

CORN PRODUCTION SHOCKS IN 2012 AND BEYOND: IMPLICATIONS FOR FOOD PRICE VOLATILITY

Steven T. Berry♣, Michael J. Roberts♦ and Wolfram Schlenker♡

September 2012

Abstract

Corn prices increased sharply in the summer of 2012 due to expected production shortfalls in the United States, which produces roughly 40% of the world's corn. A heat wave in July adversely affected corn production. We extend earlier statistical models of county-level corn yields in the Eastern United States by allowing the effect of various weather measures to vary in flexible manner over the growing season: Extreme heat is especially harmful around a third of the way through the growing season. This is the time when the heat wave hit the corn belt region of United States in 2012. As a result, the model predicts corn yields will be 23% below trend. While extreme heat was significantly above normal, climate change scenarios suggest that the 2012 outcomes will soon be the new normal.

♣ Yale University and NBER. Email: steven.berry@yale.edu

♦ University of Hawaii. Email: mjrobert@hawaii.edu

♡ University of California at Berkeley and NBER. Email: schlenker@berkeley.edu

Historically, 25% of an average year's global corn production is held in inventories to buffer weather shocks and smooth them 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% to +7% in 1961-2010, with a standard deviation of 4% (Roberts & Schlenker 2012). International trade smoothens 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% 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 driver of production shocks are weather fluctuations. An accurate model that translates weather fluctuations into production fluctuations is a crucial first step in examining food price volatility.

In this paper we extend earlier work on the effects of weather on corn production (Schlenker & Roberts 2009). We previously allowed for a highly non-linear 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-August. The main innovations of this paper 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-August, but rather use annual state-level data on plating 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 stronger production shortfalls. Fourth, we contrast 2012 to what is expected under climate change.

1 Model

We start by estimating a baseline model of yields that assumes a fixed growing season (March-August) and a constant effect of weather variables over the growing season. This

baseline replicates a the baseline model from earlier research (Schlenker & Roberts 2009). Then, following Ortiz-Bobea & Just (Forthcoming), we consider models that account for planting date and temperature effects that vary over the growing season.

1.1 Baseline Model 1

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

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

where m_{it} is growing degree days between 10 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 and p^2 are total precipitation and total precipitation squared.¹ County fixed effects c_i account for baseline differences between counties and state-specific time trends f_s account for 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 regression is constructed using daily fine-scaled weather measures on a 2.5x2.5 mile grid for the contiguous United States. We follow the same algorithm of Schlenker & Roberts (2009), but update the data until August 31, 2012. We only use counties east of the 100 degree meridian (except Florida) in the regression, because the response function might be different for highly irrigated area. The data set spans the years 1950-2012.

1.2 Model 2: Time-Varying Parameters

Model 2 allows the effect of weather variables to vary over the growing season. Ortiz-Bobea & Just (Forthcoming) extend our earlier work by separating the growing season into three subintervals, and then estimate separate (constant) coefficient for each of the subintervals. This paper 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

¹*Growing degree days* measures cumulative heat exposure above a threshold temperature and/or below an upper bound. So, degree days 10-29°C counts all temperatures below 10°C as zero, temperatures between 10° and 29° as the cumulative difference between temperature and 10 ° and temperatures at or above 29°C as 19 per unit time. 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 2 degree days, and so on. Degree days above 29°C puts the lower bound at 29°C and has no upper bound.

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 measure 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 this weather measures should hence evolve over time. Model 2 is

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

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-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 plating (time 0) to harvest (time 1).

Specifically, we construct a restricted cubic spline polynomial in time over the growing season with k knots, 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.² Because we normalize the growing season to length one, the “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 becomes the phase of the growing season $x_{dit} = \frac{d}{D_{it}}$

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

The second equality simply exchanges the order of summation. We are ultimately left with $k - 1$ 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 2 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

²Spline knots locations are as follows: $k = 3$ indicates three 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 0.2, 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 0.02; $k = 7$ knots are set at 0.02, 0.18, 0.34, 0.5, 0.66, 0.82, and 0.98.

degree days above 29°C and daily precipitation

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

Note that the effect of precipitation is modeled differently. Originally we modeled the effect of season-total precipitation and season-total precipitation squared. Here we model the effect of *daily* precipitation as well as *daily* precipitation squared, and allows these effects to vary over the growing season.

2 Data

We pair data on annual county-level corn yields with fine scaled-weather measures that were constructed on a 2.5x2.5 mile grid for the entire United States. We follow the same algorithm of Schlenker & Roberts (2009), but update the data until August 31, 2012. These data give daily minimum and maximum temperature, dew point, 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-August 31st for all counties and years. These variables were calculated using all counties east of the 100 degree meridian (except Florida). The second row of Table 1 displays the fraction of the US growing area and production that falls in these countries for the three most recent years 2009-2011. With 90% 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 eastern states closely match US overall yields as shown in the last two columns of the table.

2.1 Extreme Heat in Key Growing Regions (Model 1)

All counties east of the 100 degree meridian (except Florida) that grow corn and for which we have weather variables are displayed on the maps in Figure 1. The top graph shows 2012

anomalies of season-total degree days above 29°C, i.e., the difference between 2012 and the average from 1950-2011. The bottom graph shows 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 : the production-weighted average exposure to degree days above 29°C is 33 among all eastern counties in 1950-2011. Some highly productive counties in the Corn Belt experienced up to 110 extra degree days above 29°C. As we show below, each degree days above 29°C reduces log yields by 0.006 log points, so the effect of an extra 100 degree days above 29°C is a decrease of 60 log points.

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.³ Figure 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 1950-2011 are shown as grey dashed lines, while the outcome for 2012 is shown as a thick solid red line.

The top panel of Figure 2 shows degree days 10-29°C. Since the US had a warm spring, degree days above 10-29°C start to increase earlier than usual. The beneficial side effect is that a warmer spring allows for earlier planting. The bottom graph of Figure 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 been exceptionally hot by historic standards. At the beginning of July, the measures was slightly above normal, but by the end of July, it had superseded the hottest year among the 1950-2011 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 3 displays the cumulative season total for precipitation. It was below normal in 2012, and the only year with drier condition in the 1950-2011 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 3

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

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 models (Roberts, Schlenker & Eyer Forthcoming). 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 Planting Dates and Growing Seasons (Model 2)

The second model relaxes two assumptions: first, we no longer fix the growing season to March-August, but instead used data from the National Agricultural Statistics Service (NASS) on planting and harvesting dates. NASS reports on a weekly level what fraction of the corn area in major corn-producing states was harvested and planted. We define the beginning of the growing season as the Monday of the week by the end of which at least 50% 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% 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 4. Southern places tend to plant earlier, as they are not limited by the probability of late freezes. On top of that, Northern places have a larger intra-year cycle in solar radiation, which is an important component of crop growth that limits farmers to shift the planting date too much 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, but not the harvest date, is available for a state, we approximate the harvest date by adding the average growing season length to the reported planting date. By the 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 also have a longer growing season as shown in the bottom graph of Figure 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. The number of yearly observations we obtain for each county is shown in Figure 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 (except Florida) that report planting and/or harvest dates are summarized in the third row of Table 1. States that report planting dates account for 85% of the corn growing area and 87% of the US corn production in the

most recent three years (2009-2011).

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 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 6 displays the average daily exposure over the growing season for four weather variables: degree days 10-29°C, degree days above 29°C, precipitation, and vapor pressure deficit. We use either restricted cubic splines with 4, 5, 6, or 7 knots. The results seem fairly stable as long as we include at least 5 knots. The harvest date in 2012 are approximated by adding the average growing season length to the observed planting dates. Our weather variables are only observed until August 31, 2012 at the time of writing. In case the harvest date occurs after August 31st, we approximate the rest of the growing season by assigning each day after August 31st the average exposure in the given county.

3 Empirical Results

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

3.1 Baseline Model 1

Results for a panel analysis for eastern counties (except Florida) for the years 1950-2011 is given in Table 2. All columns use the same set of observations, but vary the time controls to capture overall trends in yields. Columns (a), (b), and (c) use state-specific restricted cubic splines with 3, 4, and 5 knots, respectively. On top of that, columns (21)-(2c) also include year fixed effects to model technological breakthroughs. The results are very stable across specifications. Moderate heat (degree days 10-29°C) is beneficial, while extreme heat (degree days above 29°C) is highly damaging. Note that 2000 degree days 10-29°C increase expected yields by as much as 100 degree days above 29°C decreases them. Moreover, a coefficient of -0.6 on degree days above 29°C implies that the 100 additional degree days that some counties experienced in 2012 are lowering expected yields by 60 log points. 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 25inches, which matches closely the estimate of optimal rainfall from agronomic studies.

The effect of the 2012 weather outcomes on expected yields are shown in Figure 7. The top graph depicts predicted deviations from the time trend in 100 log points (using the specification from column (1a) in Table 2). There is significant heterogeneity, some counties are predicted to be as much as 62% below normal, while others experience up to 28% above normal yields. Unfortunately, the more productive areas experience the yield declines. The bottom graph does not show relative impacts, but predicted *total impacts*. We set the harvest area to the average of the last three years (2009-2011) and multiply it with the predicted production shortfall per area. 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 15% production shortfall.

3.2 Model 2: Time-Varying 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 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-August). The coefficient on the two degree days variables remain rather unchanged. The precipitation variables are different: 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. Panel C summarizes the predicted decrease in total production, which is 20%.

Column (2a)-(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 we hence plot them over the growing season in Figure 8. There is considerable heterogeneity over the growing season: the most damaging effects occur during phase 0.3-0.4 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 6. Panel B of Table 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 24-25% in Panel C, which is not surprising as most of the excessive heat happened in July,

which is in the 0.3-0.4 window when extreme heat is most damaging.

A lot of media coverage focused on the concurrence of extremely hot temperatures and drought conditions. Table 4 therefore also includes an interaction term between daily degree days above 29°C and daily precipitation levels. 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. Panels (3) through (6) consecutively relax the assumption that 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 impacts for 2012 are comparable to a model where we only allow the effect of degree days above 29°C to vary over the growing season.

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 15%. If we instead average the weather measures over the actual growing season, the impacts increase to 20%, and if we allow the effect of extreme heat to vary over the growing season to vary over the growing season, the effects further increase to 24% as the heat wave hit when it is most damaging.

For comparison, a comparable model to our baseline model in Schlenker & Roberts (2009) predicted decreases of slightly more than 20% under the Hadley III climate change model by mid-century (2020-2049). 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 & 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 paper 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 current year might hence be soon the new normal.

5 Conclusions

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 15%. While some areas are severely hit, other actually have above-normal yields. Once we use the actual growing season (instead of the fixed one), the production decline goes up to 20%. If the effect of weather variables is allowed to vary over the growing season, the predicted impact further increases to 23% as the heat wave hit during a time when it is most damaging. Production shortfalls of 20% in an area that accounts for 40% of global production will have strong effects on prices.

If the climate forecasts turn out to accurate, we will experience increased variability in degree days above 29°C. The reason is that it is a truncated temperature variable. An upward shift in the mean of the variable that leaves the variance constant will increase year-to-year variability 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 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

Hansen, James, Makiko Sato, and Reto Ruedy. 2012. “Perceptions of Climate Change: The New Climate Dice.” *WorkingPaper*.

Ortiz-Bobea, Ariel, and Richard E. Just. Forthcoming. “Modeing the Structure of Adaptation in Climate Change Impact Assessment.” *American Journal of Agricultural Economics*.

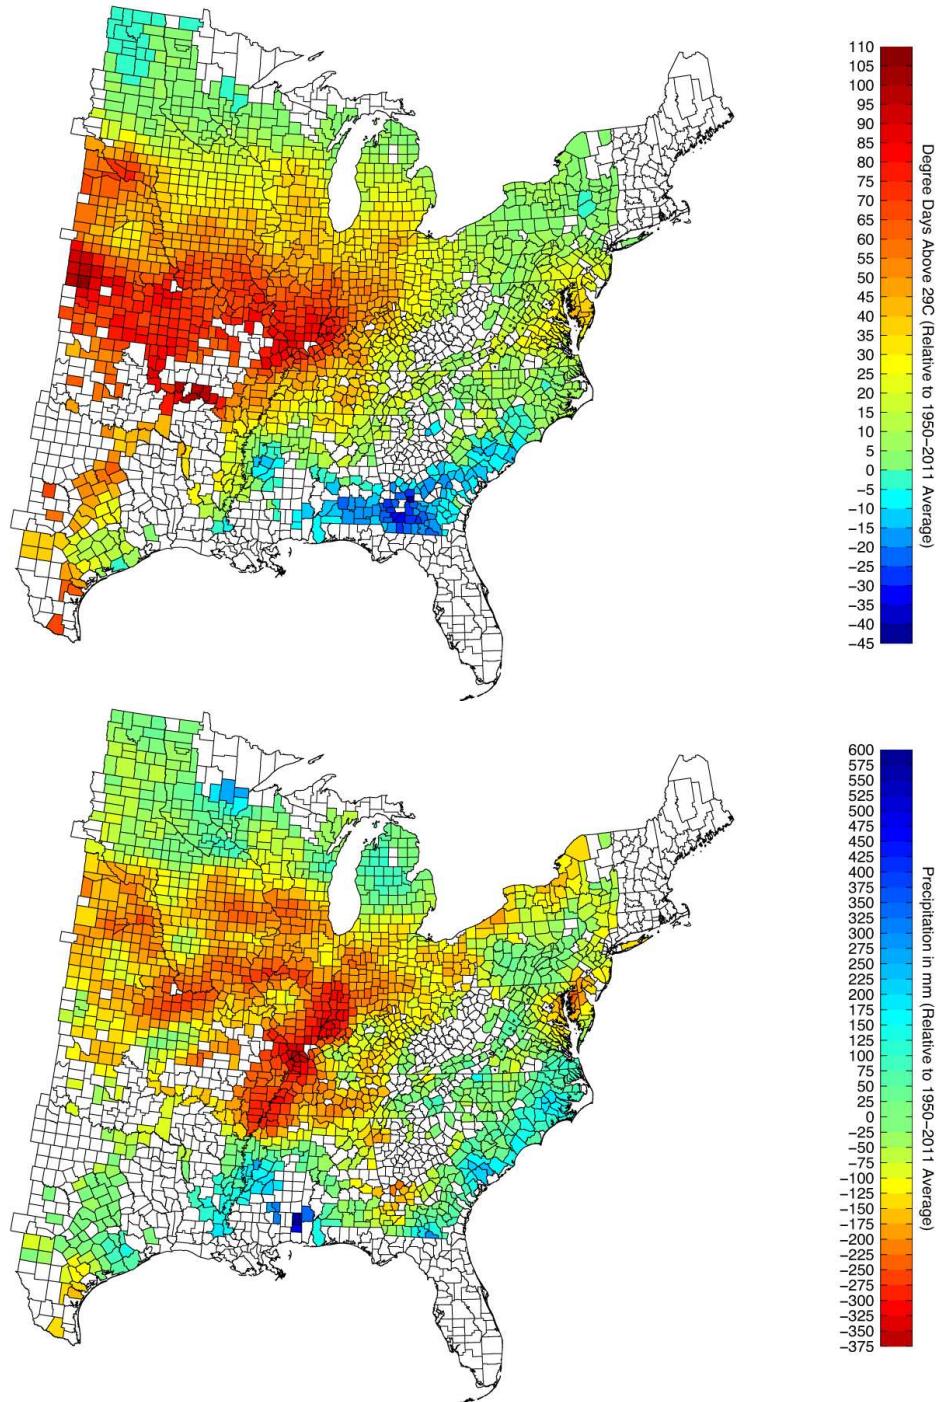
Roberts, Michael J., and Wolfram Schlenker. 2012. “Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate.” *Working Paper*.

Roberts, Michael J., Wolfram Schlenker, and Jonathan Eyer. Forthcoming. “Agronomic Weather Measures in Econometric Models of Crop Yield with Implications for Climate Change.” *American Journal of Agricultural Economics*.

Schlenker, Wolfram, and Michael J. Roberts. 2009. “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change.” *Proceedings of the National Academy of Sciences of the United States*, 106(37): 15594–15598.

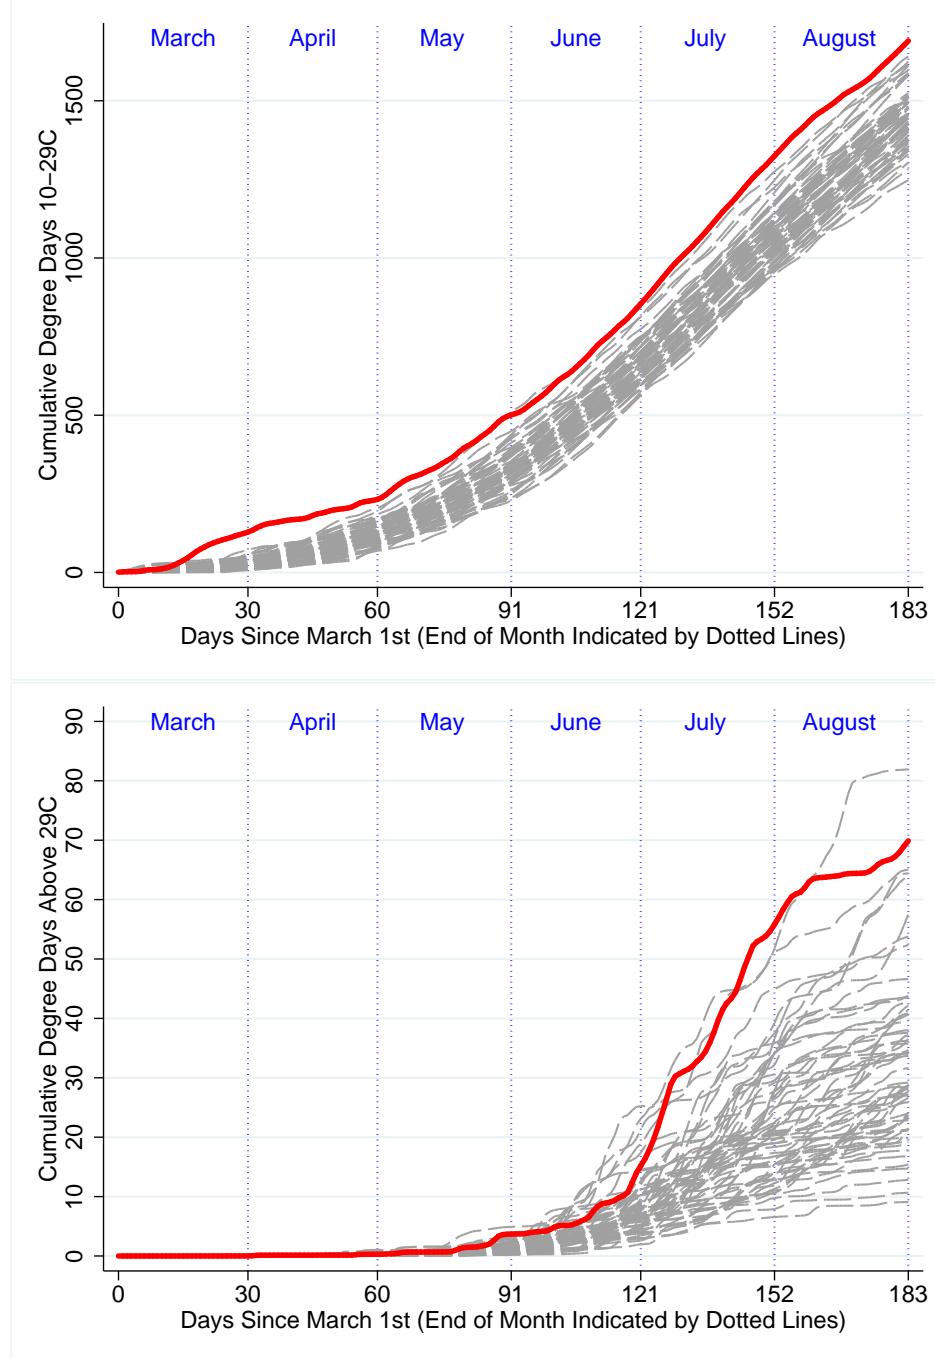
Snyder, R. L. 1985. “Hand Calculating Degree Days.” *Agricultural and Forest Meterology*, 35(1-4): 353–358.

Figure 1: Spatial Distribution of Degree Days Above 29C and Precipitation in 2012



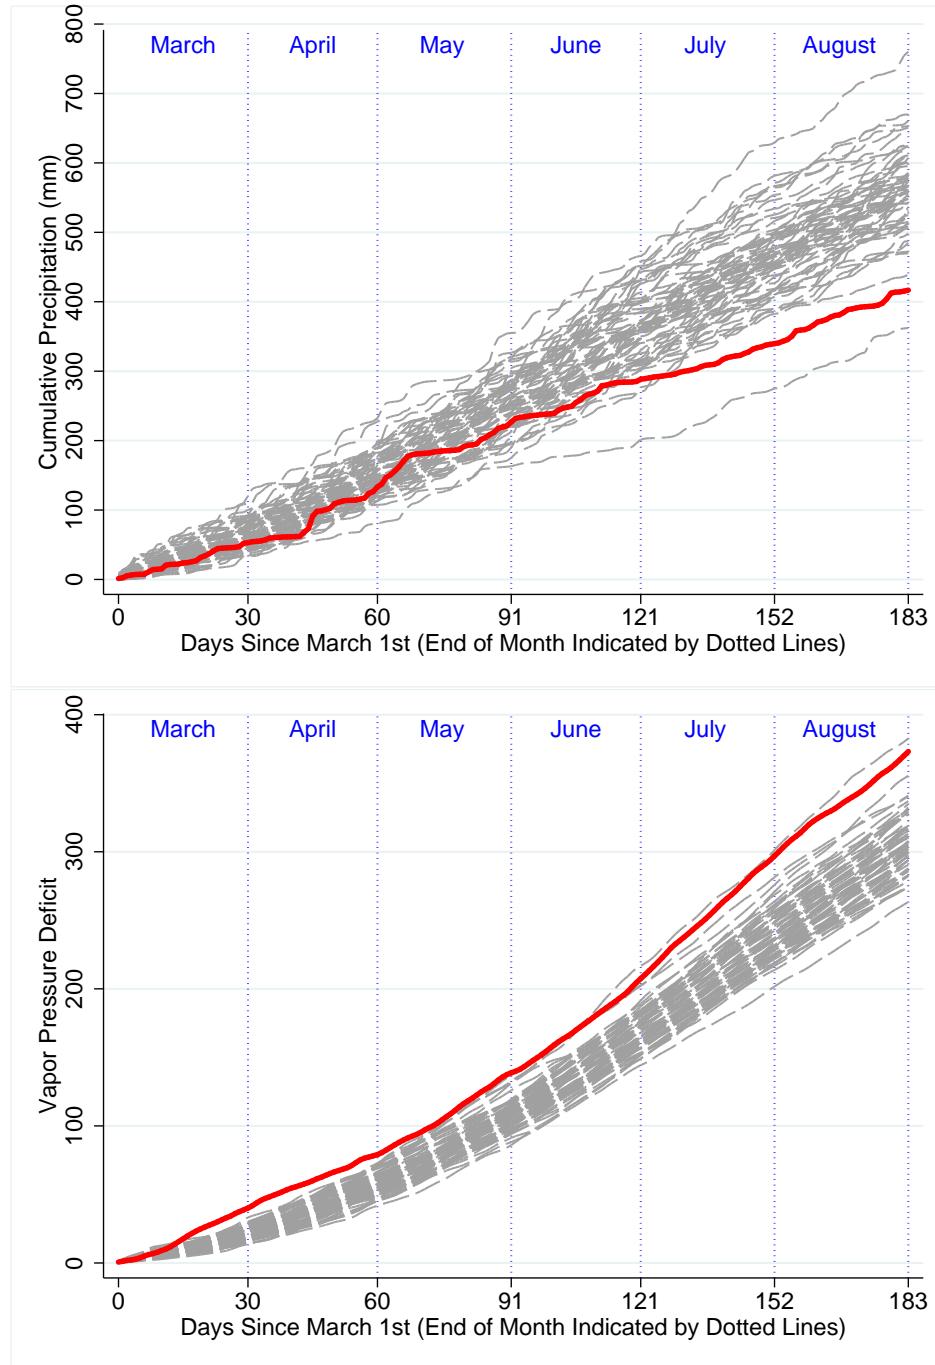
Notes: Spatial distribution of weather anomalies over the fixed 2012 growing season (March - August). Top panel shows degree days above 29°C, while the bottom panel shows precipitation totals.

Figure 2: Degree Days 10-29°C and Degree Days above 29°C in 2012 relative to 1950-2011.



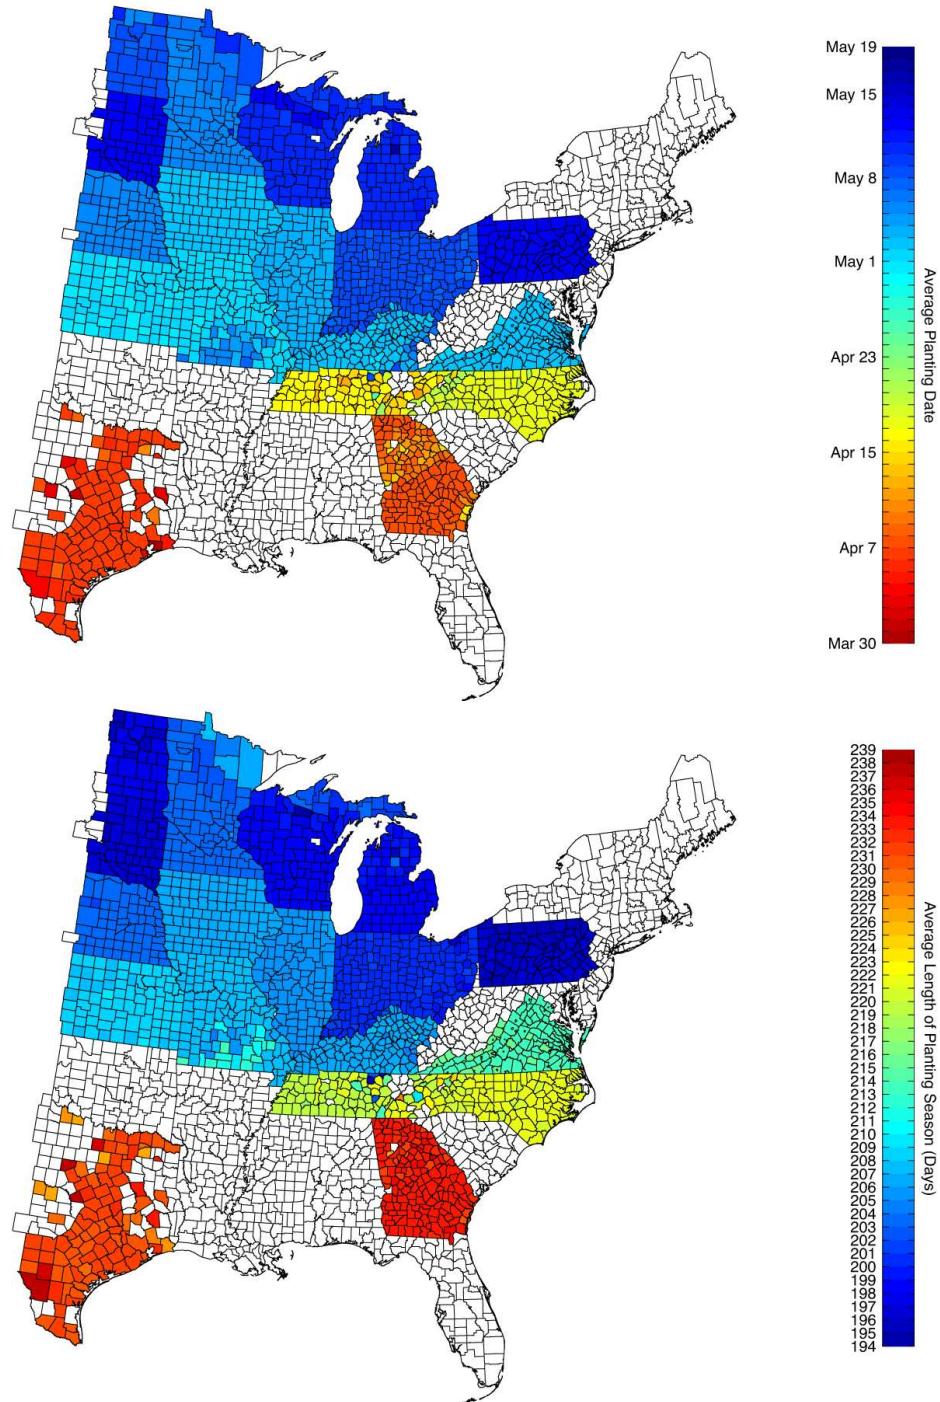
Notes: Panels show cumulative total of degree days 10-29°C and cumulative total of degree days above 29°C for Eastern United States. The weather measure is the weighted average of all counties east of the 100 degree meridian excluding Florida, where the weights are predicted yields along a trend line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

Figure 3: Precipitation and Vapor Pressure Deficit in 2012 relative to 1950-2011



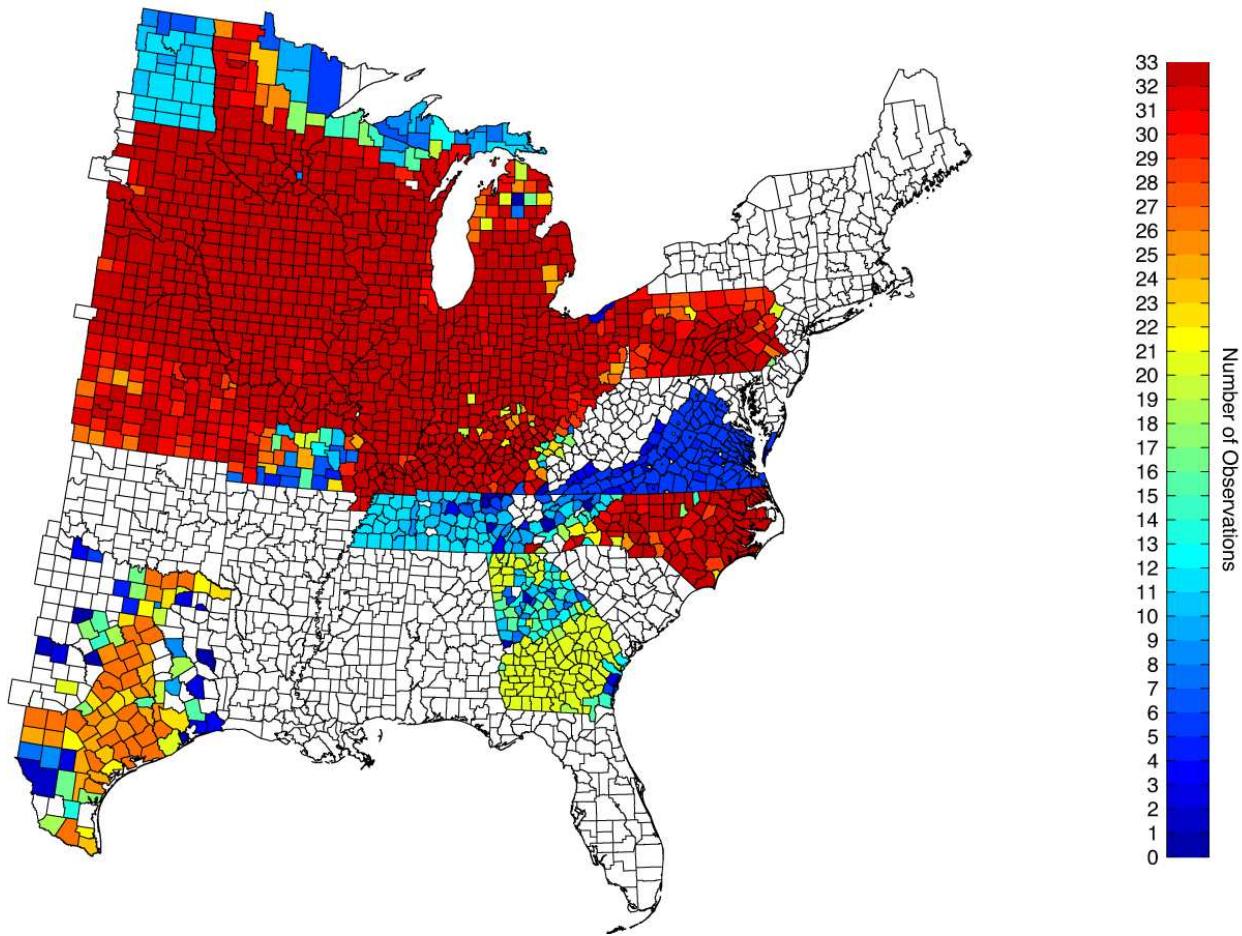
Notes: Panels show precipitation and vapor pressure deficit for Eastern United States. The weather measure is the weighted average of all counties east of the 100 degree meridian excluding Florida, where the weights are predicted yields along a trend line (restricted cubic spline with 3 knots) times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

Figure 4: Average Planting Date and Growing Season Length (1979-2011)



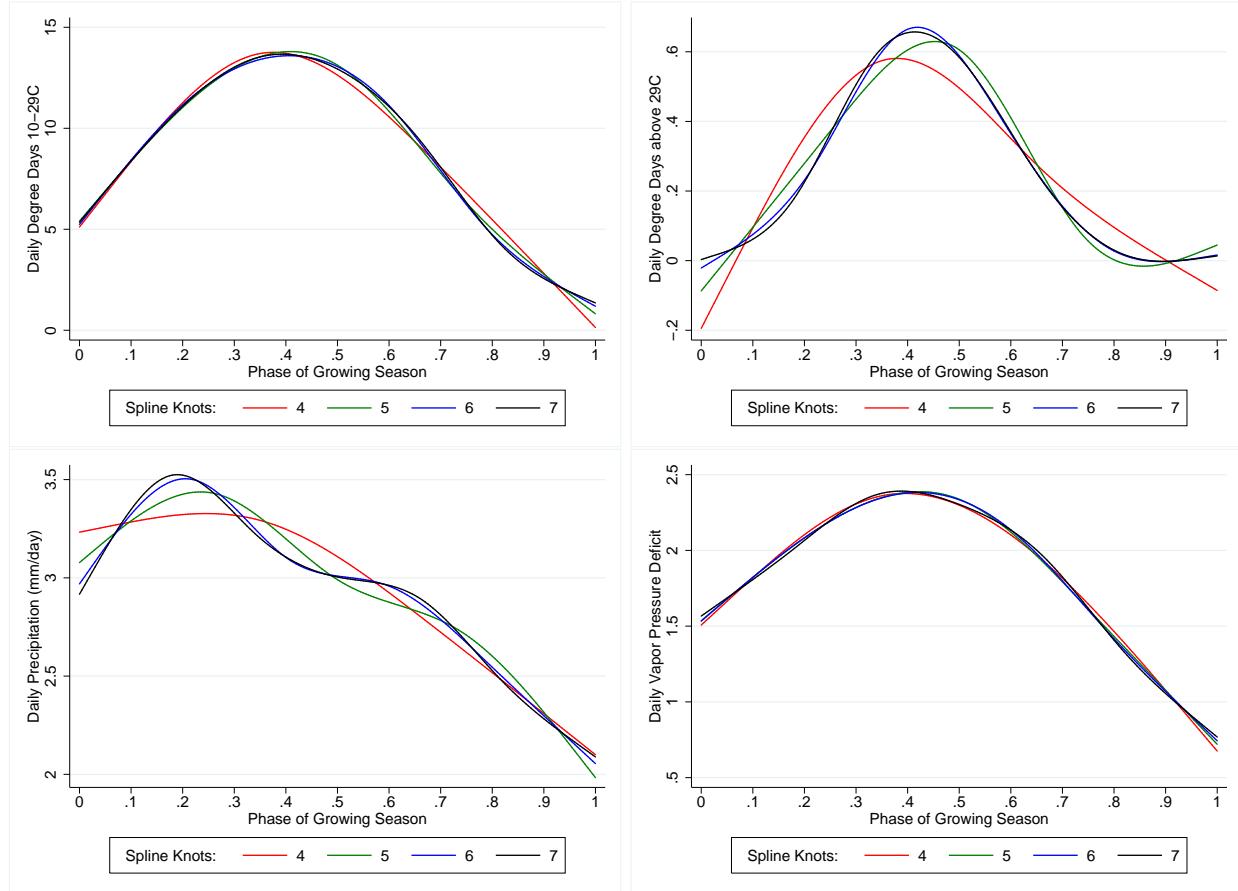
Notes: Top graph shows average planting date in 1979-2010, while bottom graph shows average growing season length. Both planting dates and growing season length are reported annually for each state. Counties within each state might have different values because they grew corn in different years.

Figure 5: Counties with Yearly State-level Data on Planting Dates (1979-2011)



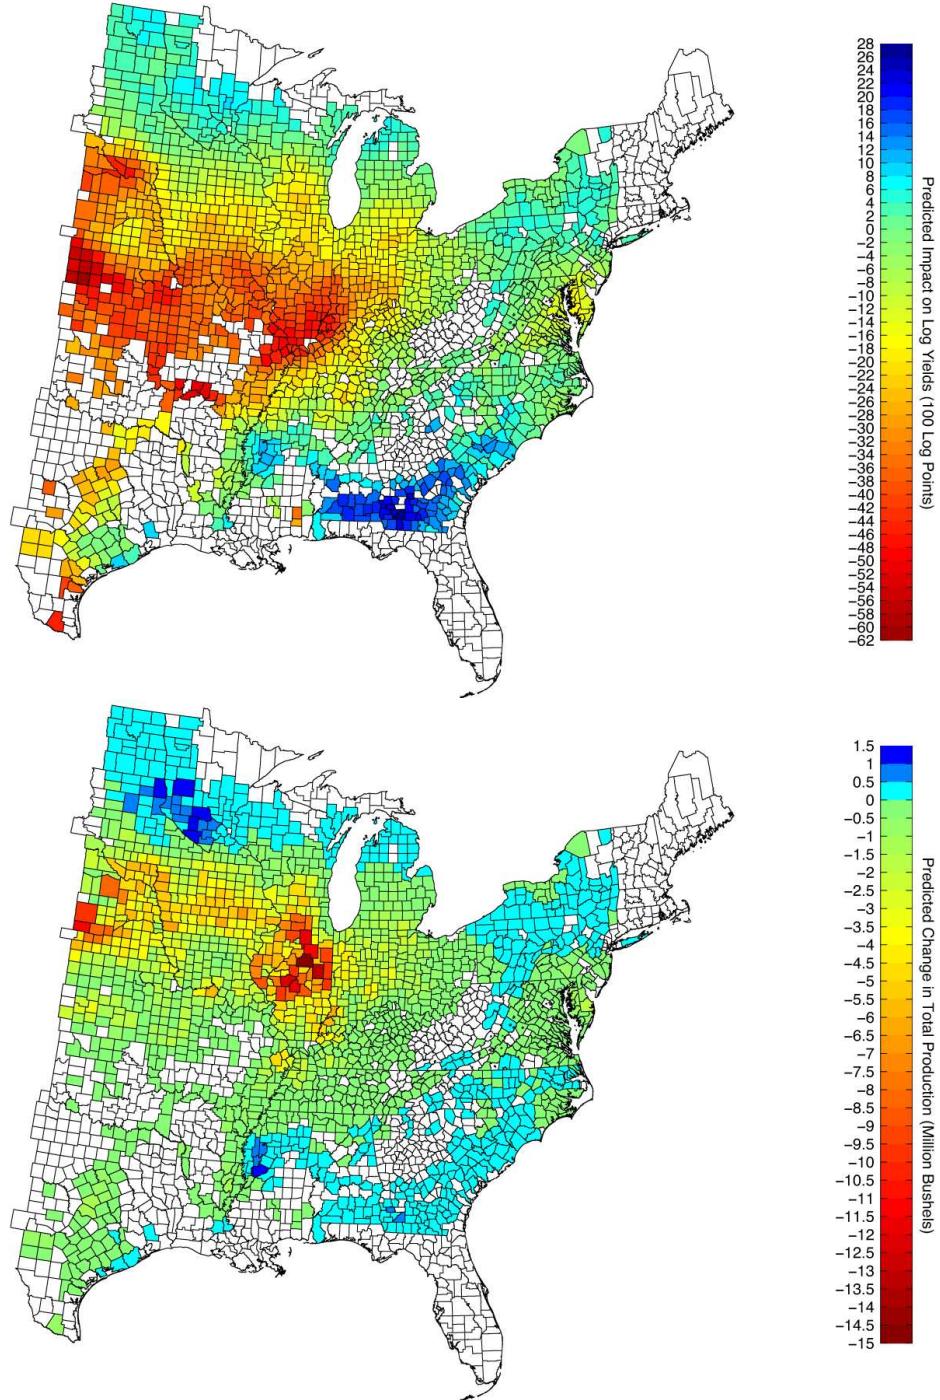
Notes: Figure displays counties for which annual planting and/or harvesting dates as well as yields were reported. Counties are color-coded by the number of yearly observations that are available for 1979-2011. Planting and harvest dates are reported on a state-level, while yields are reported for each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

Figure 6: Exposure to Various Weather Variables over the Growing Season



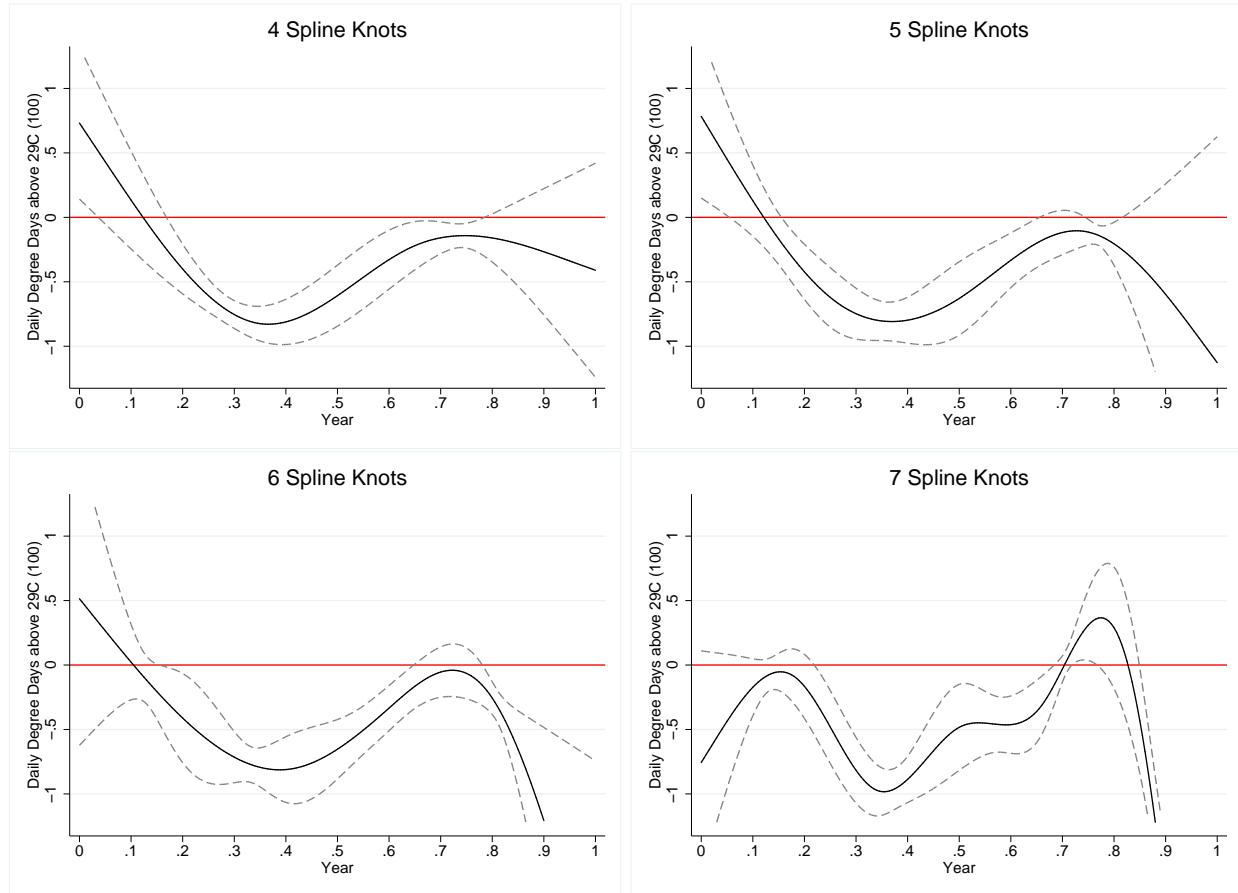
Notes: Panels show the average exposure to various weather measures over the growing season. We use year-and-state specific estimates of the National Agricultural Statistical Service (NASS) to define the growing season: The week in which the planted area exceeds 50% is the start (x-value of 0) and the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

Figure 7: Predicted Yields in 2012



Notes: Predicted yield and production impacts in 2012 by county using the regression specification in column (1a) of Table 2. The top panel shows changes in predicted yields in 100 log points, while the bottom shows predicted changes in total production (using the average area of 2009-2011 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%.

Figure 8: Effect of Degree Days above 29C as it Varies Over the Growing Season



Notes: Panels show the marginal effect of 100 degree days above 29C. 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 3.

Table 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 used in this study for the most recent three years 2009-2011. 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.

Table 2: The Effect of Weather on Maize Yields using a Fixed Growing Season March-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)
R ²	0.7734	0.7784	0.7810	0.7920	0.7955	0.7972
Observations	115205	115205	115205	115205	115205	115205
Counties	2276	2276	2276	2276	2276	2276
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 year 1950-2011 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. Errors are clustered at the state level. Stars indicate significance levels: *** : 1%; ** : 5%; * : 10%.

Table 3: The Effect of Weather on Maize Yields using Time-Varying Growing Seasons

	(1)	(2a)	(2b)	(2c)	(2d)
Panel A: Time Invariant Variables					
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)
Panel B: Joint Significance of Time Varying Variable					
$F_{\text{Degree Days Above } 29^\circ\text{C}}$		80.45	70.54	64.45	51.68
$p_{\text{Degree Days Above } 29^\circ\text{C}}$		1.28e-10	9.49e-11	7.44e-11	2.24e-10
Panel C: Impact of 2012 Weather Outcome					
Total Production Impact (%)	-20.30	-23.54	-23.86	-23.88	-25.24
R^2	0.5151	0.5366	0.5364	0.5369	0.5422
Observations	43249	43249	43249	43249	43249
Counties	1659	1659	1659	1659	1659
Spline Knots (Time Varying Var.)		4	5	6	7

Notes: Table regresses log maize yields for counties east of the 100 degree where state-level plating dates are available in 1979-2011. Counties are shown in Figure 5. Column (1a) uses the same specification as column (1a) in Table 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-August). The remaining columns (2a)-(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 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). Errors are clustered at the state level. Stars indicate significance levels: *** : 1%; ** : 5%; * : 10%.

Table 4: The Effect of Weather on Maize Yields using Time-Varying Growing Seasons and Precipitation Interactions

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Time Invariant Variables						
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)				
DDays Above 29°C X Precipitation		-32.434 (31.588)	-19.556 (25.569)			
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)		
Panel B: Joint Significance of Time Varying Variable						
$p_{\text{Degree Days Above 29°C}}$			7.88e-10	4.90e-09	2.20e-07	4.00e-09
$p_{\text{Degree Days Above 29°C X Precipitation}}$.0000632	.00216	.0158
$p_{\text{Precipitation}}$.00454	.00429
$p_{\text{Precipitation Squared}}$.000852	.00187
$p_{\text{Degree Days 10-29°C}}$.0341
Panel C: Impact of 2012 Weather Outcome						
Total Production Impact (%)	-20.30	-20.23	-23.67	-22.98	-24.24	-24.31
R ²	0.5151	0.5167	0.5370	0.5407	0.5524	0.5540
Observations	43249	43249	43249	43249	43249	43249
Counties	1659	1659	1659	1659	1659	1659
Spline Knots (Time Varying Var.)			5	5	5	5

Notes: Table regresses log maize yields for counties east of the 100 degree where state-level plating dates are available in 1979-2011. Counties are shown in Figure 5. Column (1) is the same as column (1) in Table 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 3). Column (4)-(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 10-29°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). Errors are clustered at the state level. Stars indicate significance levels: *** : 1%; ** : 5%; * : 10%.