average in our estimation sample.

Table 2.3 The effect of weather on maize yields using time-varying growing seasons

	(1)	(2a)	(2b)	(2c)	(2d)
<i>A. Time invariant variables</i>					
Thousand degree days 10–29°C	0.333*** (0.091)	0.322*** (0.087)	0.320*** (0.085)	0.317*** (0.086)	0.322*** (0.083)
Hundred degree days above 29°C	-0.591*** (0.086)				
Precipitation (m)	0.649*** (0.211)	0.622** (0.217)	0.608** (0.230)	0.589** (0.216)	0.648*** (0.222)
Precipitation (m) squared	-0.439** (0.166)	-0.392** (0.166)	-0.384** (0.173)	-0.373** (0.166)	-0.409** (0.170)
<i>B. Joint sig. of time-varying variable</i>					
$F_{\text{Degree Days Above } 29^{\circ}\text{C}}$		80.49	70.59	64.54	51.75
$p_{\text{Degree Days Above } 29^{\circ}\text{C}}$		1.27e-10	9.43e-11	7.36e-11	2.21e-10
<i>C. Impact of 2012 weather outcome</i>					
Total production impact (%)	-18.54	-20.89	-20.68	-20.80	-21.68
<i>D. Prediction error for 2012</i>					
RMSE–2012 county prediction	0.3688	0.3320	0.3333	0.3332	0.3321
Pred. error total prod. 2012 (%)	8.00	4.20	4.48	4.61	3.55
R^2	0.5151	0.5366	0.5364	0.5369	0.5422
Observations	43,249	43,249	43,249	43,249	43,249
Counties	1,659	1,659	1,659	1,659	1,659
Spline knots (time-varying var.)		4	5	6	7

Notes: Table regresses log maize yields for counties east of the 100-degree meridian where state-level planting dates are available in 1979 to 2011. Counties are shown in figure 2.5. Column (1) uses the same specification as column (1a) in table 2.2 except that it only uses counties and years for which planting dates are available and averages the weather variables over the actual growing season (instead of March through August). The remaining columns (2a) to (2d) allow the effect of degree days above 29°C to vary over the growing season. Columns differ by the number of spline knots used in the estimation of the seasonality, varying from 4 to 7 knots. The spline polynomials are shown in figure 2.8. Panel B of the table gives the F -statistics as well as the p -value for the joint significance of the *time-varying* components (not including the constant marginal effect). Panel C gives the predicted production shortfall below trend from the 2012 weather outcomes in percentage points. Panel D compares prediction for 2012 to actual observed yields. The first row shows the root mean squared prediction error of all county-level log yields, while the second row gives the prediction error of total production for all counties combined. Errors are clustered at the state level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

broadly comparable spatial pattern to the results we got when we fixed the growing season to March through August and forced the weather variables to have the same impact for all days of the growing season in figure 2.7, but the magnitude of the impacts is larger.

Panel D again compares predicted log yields and total production to the observed outcomes in 2012. While the prediction error decreases from col-

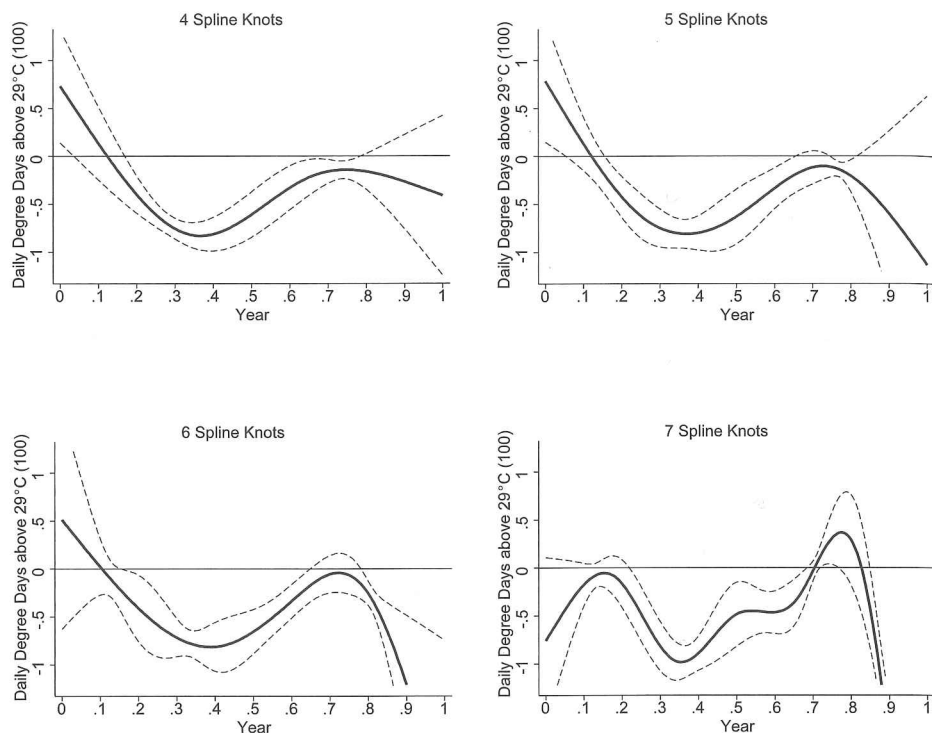


Fig. 2.8 Effect of degree days above 29°C as it varies over the growing season
Notes: Panels show the marginal effect of 100 degree days above 29°C. A reference model that fixes the effect to be the same across the growing season gave an estimate of -0.59 in column (1) of table 2.3.

umns (1) to (2d) as the models become more flexible, it is comparable in column (2d) of table 2.3 to column (1a) of table 2.2. The longer time series of the more simplistic model in table 2.2 gives a better prediction of the trend, which is counterbalanced by more accurately predicted production shortfall in table 2.3.

A lot of media coverage focused on the concurrence of extremely hot temperatures and drought conditions. Table 2.4, therefore, also includes an interaction term between daily degree days above 29°C and daily precipitation levels. The precipitation variables are different from the measures we used until now: we previously measured *growing season total* precipitation and its square. Since we are now interested how the effect varies over the growing season, we use *daily* precipitation and daily precipitation squared, which are then aggregated over the season.

The interaction is not significant in column (2), and the inclusion has almost no effect on the predicted impact of the 2012 weather outcomes in panel C. Columns (3) through (6) consecutively relax the assumption that

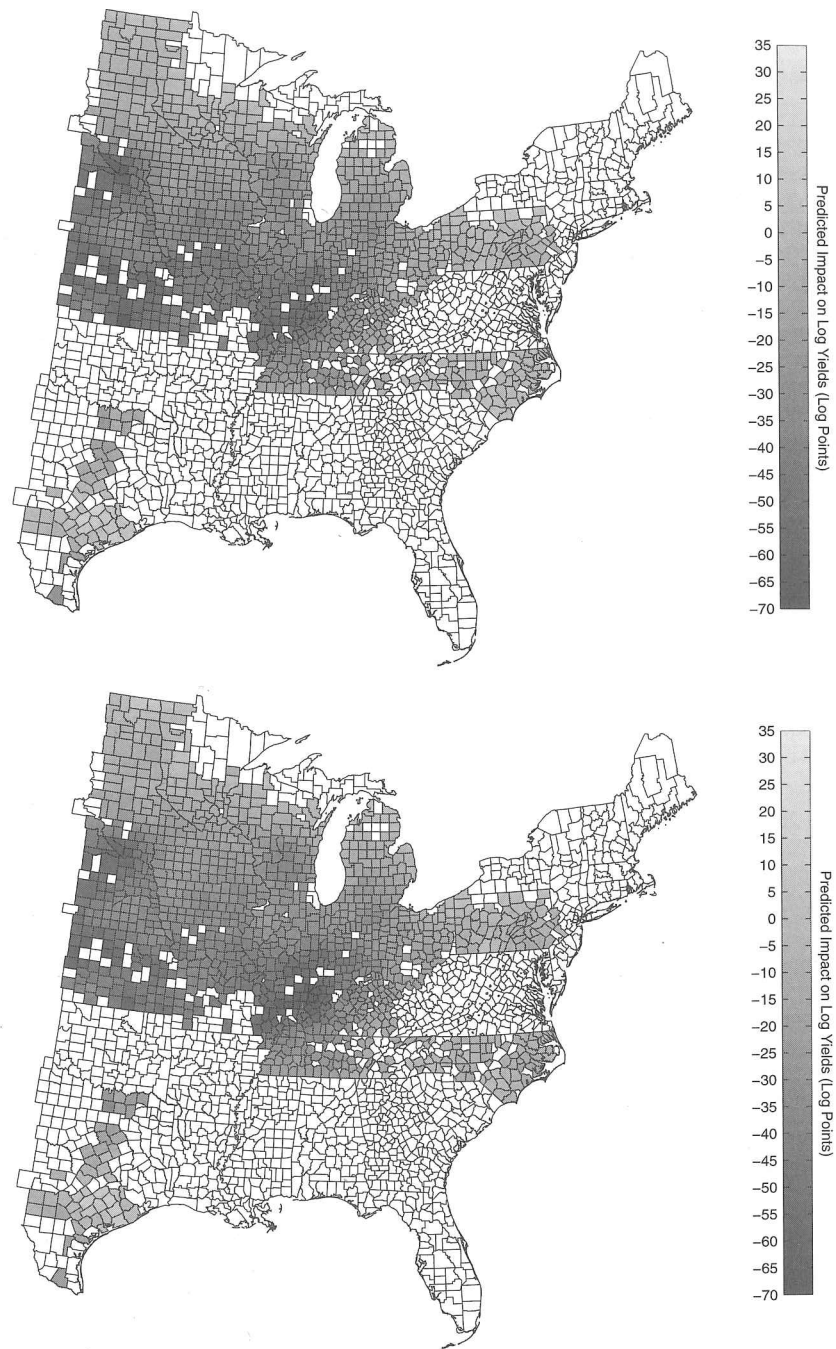


Fig. 2.9 Predicted yields in 2012 using time-varying coefficients
Notes: Both panels show changes in predicted yields in log points. The top panel uses the regression specification in column (2b) of table 2.3, while the bottom panel uses the specification in column (6) of table 2.4.

Table 2.4 The effect of weather on maize yields using time-varying growing seasons and precipitation interactions

	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Time invariant variables</i>						
Thousand degree days 10–29°C	0.333*** (0.091)	0.354*** (0.075)	0.334*** (0.074)	0.336*** (0.072)	0.313*** (0.074)	
Hundred degree days above 29°C	–0.591*** (0.086)	–0.562*** (0.107)				
Days above 29°C × precipitation		–32.435 (31.586)	–19.560 (25.565)			
Precipitation (m)	0.649*** (0.211)	0.708*** (0.207)	0.650** (0.231)	0.654** (0.237)		
Precipitation (m) squared	–0.439** (0.166)	–0.473*** (0.160)	–0.409** (0.170)	–0.415** (0.173)		
<i>B. Joint significance of time-varying variable</i>						
$P_{\text{Degree Days Above } 29^{\circ}\text{C}}$			7.88e–10	4.88e–09	2.22e–07	4.00e–09
$P_{\text{Degree Days Above } 29^{\circ}\text{C} \times \text{Precipitation}}$				0.0000619	0.00213	0.0157
$P_{\text{Precipitation}}$					0.00453	0.00426
$P_{\text{Precipitation Squared}}$					0.000857	0.00186
$P_{\text{Degree Days } 10\text{--}29^{\circ}\text{C}}$						0.0352
<i>C. Impact of 2012 weather outcome</i>						
Total production impact (%)	–18.54	–18.78	–20.79	–20.73	–22.19	–22.80
<i>D. Prediction error for 2012</i>						
RMSE—2012 county prediction	0.3688	0.3672	0.3329	0.3285	0.3328	0.3271
Pred. error total prod. 2012 (%)	8.00	8.09	4.55	4.67	2.96	1.69
R^2	0.5151	0.5167	0.5370	0.5407	0.5524	0.5540
Observations	43,249	43,249	43,249	43,249	43,249	43,249
Counties	1,659	1,659	1,659	1,659	1,659	1,659
Spline knots (time- varying var.)			5	5	5	5

Notes: Table regresses log maize yields for counties east of the 100-degree meridian where state-level plating dates are available in 1979 to 2011. Counties are shown in figure 2.5. Column (1) is the same as column (1) in table 2.3. Column (2) adds an interaction term between daily extreme heat and precipitation. Column (3) allows the effect of extreme heat to vary over the growing season (similar to column [2b] in table 2.3). Columns (4) to (6) allow the effect of other variables to vary over the season: respectively, the effect of the interaction between extreme heat and precipitation, the effect of precipitation and precipitation squared, and the effect of moderate degree days 10°C–29°C. Panel B of the table gives the p -values for the joint significance of the *time-varying* components (not including the constant marginal effect). Panel C gives the predicted production shortfall below trend from the 2012 weather outcomes in percentage points. Panel D compares the prediction for 2012 to actual observed yields. The first row shows the root mean squared prediction error of all county-level log yields, while the second row gives the prediction error of total production for all counties combined. Errors are clustered at the state level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

various weather variables are constant over the growing season (we use state-specific restricted cubic splines with 3 spline knots to capture time trends and restricted cubic splines with 5 knots to capture seasonality components of the effects of weather variables for all specifications). While the time-variant portions of all weather variables are significant as shown in panel B (p -values are generally less than 0.05), the predicted weather impacts for 2012 in panel C are comparable to a model where we only allow the effect of degree days above 29°C to vary over the growing season. The spatial distribution of impacts under the most flexible model (column [6]) is shown in the bottom graph of figure 2.9. The pattern is remarkably similar to the top graph that only allows the effect of degree days above 29°C to vary over the growing season.

The model that is most flexible in all weather variables (column [6]) has the lowest prediction error for 2012 as shown in panel D of table 2.4, suggesting that flexibility in the seasonal effects of the weather variables improves the prediction. In summary, switching from a fixed growing season (March through August) to a time-varying growing season gave larger prediction errors, but allowing the effect of the weather variables to vary over the growing season reduced it again. Both the time invariant baseline model as well as the model using time-varying parameters predicted the effects of 2012 fairly accurately.

2.4 Discussion

The 2012 heat wave resulted in significant production shortfalls. A baseline model that holds the growing season as well as the effect of the weather variables over the growing season constant gives predicted declines of 14.4 percent. If we instead average the weather measures over the actual growing season, the impacts increase to 19 percent, and if we allow the effect of extreme heat to vary over the growing season, the predicted damages increase further in magnitude up to 23 percent as the heat wave hit when it is most damaging.

For comparison, a comparable model to our baseline model in Schlenker and Roberts (2009) predicted decreases of slightly more than 20 percent under the Hadley III climate change model by midcentury (2020 to 2049). The predicted impacts from 2012 are hence predicted to become more frequent pretty soon if the climate forecasts turn out to be accurate.

Hansen, Sato, and Ruedy (2012) look at the frequency of extreme temperatures around the world and argue that it is predicted to increase significantly with climate change. The chapter finds that the United States is one of the few areas that has been “lucky” so far, in the sense that it has not seen a significant increase in observed extremes. The year of 2012 might soon be the new normal.

2.5 Conclusion

We model the impact of the 2012 heat wave/drought with two models. A baseline model keeps the growing season as well the effect of various weather measures over the growing season constant. In a new extension, we then obtain the actual growing season on a state level and allow the effect of weather to vary over the growing season. We find that the time-varying components are highly statistically significant.

The baseline model predicts overall production declines in our sample of 14.4 percent. While some areas are severely hit, others actually have above-normal yields. Once we use the actual growing season (instead of the artificially fixed one), the production decline goes up in magnitude to 19 percent. If the effect of extreme heat is allowed to vary over the growing season, the predicted damage increases further to 23 percent as the heat wave hit during a time when it is most damaging. Production shortfalls of around 20 percent in an area that accounts for 40 percent of global production will have strong effects on prices. Recall that historic global corn production shocks (deviations from a trend) ranged from -13 percent to +7 percent in 1961 to 2010.

If climate forecasts turn out to accurate, we will experience increased variability in degree days above 29°C even if the variance of temperatures remains constant. The reason behind this behavior is that degree days above 29°C are a truncated temperature variable. An upward shift in the mean of the variable that leaves the variance constant will increase year-to-year variability of degree days above 29°C as the bound of 29°C binds less frequently. Temperature fluctuations below 29°C have no effect on degree days above 29°C, while temperature fluctuations above the threshold do. An upward shift in temperatures hence shifts more mass of the probability distribution to a region where it translates into fluctuations of damaging degree days. Climate change has the potential to not only decrease average production, but also to make it more volatile. As a response, food price volatility will likely increase, even though some of the increased volatility will be buffered through higher storage levels.

References

- Boone, Christopher, Wolfram Schlenker, and Juha Siikamäki. 2013. "Ground-Level Ozone Pollution and Corn Yields in the United States." Working Paper.
- Hansen, James, Makiko Sato, and Reto Ruedy. 2012. "Perception of Climate Change." *Proceedings of the National Academy of Sciences* 109 (37): E2415–E2423.
- Ortiz-Bobea, Ariel, and Richard E. Just. 2013. "Modeling the Structure of Adaptation in Climate Change Impact Assessment." *American Journal of Agricultural Economics* 95 (2): 244–51.
- Roberts, Michael J., and Wolfram Schlenker. 2013. "Identifying Supply and Demand

Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate." *American Economic Review* 103 (6): 2265–95.

- Roberts, Michael J., Wolfram Schlenker, and Jonathan Eyer. 2013. "Agronomic Weather Measures in Econometric Models of Crop Yield with Implications for Climate Change." *American Journal of Agricultural Economics* 95 (2): 236–43.
- Schlenker, Wolfram, and Michael J. Roberts. 2009. "Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change." *Proceedings of the National Academy of Sciences of the United States* 106 (37): 15594–8.
- Snyder, R. L. 1985. "Hand Calculating Degree Days." *Agricultural and Forest Meteorology* 35 (1–4): 353–8.

Comment Derek Headey

Overview of the Chapter

In this chapter Berry, Roberts, and Schlenker extend some of their earlier work on the effects of weather shocks on US maize production. A key motivation for their chapter—and the link to the broader theme of this book—is that the United States is a major producer and exporter of maize, such that production shocks in the United States are a potential driver of maize price volatility, which may have important ramifications for the world's poor.¹ The main technical innovations of this chapter are that they now allow the effect of various weather measures to evolve over the growing season, and that the growing season is made more location specific. This new and improved model is then applied to the 2012 growing season, when large parts of the US maize belt experienced a severe heat wave and drought. Strikingly, their improved model predicts yield declines of up to 24 percent. In their concluding remarks they note that some climate change models predict that these kinds of heat spells/droughts may well be the new normal in the US maize belt.

My comments will be confined to four areas: a few technical issues, a quick look at whether their predictions came true, some discussion and exploratory analysis of the impact of US maize production on international prices, and some policy and programmatic implications of their model and results.

Some Technical Issues

Technically, the chapter is strong. The authors build on much simpler attempts to model weather with production outcomes, with a particular

Derek Headey is a research fellow of the International Food Policy Research Institute (IFPRI).

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see <http://www.nber.org/chapters/c12807>.ack.

1. Maize is the most important staple food in Africa, and a major crop in Latin America.