

THE EVOLUTION OF HEAT TOLERANCE OF CORN: IMPLICATIONS FOR CLIMATE CHANGE

Michael J. Roberts♣ and Wolfram Schlenker♠

Abstract

We present new evidence on the relationship between weather and corn yields in Indiana between 1901 and 2005, extending earlier results for corn from 1950-2005. Indiana, a major corn-producing state has the best coverage of daily weather records for the early half of the 20th century. The effects of precipitation and extreme heat are shown to evolve over time as new seed varieties, supplemental irrigation systems and management practices are introduced. In particular, we find the detrimental effects of either too much or too little precipitation seems to have steadily diminished over time. In contrast, the evolution of tolerance to extreme heat is highly nonlinear, growing with the adoption of double-cross hybrid corn in the 1940's, peaking around 1960, and then declined sharply as single-cross hybrids came online. Corn in Indiana is most sensitive to extreme temperatures at the end of our sample. Since climate change models predict an increase in extreme temperatures, the big question is whether the next breeding cycles can increase both average yields and heat tolerance simultaneously as in the period 1940-1960, or whether an continued increase in average yields can only be achieved at the expense of more sensitivity to extreme heat as in the period from 1960 onwards. Finally, we discuss these impacts in relation to possible distortionary effects of current agricultural subsidies in the United States.

♣ Department of Agricultural and Resource Economics, North Carolina State University, Box 8109, Raleigh, NC. Email: michael_roberts@ncsu.edu.

♠ Department of Economics and School of International and Public Affairs, Columbia University, 420 West 118th Street, Room. 1308, MC 3323, New York, NY 10027. Email: wolfram.schlenker@columbia.edu.

1 Introduction

With evidence accumulating that greenhouse gas concentrations are warming the world's climate, there is growing interest in the potential impacts that may occur under different warming scenarios, and on how economies might adapt to changing climatic conditions. Agriculture is of particular interest due to the fact that climate is a direct natural input in the production process. Agriculture in developed nations, and particularly in the United States, has received considerable attention. This attention may derive from the fact that wealthier nations produce a disproportionate share of the world's agricultural commodities, at least partly due to their relatively more temperate climates. Accordingly, climate change impacts on agriculture in developed nations, and particularly the United States, the world's largest producer, have broad implications for food supply and prices worldwide.

In recent research, we conducted detailed statistical analyses of the relationship between weather and crop yields of corn, soybeans, and cotton in 1950-2005. These crops are among the four largest U.S. crops, all of which are important for world commodity prices (Schlenker and Roberts 2009). Corn and soybeans are two of the world's four key staple commodities that comprise about three quarters of calories produced worldwide (rice and wheat are the other two). The U.S. produces about 40 percent of world production in these two crops, making it, by far, the world's largest producer and exporter. While less important for global food supply, cotton is grown in the warmer Southern areas of the United States and might be better suited to warmer temperatures.

We found that yields of all three crops grow roughly linearly in temperature up to a threshold, above which yield growth declines sharply. The threshold varies by crop: 29C (84F) for corn, 30C (86F) for soybeans, and 32C (90F) for cotton. For all three crops, the slope of the decline above the optimum temperature for yield growth is significantly steeper than the incline below the optimum temperature. Cumulative exposure to cumulative heat above the threshold is the strongest single predictor of yield outcomes. One implication is that a modest amount of warming could change, markedly, the best locations for growing these key crops.

In this paper we extend the analysis and construct a fine-scaled weather data set for the entire 20th century in Indiana. This prolonged period covers weather extremes of the 1930s that led to the Dust Bowl and includes observation both before and after the Green Revolution, allowing to examine how the relationship between weather and corn yields evolved over time as new seed varieties (double- and single-crossed hybrids) were introduced. Historic adaptation to weather extremes, or the failure to do so, can give valuable insights on

how difficult it is to adapt to conditions that are predicted to become more frequent under climate change.

If we hold growing areas fixed and extrapolate from this relationship what yields might be as the climate warms, projected impacts are quite severe: average yields decrease by 30-46% before the end of the century under the slowest (B1) warming scenario and decrease by 63-82% under the most rapid warming scenario (A1FI) under the Hadley III climate model. These projected declines are driven by strongly negative yield growth when temperatures exceed 29-32 degrees Celsius combined with the sharp increase in the projected frequency of these extreme temperatures.

There are several reasons why these projected damages might overstate actual potential damages. As the climate warms, agricultural production will work to adapt to warming in a ways that mitigate losses and exploit potential gains as much as possible. The most difficult economic questions pertain to how large these adaptation possibilities may be.

One form of adaptation would be to change the locations where particular crops are grown. As climates change, so will geographical comparative advantages. We should not expect crops to be grown the same locations as they are grown today. Ascertaining the potential impact of climate changes therefore calls for analysis of yield potentials of major crops across the globe, even in places where agricultural production does not exist today. Such analysis can be quite complex and requires some strong assumptions about the potential suitability of many crops in different climates and soils. For example, there is uncertainty about soil dynamics in the Tundra, a region that is currently too cold to farm but might become farmable under warming. Chapin et al. (1995) conduct experiments of soil changes in Alaska and find that the 3-year response in experimental plots are a bad predictor of 9-year changes in experimental plots. The authors emphasize the difficulty of predicting long-term changes using short-term heat waves.

A recent study by the International Food Policy Research Institute (Nelson et al. 2009) conducts a comprehensive, worldwide analysis that incorporates shifts in growing locations. Given its inherent complexity, many assumptions enter into their model that are difficult to evaluate. The amount of uncertainty surrounding their projections is probably unquantifiable. But this is the most recent, careful, and comprehensive study to date. The study predicts significant declines in commodity production and increases in commodity prices stemming from global warming. Calorie availability will not only be less than the no-climate-change scenario, but less than availability in 2000. South Asia will be hit particularly hard as yields for rice and wheat decrease significantly.

In neither our earlier work nor in this paper do we attempt such a comprehensive analysis. Rather, by focusing on major crops in the U.S., a climatically diverse country that generates the world's largest agricultural output and exports the most, we examine forms of potential adaptation observable in historical data. These historical adaptations (or lack thereof) may provide some insight into the scope and nature of potential adaptations that may be available as climate changes.

Finally, we briefly discuss our impacts in the context of current agricultural subsidies in the United States. Some might argue that negative impacts to U.S. grain production will be welfare enhancing because they will offset positive production distortions stemming from subsidies. Journalists and authors like Michael Pollan have popularized the idea that our diet rich in corn-based processed foods and corn-fed beef, dairy, poultry and pork stems in part from the fact that corn production is so heavily subsidized.

While clear evidence on the net influence of subsidies on agricultural production is lacking, most subsidies today are more-or-less "decoupled" from production decisions and are therefore probably small in magnitude. An exception is recent subsidies for ethanol production, which has clearly stimulated demand for corn and caused commodity prices to increase. But with about a third of U.S. corn production going toward fuel, this has reduced the quantity of grains going toward food production. While the evidence is mixed, there is little reason to believe food production would be much less than it is currently in the absence of U.S. agricultural policies. Indeed, we show that U.S. conservation policies may reduce U.S. commodity production more than subsidies increase it.

2 Implications of Earlier Findings for Adaptation

Our earlier research found the same nonlinear relationship between yield growth and temperatures, described in the introduction, when the analysis is narrowed to consider only cooler northern U.S. states or only warmer southern U.S. states. This evidence supports the idea that the nonlinear temperature relationship is a generalizable phenomenon. Adding to this evidence, we found the same relationship if we examined only the early half of the sample (1950-1977), or only the latter half of the sample (1978-2006). This was particularly surprising given the significant increase in average yields.

These comparisons suggest that innovations since 1950, while increasing average yields approximately three-fold, did not increase relative heat tolerance. And since most regions of the U.S. currently have temperature distributions that are warmer than optimal, there has

existed at least some incentive to breed or engineer more heat tolerance into plants. Note that our earlier examination of heat tolerance over time was relatively crude: we merely split the sample into an earlier and later period. A key focus of new analysis presented below is to examine the evolution of heat tolerance more thoroughly and over a longer time period.

The stability the nonlinear temperature-yield relationship over different subsets of the data helps to provide powerful evidence of a causal link. This is particularly true as each specification includes county fixed effects to control for time-invariant heterogeneity of soils and farming practices. While cross-sectional variation in temperatures may be associated with other factors correlated with geography, county fixed effects purge this variation from the regression. Remaining variation in weather outcomes over time are arguably random from the vantage point of farmers, and thereby constitute a viable natural experiment. The stability of results combined with strong exogeneity of weather variations in a fixed location are what make the empirical results persuasive.¹

While correlations between time-series weather variations and economic outcomes are persuasively causal, a problem with focusing only time-series variations is that they cannot account for adaptation. When farmers operate in a different *climate*, the set of adaptation strategies will be very different compared to unanticipated changes in *weather*.

One might be tempted to interpret short-run response to weather as a useful lower bound of the impact stemming from climate change. The idea is that, with climate change, adaptations would mitigate damages and exploit new opportunities, and such adaptations are simply not available in the short run. Thus, the argument goes, adaptation necessarily improves the outcome relative to the short-run response to weather. In our view, such an inference is incorrect. It is true that some decisions are available in the long run that are unavailable in the short run. But the converse is also true. For example, an aquifer with limited replenishment may provide irrigation water to help a farmer cope with a temporary drought, but may be insufficient for maintaining crop production if precipitation were permanently reduced.²

The effects of changing climate conditions are better captured by cross-sectional comparisons. The potential downside is that cross-sectional comparisons are more easily confounded

¹Deschênes and Greenstone (2007) recently proposed to use year-to-year variation in weather to estimate the relationship between profits or yields and weather. They find that agricultural profits and yields are independent of weather. However, their weather data set contains many irregularities and their profit measure, which is the difference between sales in a given year minus expenditures, does not account for storage behavior that smooths profits between periods. Once the data errors are corrected, projected climate change effects on yields are again unambiguously negative (Fisher et al. 2009).

²Other examples are provided in Fisher et al. (2009).

by unobserved factors that happen to correlate with location. Since many economic and social factors correlate spatially with geography, and climate itself is spatially correlated with geography, there is a distinct possibility that any observed association between climate and an economic outcome is not causal, but rather reflects the influence of some unmeasured factor associated with location and climate.

Considering the strengths and weaknesses of both cross-sectional climate variations and time-series weather variations, in our view it is important to consider both. And in this respect perhaps the most compelling finding of our earlier research is that both methods of comparison give very similar results. We isolated the pure cross-sectional relationship between climate and yields by pairing the average distribution of temperature and precipitation outcomes with each county's average deviation from the nationwide U.S. yield.³ To isolate the pure time-series we paired the nationwide average yield with the crop-area-weighted average weather distribution in each year.

Both of these methods of identification show the same distinctly non-linear relationship between temperature and yield growth described in the introduction. While the cross-sectional relationship may be potentially confounded by omitted variables, it is robust to inclusion or exclusion of controls for soils and other factors. Moreover, we also find it especially unlikely that unobserved confounding factors would happen to align in such a manner that would give rise to the same nonlinear relationship as observed in the time series relationship that is identified with presumably random weather fluctuations. The fact that these relationships are similar suggests that, at least historically from 1950 to 2006, there has been little scope for adaptation conditional on the locations where these key crops were grown. This finding is consistent with some earlier work using the hedonic approach, which considers cross-sectional variations in climate to land values (Schlenker et al. 2006).⁴ The

³Subtracting each year's nationwide yield from each county's yield removes the aggregate upward trend, which is substantial.

⁴Mendelsohn et al. (1994) first introduced the Ricardian method to measure the effects of climate change on agriculture by estimating a cross-sectional relationship between county-level farmland values and climatic variables in the United States. The predicted impact of changing climatic variables depend largely on the set of weights. Under the cropland weights (fraction of a county that is cropland) the predicted impacts are severely negative, and under the crop-revenue weights (the value of agricultural production sold) the effects are beneficial. The reason why the results diverged under various weights is access to highly subsidized irrigation water rights in the Western United States. These subsidized water rights capitalize into farmland values (Schlenker et al. 2007). Since access to subsidized water rights is correlated with temperature, an increase in temperature implicitly assumes an increase in subsidies, which should not be counted as a societal benefit. The crop-revenue weights aggravate the problem because highly irrigated counties in the Western United States account for a large share of overall revenues, yet the fraction of the county that is cropland (cropland weights) is small. Schlenker et al. (2005) show that if the analysis is limited to rainfed agriculture, the results converge and become unambiguously negative under both sets of weights.

hedonic approach also accounts for crop switching in response to climate change.

3 The Evolution of Weather-Yield Relationships over the 20th Century

Our earlier work found little evidence of adaptation to warmer temperatures between 1950 and 2006. In this paper we extend the analysis to include the earlier and potentially more interesting period between 1901 and 1950. Our focus on this period is motivated in large part by Sutch’s (2008) research. Sutch argues that the adoption of hybrid corn, one of history’s most remarkable and well documented technological revolutions, was precipitated in part by the extreme weather events of the 1930s. In particular, he argues that hybrid corn demonstrated particularly high yields relative to open-pollinated (non-hybrid) corn during 1934 and 1936, which (by our own key crop-related temperature measures) remain the most extreme on record. Thus, it could be that our earlier analysis did not look back far enough to the timing of the key innovation leading to the green revolution.

Specifically, in this paper we examine a panel of corn yields from 1901-2005, a time period that includes a full 35 years before the beginning of the green revolution as well as some 70 years after the first adoption of hybrid corn. Our analysis focuses on the state of Indiana, which sits in the middle of the so-called “Corn Belt” and is the nation’s third largest corn growing state. Our focus on Indiana is mainly due to data availability: it turns out that Indiana has the most comprehensive record of detailed daily weather records in the station data maintained by the National Climatic Data Center. Detailed daily weather data are necessary to estimate the effect of the entire distribution of temperature outcomes on yields. The data accounts for variations in temperatures, both within and across all days of each growing season. This detail facilitates correct identification of nonlinear temperature effects, which can be diluted from measurement error, or if temperatures are averaged over time or space. The key focus of our analysis is to examine how heat tolerance and drought tolerance has changed over time, with some particular focus on the time period following the great heat waves of 1934 and 1936 and subsequent widespread adoption of hybrid corn.

3.1 Data: A Century of Yields and Weather in Indiana

Figure 1 shows corn yields in Indiana over the 20th century. These yield data are publicly available from the U.S. Department of Agriculture’s National Agricultural Statistical Service

(USDA-NASS). All of our data sources are described in further detail in the data appendix. The graph shows the average yield in the state for all years between 1901 and 2005 as black diamonds. For years after 1928, when county-level data becomes available, a box plot shows the range and inter-quartile range of yields across counties in Indiana.

Before 1940 there was no discernible trend in yields. This is true even if one were to extend the time series back many decades before 1901, the earliest year shown in the figure. Around 1940 yields started a sharp upward trend that appears ongoing even today. Typical yields in Indiana were between 30 and 40 bushels per acre before 1940, yet today, a typical Indiana farmer can expect 150 to 160 bushels per acre. Yield variance increased along with typical yields, so we model the natural log of yield per acre.

As discussed in great detail by Sutch (2008), the beginning of the upward trend in yields began around the time when many key events occurred simultaneously. The Great Depression in the 1930s was followed by the onset of World War II in 1938, which caused large fluctuations in commodity prices. At least equally important was the early adoption of hybrid corn, starting in Iowa and quickly expanding to Illinois, Indiana, and beyond. The superior yields of hybrid corn was discovered in 1918 but it was not until later, perhaps after 1936, that seed production became commercially viable and high-yielding enough for farmers to adopt.

Also, in the decade before 1940, the Midwest, including Indiana, experienced both the hottest and driest temperatures on record for the growing-season months between March through August, shown in Figure 2 and 3. The former shows yearly weather shocks in extreme heat (degree days 29°C , further described below) over the growing season. The latter shows precipitation deviations from average climatic conditions. The decade of poor weather in the 1930s was most accentuated in the two drought years of 1934 and 1936, which brought about the great Dust Bowl, an event of massive wind erosion in states west and south of Indiana. In those years average yields in Indiana were just 27.6 and 25.6 bushels per acre, two of the three worst yields on record for the state during the 20th century. Note that drought years also showed the largest exposure to extreme heat. Indiana still fared much better than states west and south of Indiana. Iowa harvested just 60 percent of its planted acreage in 1934, an all time low, and Dust Bowl states of Nebraska and Kansas lost nearly all of their corn plantings in these years.

It is interesting to note that more recently, and particularly in the last two decades, the weather has been good for corn yields. This is in sharp contrast with what climate models project in the decades to come. Under the slowest-warming scenario (B1) in the Hadley III

model, average projected extreme heat for the years 2070-2099 is predicted to increase by 103 degree days above 29°C compared to the 1960-1989 baseline. It is added as a horizontal line in Figure 2. While the B1-scenario assumes we curb CO₂ emissions sharply in the near future, the predicted increase is still worse than the worst of the Dust Bowl years. Average projected increase in extreme heat under fast-warming A1FI scenario is way off the chart at 310 additional degree days above 29°C.

Construction of the weather variables presented in Figure 2 and 3 is further detailed in the appendix. We construct these data from daily individual weather stations in Indiana. Geographical interpolation is achieved by linking it with the PRISM weather data sets, which gives monthly observations on a 2.5x2.5 mile grid for the entire United States. Indiana is the only state in the U.S. for which the National Climatic Data Center of the National Oceanic and Atmospheric Administration reports having more than three weather stations in the early part of the century. The availability of good, fine-sale weather data is essential for identifying non-linear weather effects since these effects can be diluted with measurement error or if values are averaged over time and space. The geographical locations of weather stations in Indiana that we use to construct our data set for each 25-year period are shown in Figure 4.

The challenge for a regression model that relates yields to weather outcomes is in mapping an entire season of temperature and precipitation outcomes to a single yield outcome. We achieve this by assuming temperature effects on yields are cumulative over time and that yield is proportional to total exposure. This implies temperature effects are additively substitutable over time. That is, we sum the daily outcomes associated with each temperature over all days of the growing season. The benefit of assuming additive separability is that it allows us to keep the underlying relationship between temperatures and yields fully flexible.

Earlier work has shown that there are three weather variables that give the best out-of-sample predictions of corn yields: (i) total precipitation p_{it} in county i in year t ; (ii) degree days above 29°C (dd_{it}^H), which captures the harmful effects of high temperatures; and (iii) degree days between 10°C and 29°C degrees (dd_{it}^M), which measures the beneficial effects of moderate temperatures (Schlenker and Roberts 2009). Each measure is simply a truncated integral over the temperature distribution within each day, and then summed over all days in the growing season, as given below.

Degree days above 29°C (high temperature measure) are defined as:

$$dd_{it}^H = \sum_{j=\text{March 1st}}^{\text{August 31st}} \int_{T=29}^{\infty} (T - 29)h_{itj}(T)dT$$

where T is temperature (in degrees Celsius) and $h_{itj}(T)$ is the estimated density of time at each degree during day j in year t in county i . Since the measure is sensitive to geographic variation in temperatures, as well as variations within and across all days of the growing season, we spend considerable care in estimating $h_{itj}(T)$. Further details are given in the data appendix.

The second temperature measure is degree days between 10°C and 29°C (moderate temperature measure) are defined as:

$$dd_{it}^M = \sum_{j=\text{March 1st}}^{\text{August 31st}} \int_{T=10}^{29} (T - 10)h_{itj}(T)dT.$$

3.2 Regression Model

In this paper we take as given the two temperature measures that our earlier work found to be the best predictor of corn yields in 1950-2005. Our focus is to explore how the relationship between yields and weather has changed over the 105 years from 1901 to 2005. We use a flexible restricted cubic spline model that allows temperature and precipitation associations to change smoothly and flexibly over time. Specifically, the regression model is:

$$y_{it} = \beta_0 dd_{it}^M + \beta_1 dd_{it}^H + f_p(p_{it}) + f_t(t) + f_M(t) * dd_{it}^M + f_H(t) dd_{it}^H + f_{t2}(t) f_{p2}(p_{it}) + c_i + \epsilon_{it}$$

where y_{it} denotes the natural log of yield in county i and year t , dd_{it}^M and dd_{it}^H are the degree day measures described above, p_{it} is total precipitation. The functions $f_x(\cdot)$ are cubic splines of time or precipitation.⁵ We also include separate intercepts for each county (i.e., fixed effects, denoted c_i) to account for unobserved time-invariant heterogeneity, like soil quality. Since we combine state-level averages before 1929 with county-level averages starting in

⁵In the baseline model each of the spline functions is approximated using 5 knots, located at the 0.05, 0.275, 0.5, 0.725 and 0.95 quantiles of the empirical distribution of the relevant explanatory variables. For the time trend, knot locations are 1932, 1949, 1967, 1984, and 2001. The early knot in the time trend is due to the fact that we have only state-level observations prior to 1929, and thus fewer data points per year than after 1929 when we have county-level observations. To check the stability of the results to specification, we also estimated models with 3, 4, 6, and 7 knots for each spline function below.

1929, we use the corn acreage as weights in the regression equation to make the two set of aggregation measures comparable. Estimation of restricted cubic spline models is easily done using ordinary least squares. Since the errors within each year are likely correlated in space, we adjust our standard errors to account for this (clustering the errors by year) and possible heteroscedasticity using the Huber-White method.

3.3 Results

Figure 5 shows the effects of each of the four variables: time, precipitation, dd^M and dd^H while holding all other three variables at their median values. These results are characteristically similar to what we found in our earlier work that focused on the period from 1950-2005: there is a sharp upward trend in yields over time as shown in the top left panel. Yields have an inverted-U shape with rainfall as shown in the top right panel. Yields increase gradually with temperate degree days between 10°C and 29°C as shown in the bottom left panel. Finally, yields decline sharply with extreme heat, measured as degree days above 29°C, as shown in the bottom right panel. Since all regressions include county fixed effects, the graph will be shifted up-or down by the county-specific intercepts. We hence normalize each graph and display impacts relative to optimal outcome of each variable in question. For example, by how much do yields decline if precipitation deviates from the optimum for the season. All four panels of Figure 5 use the same scale on the y-axis to make the contribution of each variable comparable across plots. The time trend is responsible for the largest effect followed by degree days above 29°C.

These median-value predictions, however, do not show how these relationships have changed over time. We explore how these relationships change over time in Figure 6 for precipitation, Figure 7 for extreme heat dd^H , and Figure 8 for moderate temperatures dd^M . Each of these figures plots the relationship of the three weather variables at fifteen points in time.⁶

The effects of all three weather variables have shifted markedly over time. Figure 6 shows that the influence of precipitation nearly vanished by 2000. Deviations from the optimal precipitation levels have limited effects on yields. We believe that two explanation are most likely responsible for the fact that yields are no longer directly linked to rainfall during the growing season. First, a lack of precipitation in the growing season might be counter-balanced with irrigation. Continued mechanization of agriculture has led to the gradual expansion of

⁶We report the confidence bands obtained from the R package *Design* when the errors were clustered for all observations in a year.

pivot irrigation systems that can provide supplementary water during especially dry periods. While only a minority of corn fields in Indiana have pivot irrigation systems, the ones that do are probably more prone to dryness or have sandier soils. Second, seed companies may have bred increased drought tolerance into corn plant varieties. While climate models vary considerably in their predictions for precipitation changes, with some forecasting increases and others decreases, evidence from weather and yields in Indiana suggest this may be of little economic consideration.

The evolution of heat tolerance in Figure 7 differs from that of precipitation. Heat tolerance increased until 1960 followed by a decline after 1960. Figure 9 shows the *marginal* effect of extreme heat, i.e., the slope of the regression line in Figure 7 over all years in our sample. The negative influence of an additional degree day above 29°C is lowest around 1960 and most damaging in recent years when corn varieties were optimized for maximum average yields. The magnitude of the negative coefficient on dd^H is nearly three times as large in 2000 as it is in 1960, and about twice as large in 1901 as compared to 1960. This result is qualitatively insensitive to how many knots we use in the spline once we include at least 4 knots to make the model flexible enough to capture the nonlinearities. Figure 10 replicates this analysis for the marginal effect of moderate temperature as measured by degree days between 10 and 29°C.

Estimated slopes in the early years of Figure 9 should be interpreted with some caution because there are much fewer data points before 1929 as only state level data are available. Our spline model places more emphasis where there are more data and linearizes the model in the tails of the data. Closer inspection of the data do suggest that much of the increase in heat tolerance actually took place between 1940 and 1960, rather than being a steady smooth trend up from 1901. This interpretation would be consistent with the relatively stable farming technologies between 1901 and 1936 and rapid technological progress after 1940. This would also be consistent with Sutch's historical account of the adoption of hybrid corn.

The most interesting and relevant finding that speaks to implications for climate change is the sharp decline in tolerance to extreme heat since 1960. This finding is a powerful counterpoint the apparent increase in drought tolerance. Under the latest climate change models, a sharp rise in maximum temperatures is predicted to significantly increase the occurrence of temperatures above 29°C. Since degree days above 29°C are a truncated temperature variable, modest shifts in the temperature distribution can have a large relative influence on this temperature measure. For example, a 1°C warming from 29.5°C to 30.5°C triples degree

days above 29°C. The historic average number of degree days above 29°C is 25 in Indiana. Under the Hadley II model (IS92a scenario) the number is predicted to increase by 19 at the end of this century. Under the much warmer Hadley III model, degree days above 29°C are projected to increase by 103 under the slow-warming B1 scenario. Thus, even under the slowest-warming scenario, typical weather outcomes in the latter part of this century will be far worse than the worst drought years in the historical record, 1934 and 1936 (refer to Figure 1). Under the fastest-warming A1FI scenario, degree days above 29°C are projected to increase by 310, making the measure in a typical year about 3.5 times worse than the worst year on record.

Finally, the relationship between the precipitation and log yields is significant, but the interaction between time and precipitation has a p-value of 0.06. For all factors besides precipitation, both the main effects and their non-linear interactions with time, are significant at the 5% level, suggesting that the relationship was not stable over the century but evolved. A summary of significance tests is reported in Table 1.

3.4 Discussion of New Results

The last section extended earlier research on the link between weather and yields by examining how key weather variables are associated with corn yields in Indiana over a longer time period 1901-2005. We use restricted cubic spline regressions to let the effect of precipitation, moderate heat, and extreme heat evolve smoothly over time in a flexible way. Results for each variable, while holding all other variables constant at their median observed outcomes, are comparable to earlier results we obtained for a model using county-level corn yields for all counties east of the 100 degree meridian in the years 1950-2005.

The median association, however, obscures significant evolution of precipitation and temperature effects over time, effects that we had not examined in our earlier research. The overall influence of precipitation during the growing season has diminished with time.⁷ We hypothesize that attenuation of precipitation effects stems from increased use of supplemental irrigation and possibly the development of more drought tolerant seed varieties and cropping systems that have increased planting densities and canopy cover of the soil.

The evolution of temperature effects looks rather different from that of precipitation

⁷A cross-validation analysis shows the fine-scale precipitation data to be less accurate than the fine-scale temperature data. Since error in an explanatory variable causes attenuation bias, it is likely that precipitation is more important in reality than our regressions imply. But since the data are likely *more* accurate in the recent period as compared to the earlier period, attenuation bias cannot explain the general trend of decreasing importance of precipitation.

effects. The evolution of heat tolerance over time is nonlinear, increasing sharply between about 1940 and about 1960 and then declining. We found corn in Indiana to be most sensitive to extreme heat in the more recent years of our sample. The later decline in heat tolerance might be due to the fact that maximizing corn plants for average yields also makes them more sensitive to suboptimal growing conditions. It is interesting to note that the key turning points in evolution of heat tolerance align almost perfectly with the adoption of double-cross hybrid corn (around 1940) and single-cross hybrid corn (around 1960). It is also notable that, from inspection of Richard Stuch's figure showing U.S. aggregate corn yields from 1866-2002 and our own Figure 1, yields became noticeably more variable as corn transitioned from double-cross to single-cross varieties, a pattern that could be indicative of greater heat sensitivity.

Why did we find relative heat tolerance to be stable in our earlier study and not in this one? We believe there are several interrelated reasons. First, our earlier study began in 1950, awhile after first adoption of hybrid corn and growth in heat tolerance, but well before hybrid corn had been universally adopted in all states. Second, we simply split the sample into two sub-periods, 1950-1977 and 1978-2005, while pooling all states east of the 100th median. Since different states adopted hybrid corn at different times, and heat tolerance grew and then declined, our regressions would have picked up average heat tolerance in each sub-period. When pooling all states, it is likely that average heat tolerance was about equal in these two sub-periods. Note that Indiana was relatively early on the adoption curve for hybrid corn.

Our new findings have mixed implications for climate change impacts: on the one hand, sensitivity to extreme heat is highest at the end of the sample and the one feature all climate models agree on is that these extreme heat events are likely to increase, even though the size of the increase varies tremendously between model and emission scenarios. On the other hand, there was a period between 1940-1960 when both heat tolerance and average yields increased at the same time. The question is whether recent increases in yields could only be achieved by making plants less heat resistant, or whether future breeding cycles can increase both heat tolerance and average yields at the same time.

4 Current Policy Distortions and Climate Impacts

One argument, notably popularized by Michael Pollan, is that we currently produce too much corn and soybeans, and this contributes to obesity and other problems. The claim is that

these commodities are used to produce an overabundance processed foods like corn syrup and of animal-based food products like meat, eggs, and cheese. All of these foods are unhealthy if over-consumed. Many experts do point to excessive sugar consumption (particularly corn syrup in sodas) and marbled fatty meat from corn-fed animals as important contributors to the obesity epidemic in the United States and other developed nations. But Pollan (and others) go much further (*New York Times Magazine*, October 12, 2003):

The rules of classical economics just don't seem to operate very well on the farm. When prices fall, for example, it would make sense for farmers to cut back on production, shrinking the supply of food to drive up its price. But in reality, farmers do precisely the opposite, planting and harvesting more food to keep their total income from falling, a practice that of course depresses prices even further. What's rational for the individual farmer is disastrous for farmers as a group. Add to this logic the constant stream of improvements in agricultural technology (mechanization, hybrid seed, agrochemicals and now genetically modified crops – innovations all eagerly seized on by farmers hoping to stay one step ahead of falling prices by boosting yield), and you have a sure-fire recipe for overproduction – another word for way too much food. All this would be bad enough if the government weren't doing its best to make matters even worse, by recklessly encouraging farmers to produce even more unneeded food. Absurdly, while one hand of the federal government is campaigning against the epidemic of obesity, the other hand is actually subsidizing it, by writing farmers a check for every bushel of corn they can grow. We have been hearing a lot lately about how our agricultural policy is undermining our foreign-policy goals, forcing third-world farmers to compete against a flood tide of cheap American grain. Well, those same policies are also undermining our public-health goals by loosing a tide of cheap calories at home.

If Pollan is correct, then perhaps climate change, by reducing corn and soybean yields, would actually improve welfare by making unhealthy meats and processed foods more expensive. In terms of economic welfare, the claim is that farmers currently produce more than socially optimal because crops are subsidized and because farmers fail to maximize profits.

To economists, there is an obvious logical problem with Pollan's argument: the first paragraph in the quote contradicts the second. If farmers produce more when prices are lower then how could they also produce more when the government writes "farmers a check

for every bushel of corn they can grow,” which effectively increases the price. There is a difference between a shift in the supply curve, stemming from technological advance, and a movement along the curve. There is, in fact, strong evidence supporting the idea that commodity supply curves slope upward (e.g., see Roberts and Schlenker (2009)).

So let’s consider the case of agricultural subsidies. The critical issues concern (1) how much more farmer produce as a result of subsidies, and (2) whether the effect of subsidies causes enough of a difference in food prices to affect consumers’ food choices. On both counts, there is little reason to believe agricultural subsidies have much influence on the obesity crisis. In response to a question about whether there is evidence to support the claim that agricultural subsidies contribute to obesity, Dan Sumner wrote (*Freakonomics* blog at the *New York Times* July 24, 2008. For analysis related to this quote, see Alston et al. (2007)):

The reasoning is that, although farm subsidies programs have made the price of corn and soybeans slightly cheaper for buyers in the U.S., the accompanying trade policies have raised the prices of sugar and dairy products. Furthermore, farm costs comprise such a small fraction of the retail price, the small farm price effects have tiny retail price impacts. Finally, in rich countries such as the U.S., buyers respond little to any food price declines or increases.

We might also note that Australia, which has no farm subsidies, just passed the U.S. as the fattest nation.

With respect to the first issue, the size of the production response to subsidies, we need to correct Pollan’s notion that U.S. farmers receive a check from the government for every bushel of corn they grow. This assertion is false. In reality, the nature of subsidies is quite complex, has changed a lot over time, and, particularly since the 1996 Federal Agricultural Improvement and Reform (FAIR) Act, typically has only oblique connections to the quantities that farmers produce. Since the FAIR Act, the bulk of payments have been (with a few important exceptions) “decoupled” from farmers decisions about which crops they choose to grow. Instead, farmers have been paid according to their historical “base” acreage in subsidized crops. And until 2002, payment rates were set according to yields farmers received in the early 1980s.⁸ The exceptions were that farmers had to keep their land in some kind of agricultural use (pasture and fallow cover crops do count as

⁸Payment rates were reset for some farmers in 2002, but they were still based on historical yields, not the yields farmers obtained in the year they received payment.

an agricultural use) and could not plant fruit or vegetable crops. The first restriction is innocuous and largely unenforceable. And leading experts believe the land on which grains have been historically produced would not be well suited to fruit for vegetable cultivation (Babcock 2006). But this particular restriction has been at the center trade-related debates about the impact of subsidies, particularly cotton.

While empirical research on the effects of agricultural subsidies continues, and no particularly compelling natural experiment has yet been found, the current consensus appears to be that total production impacts stemming from traditional subsidies are modest. There may, however, be a link between subsidy rates and the rate of farm size growth (Roberts and Key 2008).

The component of subsidies most connected to production quantities stems from loan deficiency payments. Historically these “loans” were nominally designed to help farmers better time the sale of their commodities relative to market conditions. But if farmers could not find a price for their commodities above the loan rate paid by the government, the government would effectively purchase commodities at the loan rate. Today the government no longer takes farmers’ production quantities as collateral; instead, in the event price falls below the loan rate, they simply pay farmers a deficiency payment equal to the difference between the loan rate and price the farmer received. Today loan rates are quite low relative to market rates and so the expected amount of this subsidy is small.

Another important feature of production-based subsidies is that they have generally been counter-balanced with required set-asides or payments that compensate farmers for retiring land from crop production since their inception in the 1930s. Demand for commodities is very inelastic, so the easiest way to support farm incomes and land values is to restrict the quantity produced. Before the FAIR Act there was some paid land retirement for conservation uses, plus farmers were often required to “set aside” (not plant) some portion of their base acres in order to receive production-based subsidy payments. By the time of the FAIR Act, land retirement was completely separated from other subsidies and focused more squarely on conservation. Today all land retirement, about 34 million acres worth, is part of the Conservation Reserve Program.

The goals of set aside and land retirement programs have been to stabilize commodity prices (or keep them higher) and for soil conservation.⁹ The critical issue to note is that these policies have always worked in tandem with production-based subsidies, and have

⁹Since the mid 1980s, the conservation goals have expanded to include other environmental objectives besides soil erosion, like improved water quality and wildlife preservation.

acted strongly to reduce production, perhaps even more than production-based subsidies (particularly the "decoupled" variety) have acted to increase production.

The data illustrated in Figure 11 provide an indication of the overall importance of set asides and land retirement programs. The figures plot total acres harvested in the United States from 1910 to 2006 together with failed cropland (crops planted but never harvested) and land in set aside and land retirement programs. Note that total cropland harvested has been relatively stable for nearly 100 years. The relative stability of acreage in production stands in contrast large fluctuations and a general trend downward in real commodity prices. While evidence strongly indicates supply is upward sloping, it is also very inelastic. We also note that most variation in cropland harvested is associated with equal and opposite changes in land placed in set-aside and retirement programs. While one cannot be sure about the causal link, the pattern suggests that without compensated land retirement there would be even more grain production. Another way to judge the magnitude of set-asides and conservation is to compare them to failed acreage (planted but not harvested) during the dust bowl years of 1934 and 1936.

Data from the National Resources Inventory (NRI), which tracks land use over time on about 800,000 fixed points throughout the United States, show that specific land use changes match the more recent aggregated fluctuations in Figure 11. Indeed, the NRI data show that transitions into and out of the Conservation Reserve Program constitute, by far, the largest component of land use changes in the United States between 1982 and 1997.¹⁰ Over this time frame, net conversions of cultivated cropland to less-intensive uses (conservation, uncultivated cropland, pasture or range) were ten times conversions of cropland to urban development (Lubowski et al. 2006a, Lubowski et al. 2006b).

With regard to how much commodity prices affect food demand, it is important to recognize how little commodities prices factor into the retail prices of food in developed nations. For example, about six pounds of corn are used to produce each pound of beef (Cornell University Cooperative Extension n.d.). Six pounds of corn amounts to a little more than one tenth of a bushel, which today costs less than four dollars. So at today's prices, the corn in a pound of beef costs about 40 cents, or 10 cents for the corn in quarter-pound hamburger. If corn prices were to triple to 12 dollars a bushel, the price of a hamburger would rise just 20 cents, all else the same. Clearly most of what goes into the retail price of most food products are the processing, transportation and retailing costs, which are not directly influenced by food commodity prices. Moreover, in the United States and other

¹⁰NRI data after 1997 have not been released by USDA's National Resource and Conservation Service.

developed nations, consumers tend to respond to lower prices by substituting food they eat at home with food eaten at restaurants, which does not necessarily affect how much they eat.

Today, the largest distortion in crop production likely comes from a 51 cents per gallon subsidy on ethanol, mandates for continued growth in ethanol production, coupled with restrictions on ethanol imports. Ethanol now consumes about a third of the U.S. corn crop, 13% of the world's corn, and about 5% of the world caloric production of corn, soybeans, wheat and rice combined (Roberts and Schlenker 2009). U.S. ethanol subsidies were one of several factors blamed for the nearly four-fold rise in worldwide commodity that occurred between the winter of 2006 and summer of 2008. Over this time frame, corn prices rose from less than \$2/bushel to nearly \$8/bushel, and commensurate price increases were seen for wheat, soybeans, and rice. While these subsidies are widely believed to have increased commodity prices and production, they have reduced the amount of commodities going toward food production. Thus, while ethanol subsidies likely have negative welfare impacts, climate-change impacts lowering yields would not attenuate these impacts.

The picture is entirely different in less developed nations. Almost half of the world, around 3 billion people, live on less than \$2.50 per day. Consequently, they spend a large portion of their income on food and buy raw commodities rather than processed foods or meat. To the poor, a four-fold increase in commodity prices translates into nearly the same proportional influence on prices they pay for food. In poor countries demand response to the price spike was more elastic than that it was in rich countries, sometimes to the detriment of meeting basic caloric needs. So while the recent price increases had little influence food consumption in the United States, high prices did contribute to problems of hunger and civil conflict in the poorest nations.

On a related note, many have claimed that agricultural subsidies in rich countries (and low commodity prices) are helping to keep the developing world poor. Jagdish Bhagwati regards this view as a "pernicious fallacy." He writes (Bhagwati 2005):

The removal of subsidies is desirable, as it promises aggregate income gains, and many economists have therefore campaigned against them for nearly four decades. But Oxfam and the heads of several international aid institutions have now added the twist that the removal of these subsidies will also help the poorest countries known as the "least developed countries." This is dangerous nonsense. The economists Alberto Vales and Alex McCalla have shown that as many as 45 LDCs, out of 49, are net food importers; and as many as 33 are net importers

of all agricultural products together. As prices rise with the removal of subsidies, surely the importers will be harmed, not helped, except in the singular cases where the importers switch to becoming significant exporters. The mistake on the part of Oxfam, the leadership at the World Bank and countless journalists who have followed in their wake is easily understandable but not to be condoned: They have simply assumed that what is true of cotton subsidies in the rich countries, because there are four African countries that rely on cotton exports, must be true more generally. This is the fallacy of using an unrepresentative sample!

And for the poor, cotton would seem to be somewhat less essential a commodity as compared to food commodities. We should keep in mind that inexpensive food historically has been the precursor to growth and prosperity. Indeed, it is inexpensive food that laid the foundation for the industrial revolution and escape from Malthusian cycles (Clark 2007). It is therefore hard to see how higher food prices could be good for developing nations where many still live near the fringe of survival. The recent spike in commodity prices made this connection between food prices and welfare in less developed nations painfully vivid. Thankfully, the prices have fallen since the summer of 2008.

U.S. agriculture, mainly due to its tremendous endowment of fertile cropland and efficient production systems, has a major influence on world prices of basic agricultural commodities. It is also likely that current subsidies to US agricultural producers cause many kinds of distortions to agricultural commodities markets. But placed in broader context, it is hard to see how the US would produce significantly less than it currently does if subsidies were diminished or removed. All of this is to point out that the importance of the United States to world food production is not simply an artifact of its subsidies. This means negative climate impacts to crop production in the United States would in no way be welfare improving. And because the United States is the world's largest exporter, and most poor countries are food importers, the burden of climate change impacts on US production will fall disproportionately on the least developed nations.

5 Conclusions

Since the late 1930s when US farmers began using hybrid corn, commercial fertilizers and other modern farming techniques, average crop yields in the United States and around the world have grown tremendously. Today corn yields in the United States equal more than four times the best yields of the 1930s. Yields of most other staple commodities have more

than tripled. Over the same time period, world population grew slightly less than three fold. Higher yields have brought lower commodity prices, which have relieved hunger in less developed nations and have fed a growing (and likely unhealthy) appetite for meat and processed foods in rich countries. Yield growth has probably also attenuated expansion of cropping areas and deforestation. Recent adoption of genetically modified seeds have spurred yield gains in developing nations that have adopted them (Qaim and Zilberman 2003) and may hold promise for further yield gains in both developed and developing nations.

But global warming now poses a significant threat to crop yields. Crop scientists have long predicted that warming will cause yield declines in tropic and subtropic regions of the world. Climates in these regions are already too warm for optimal growing conditions for most crops so further warming will not help. More recent evidence suggests warming will also harm yields in more temperate regions where current production is greatest. Our previous statistical analysis of the United States, by far the world's largest producer and exporter of agricultural commodities, is dismal. Holding growing areas fixed (an important caveat) we predict yield declines of 38-46% for soybeans and corn between 2070-2099 under the Hadley III slow warming scenario (B1, which presumes sharp reductions in CO₂ emissions), and declines of 75-82% under the Hadley III fast-warming scenario (A1FI, which presumes the fastest growth in CO₂ emissions). Projected declines in medium term (2020-2049) are also substantial, 18-23% under the slow-warming scenario and 22-30% under the fast-warming scenario. The largest driver behind these reductions is the predicted increase in very hot temperatures. It is important to note that these predicted declines are relative to what yields would be without climate change, not necessarily in comparison to what yields are today. They also hold growing locations fixed and do not account for CO₂ fertilization, which may increase yields.

One way of adapting to warmer climates will be to change the locations where crops are grown. Corn and soybean production is likely to shift northward toward traditional wheat growing regions, and wheat (perhaps) to areas that were not previously cropped. Given the world's currently most productive areas are predicted to be harmed significantly, it is not clear how much of these losses may be mitigated by crop switching.

A team of researchers at the International Food Policy Research Institute, led by economist Gerald Nelson, recently developed the most comprehensive analysis to date (Nelson et al. 2009). Their model accounts for yield effects, crop switching, trade and price effects throughout the world, but takes population and GDP as exogenous to agriculture, and does not account for sea-level rise, which could be important for rice production in south Asia.

They predict that by 2050 calorie availability "will not only be lower than in the no-climate-change scenario—it will actually decline relative to 2000 levels throughout the developing world."

Thus, at present, it would appear that technological solutions, in addition to crop switching, will be necessary to overcome anticipated impacts from global warming. It is in this vein that we have explored historical innovation as it relates to heat tolerance. In particular we examined the evolution weather effects on corn yields in Indiana, and how these effects have changed over time with adoption of new crop varieties and farming techniques.

Where sensitivity to extreme heat is critical determinant of corn yields, in recent research we have found this sensitivity to be generally similar in warmer southern states as it is in cooler northern states. Moreover, we found no evidence that warmer areas have adapted to warmer-than-optimal climates: the cross-section of yields and climate matches the link between yields and weather in a fixed location.

In this paper we present new evidence here that may be somewhat more encouraging. We find that, following the dust bowl years—the hottest, driest and lowest-yielding years on record—heat tolerance in corn grew markedly until about 1960. After 1960, however, heat tolerance declined, even though average yields continued their steady rise. At the end of our sample in 2005, corn appears to be less tolerant to extreme heat than it was in the 1930s.

The key question is whether plant scientists and seed companies can continue to breed or engineer crops that have both greater yield potential and greater tolerance to extremes. At present these prospects seem uncertain, and greater agricultural productivity investments would seem prudent. The private sector may foresee higher future commodity prices and thus engage in these investments on their own. There may also be a role for public sector investments in basic research, particularly since these have been the source of critical innovations in the past. Such innovations have important positive spillovers that can lead to suboptimal private investment.

On the demand side, we believe it important to recognize that global income inequality is a critical obstacle to adaptation. The issue is not so much whether it will be technically feasible to feed the world's population, we see little doubt that there will be. But when median incomes of the richest nations are hundred times those of the poorest nations, it is easy to see how lower commodity supply combined with, say, a taste for meat in rich countries, could drive prices of staple commodities to the point that the poorest simply cannot afford to survive. Despite the necessity of food, demand response of the poor is larger than that of rich due to a much larger income effect.


There is no market failure or malthusian cycle in this story. It's simply a matter of income inequality. If incomes were not so divergent, prices would simply rise until enough people substituted to a presumably more healthy diet with less meat. The main reason climate change impacts on agriculture pose such a great threat lies not just in the size of potential production impacts, but also because massive income inequality limits potential adaptation on the demand side of the market. The greatest hope is an uncertain one: that technological change will obviate the need for behavioral change.

References

- Alston, Julian M., Daniel A. Sumner, and Stephen A. Vosti**, “Farm Subsidies and Obesity in the United States,” *Agricultural and Resource Economics Update - Gian-nini Foundation of Agricultural Economics, University of California Berkeley*, November/December 2007, 11 (2), 1–4.
- Babcock, Bruce A.**, “Cheap Food and Farm Subsidies: Policy Impacts of a Mythical Connection,” *Iowa Ag Review - Center for Agricultural and Rural Development*, Spring 2006, 12 (2), 1–3.
- Bhagwati, Jagdish**, “Reshaping the WTO,” *Far Eastern Economic Review*, January/February 2005, 168 (2), 25–30.
- Chapin, F. Stuart, Gaius R. Shaver, Anne E. Giblin, Knute J. Nadelhoffer, and James A. Laundre**, “Responses of Arctic Tundra to Experimental and Observed Changes in Climate,” *Ecology*, April 1995, 76 (3), 694–711.
- Clark, Gregory**, *A Farewell to Alms: A Brief Economic History of the World*, Princeton, NJ: Princeton University Press, 2007.
- Cornell University Cooperative Extension**, “On average, how many pounds of corn make one pound of beef?” <http://www.extension.org/faq/4027>.
- Deschênes, Olivier and Michael Greenstone**, “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather,” *American Economic Review*, March 2007, 97 (1), 354–385.
- Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker**, “Agriculture and Climate Change Reconsidered,” *Department of Agricultural and Resource Economics, UCB. CUDARE Working Paper 1080.*, May 2009.
- Lubowski, Ruben N., Marlow Vesterby, Shawn Bucholtz, Alba Baez, and Michael J. Roberts**, “Major Uses of Land in The United States, 2002,” Technical Report, Economic Information Bulletin Number 14, Economic Research Service, United States Department of Agriculture May 2006.
- , **Shawn Bucholtz, Roger Claassen, Michael J. Roberts, Joseph C. Cooper, Anna Gueorguieva, and Robert Johansson**, “Environmental Effects of Agricultural Land-Use Change: The Role of Economics and Policy,” Technical Report, Economic Research Report Number 25, Economic Research Service, United States Department of Agriculture August 2006.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw**, “The Impact of Global Warming on Agriculture: A Ricardian Analysis,” *American Economic Review*, September 1994, 84 (4), 753–771.


- Nelson, Gerald C., Mark W. Rosegrant, Jawoo Koo, Richard Robertson, Timothy Sulser, Tingju Zhu, Claudia Ringler, Siwa Msangi, Amanda Palazzo, Miroslav Batka, Marilia Magalhaes, Rowena Valmonte-Santos, Mandy Ewing, , and David Lee, “Climate Change: Impact on Agriculture and Costs of Adaptation,” Technical Report, International Food Policy Research Institute September 2009.
- Qaim, Martin and David Zilberman, “Yield Effects of Genetically Modified Crops in Developing Countries,” *Science*, 7 February 2003, *299* (5608), 900–902.
- Roberts, Michael J. and Nigel Key, “Agricultural Payments and Land Concentration: A Semiparametric Spatial Regression Analysis,” *American Journal of Agricultural Economics*, August 2008, *90* (3), 627643.
- and Wolfram Schlenker, “World Supply and Demand of Food Commodity Calories,” *American Journal of Agricultural Economics*, December 2009, *91* (5).
- Schlenker, Wolfram and Michael J. Roberts, “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change,” *Proceedings of the National Academy of Sciences*, September 15 2009, *106* (37), 15594–15598.
- , W. Michael Hanemann, and Anthony C. Fisher, “Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach,” *American Economic Review*, March 2005, *95* (1), 395–406.
- , — , and — , “The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions,” *Review of Economics and Statistics*, February 2006, *88* (1), 113–125.
- , — , and — , “Water Availability, Degree Days, and the Potential Impact of Climate Change on Irrigated Agriculture in California,” *Climatic Change*, March 2007, *81* (1), 19–38.
- Snyder, R. L., “Hand Calculating Degree Days,” *Agricultural and Forest Meteorology*, October 1985, *35* (1-4), 353–358.
- Sutch, Richard C., “Henry Agard Wallace, the Iowa Corn Yield Tests, and the Adoption of Hybrid Corn,” *NBER Working Paper 14141*, June 2008.

Figure 1: Average Yields in Indiana 1901-2008




Notes: Graph shows history of corn yields in Indiana. State level averages are shown as diamonds. The range of yields among Indiana's counties is shown as boxplots: The box give the 25%-75% quartile range, the median is shown as a solid line, and whiskers extend to the minimum and maximum. A locally weighted regression (bandwidth of 10 years) is shown as black line.

Figure 2: Shocks in Extreme Heat in Indiana 1901-2005



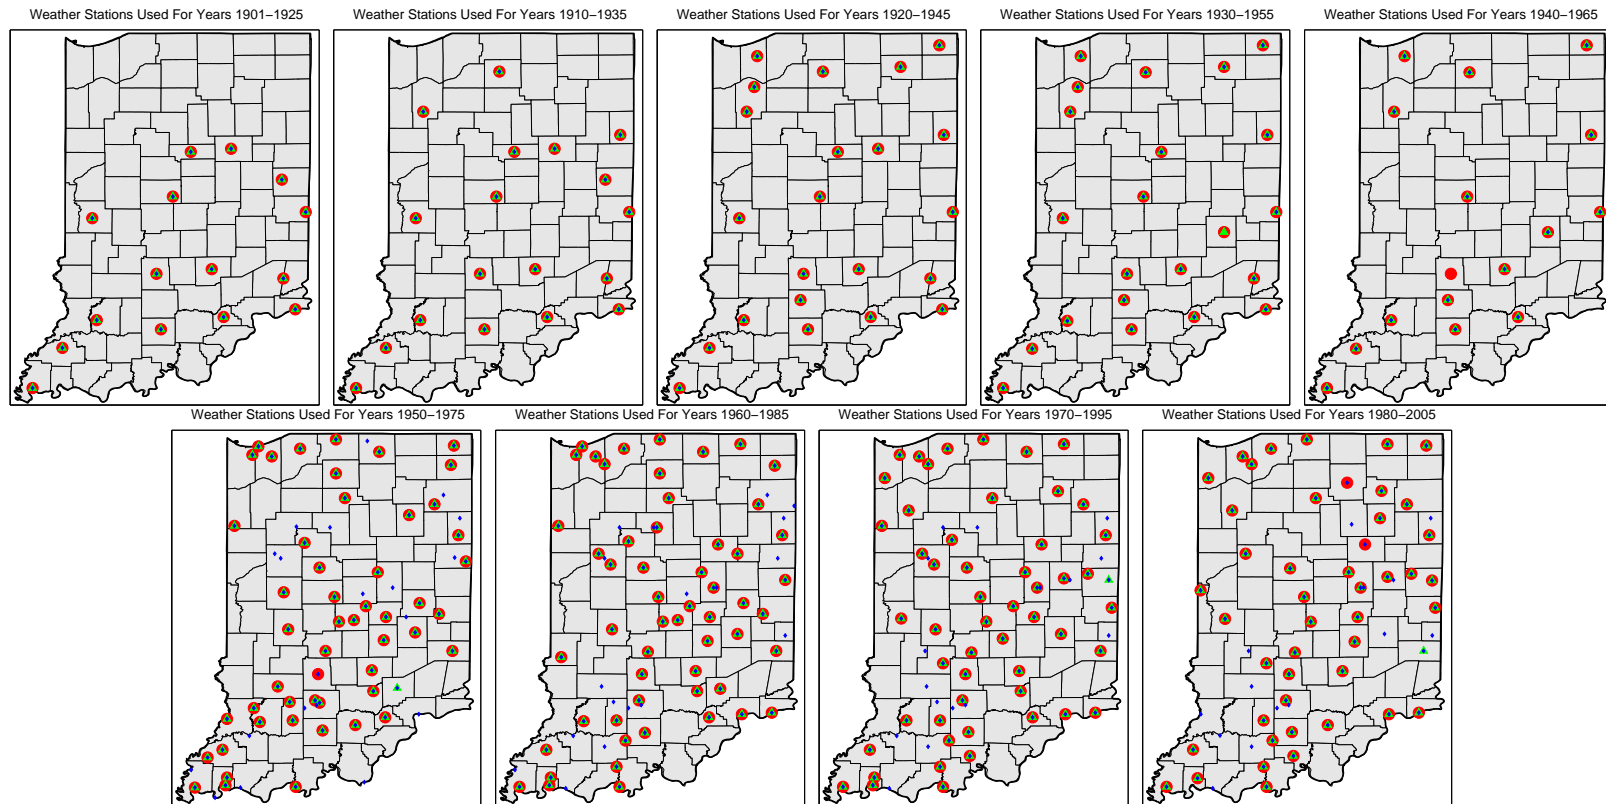
Notes: Graph shows weather shocks (deviations from averages) for degree days above 29C during the growing season March-August. State level averages are shown as diamonds. The range of weather shocks among Indiana's counties is shown as boxplots: The box give the 25%-75% quartile range, the median is shown as a solid line, and whiskers extend to the minimum and maximum. A locally weighted regression of the state averages (bandwidth of 10 years) is shown as black line.

Figure 3: Precipitation Shocks in Indiana 1901-2005




Notes: Graph shows weather shocks (deviations from averages) for total precipitation during the growing season March-August. State level averages are shown as diamonds. The range of weather shocks among Indiana's counties is shown as boxplots: The box give the 25%-75% quartile range, the median is shown as a solid line, and whiskers extend to the minimum and maximum. A locally weighted regression of the state averages (bandwidth of 10 years) is shown as black line.

Figure 4: Weather Station in Indiana By Time Period.



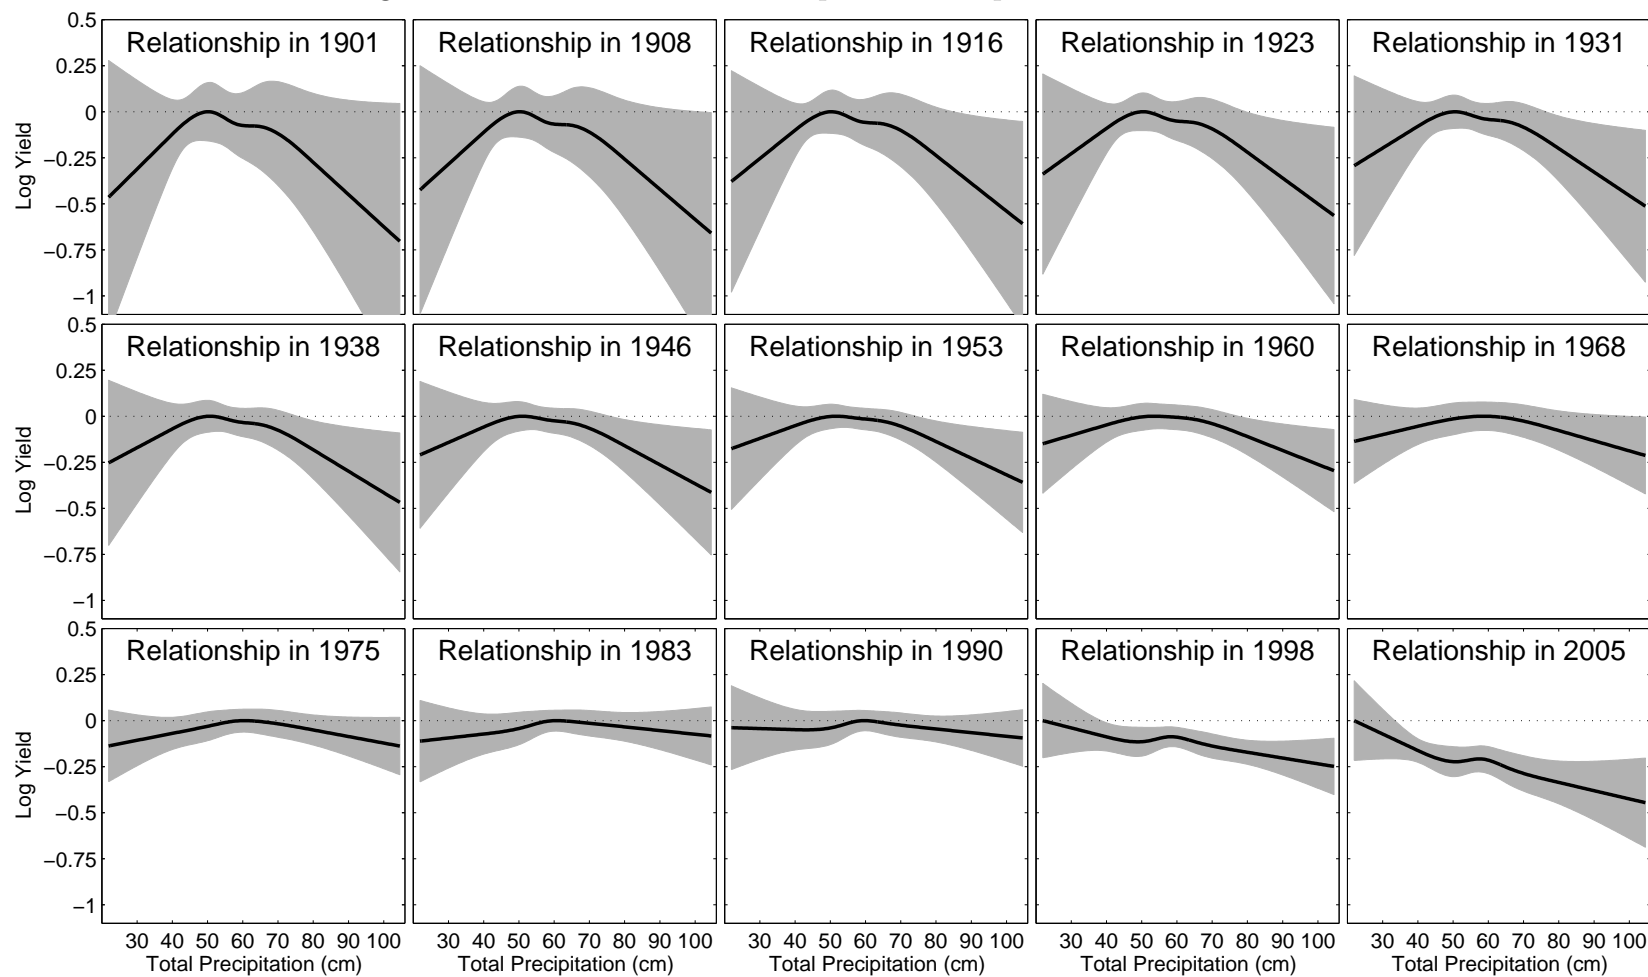
Notes: Weather stations used in interpolation are displayed as red dots for minimum temperature, green triangles for maximum temperature and blue diamonds for precipitation.

Figure 5: Regression Results at Median Values



Notes: Graphs display regression results of one variable while keeping other variables at median outcomes (shown as vertical line in each graph): top left panel shows yield trend, top right panel shows effect of precipitation, bottom left panel shows effect of degree days 10-29°C and bottom right panel shows effect of degree days above 29°C. Graphs are normalized relative to the optimal outcome of a variable.

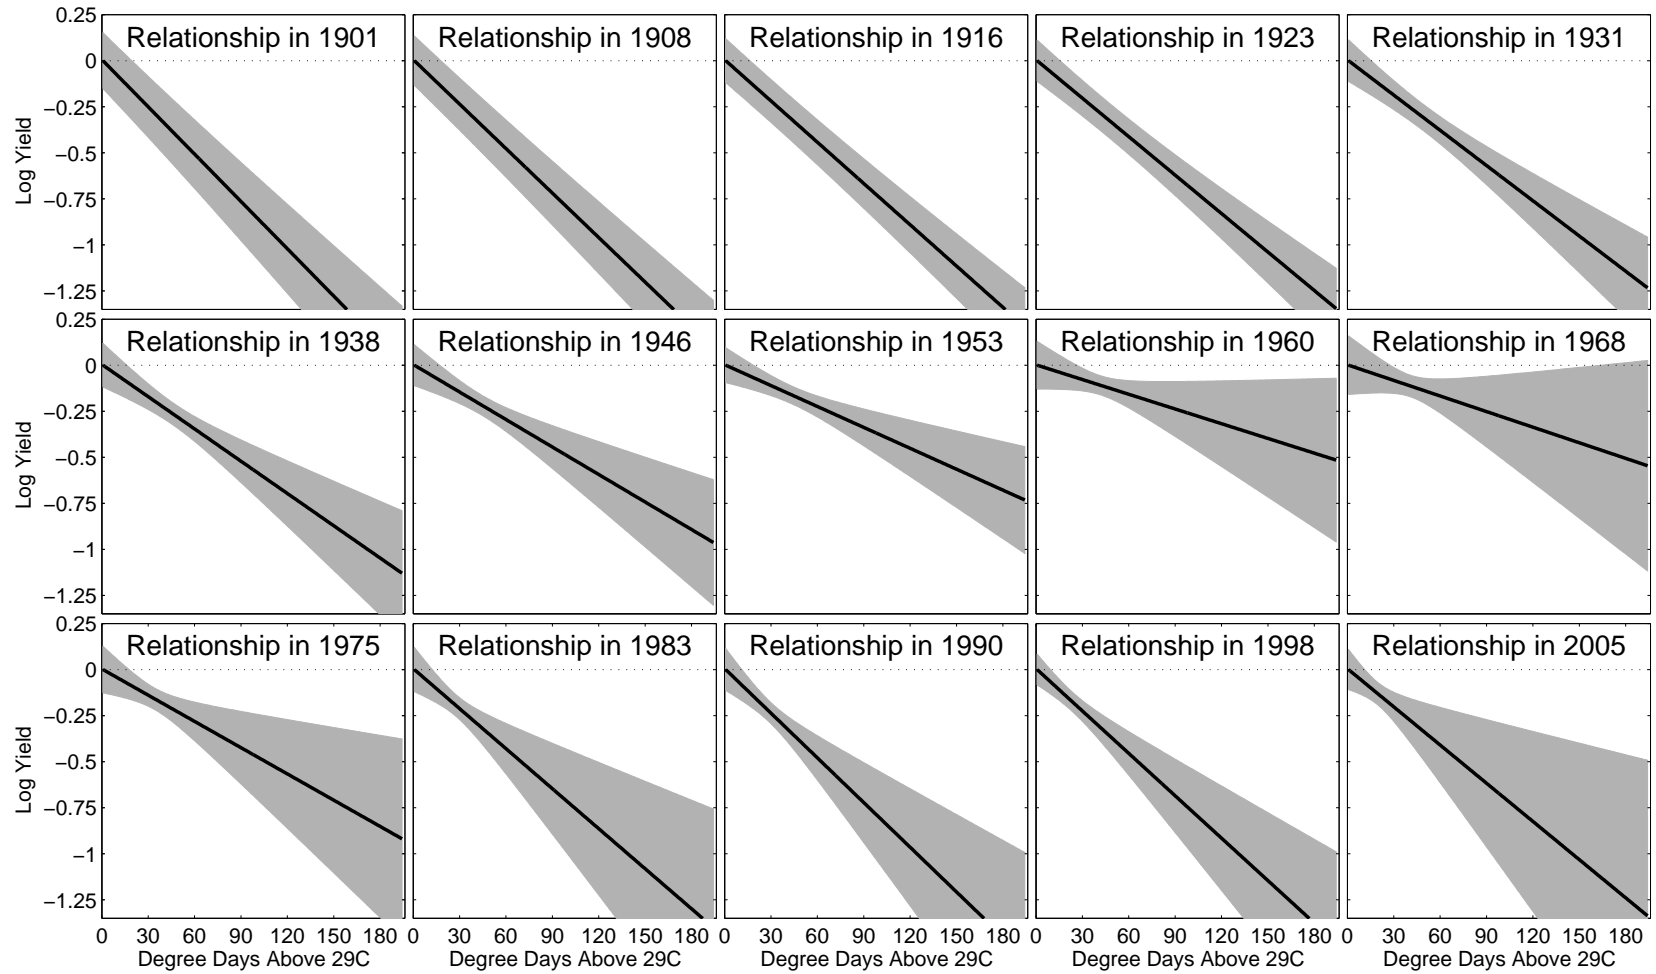
Figure 6: The Evolution of the Impact of Precipitation on Corn Yields



30

Notes: Graphs displays the effect of total precipitation during the growing season on log yields at fifteen periods in time. Graphs are normalized relative to the best value of each variable.

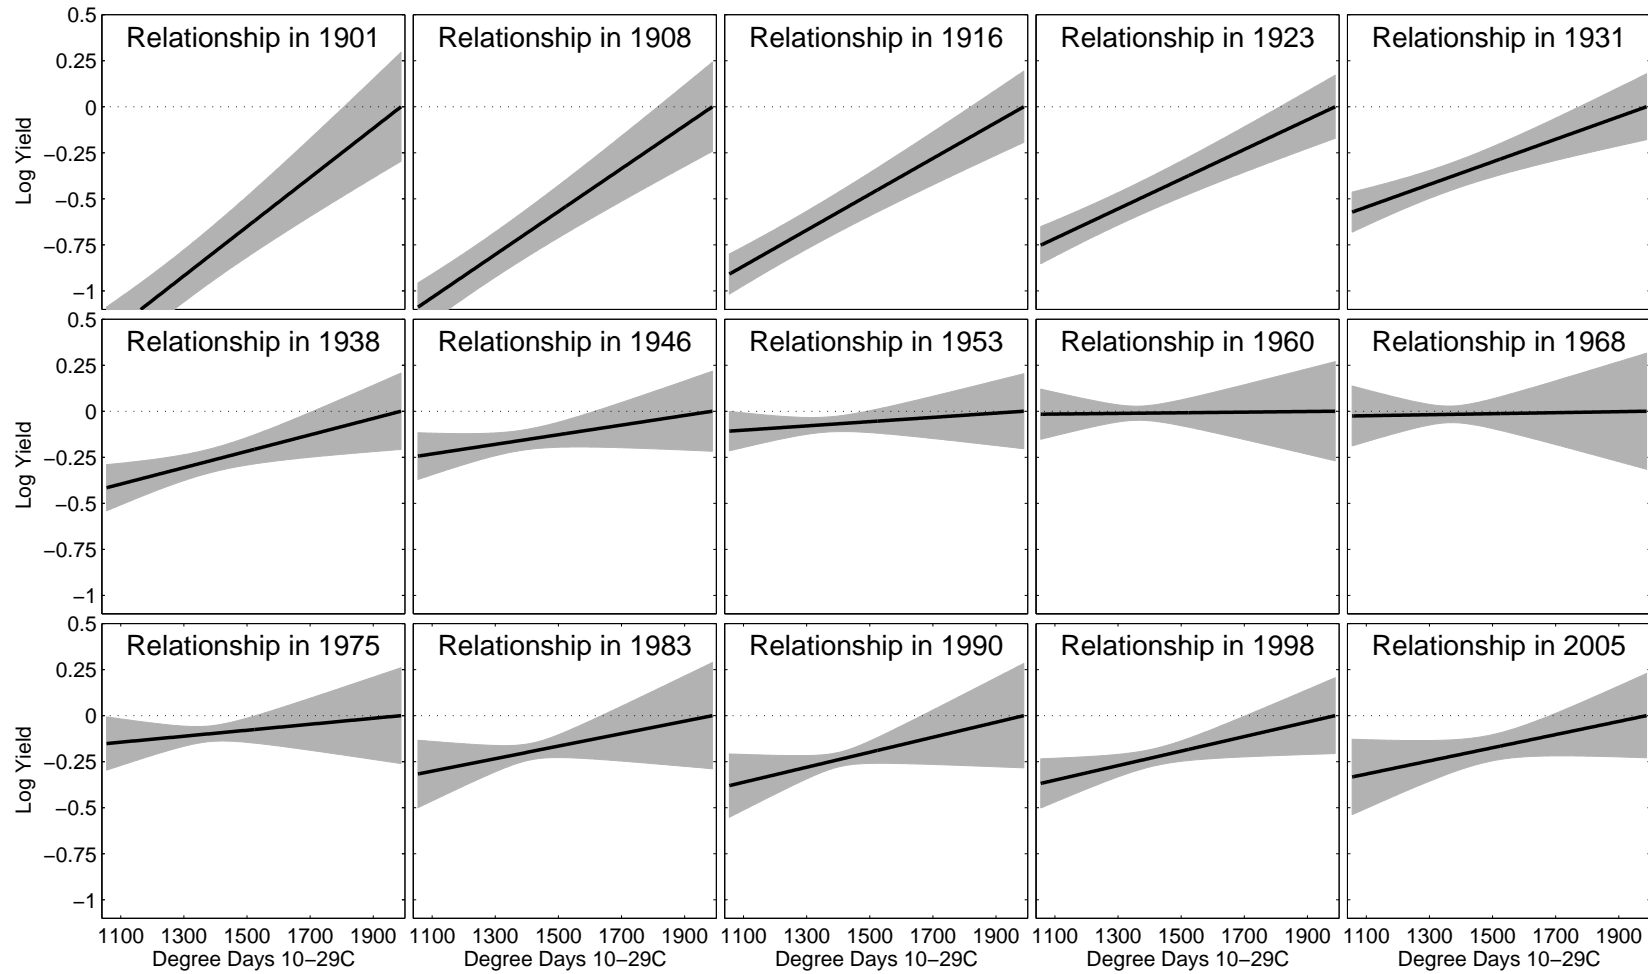
Figure 7: The Evolution of the Impact of Extreme Heat on Corn Yields



13

Notes: Graphs displays the effect of extreme temperatures during the growing season on log yields at fifteen periods in time. Graphs are normalized relative to the best value of each variable.

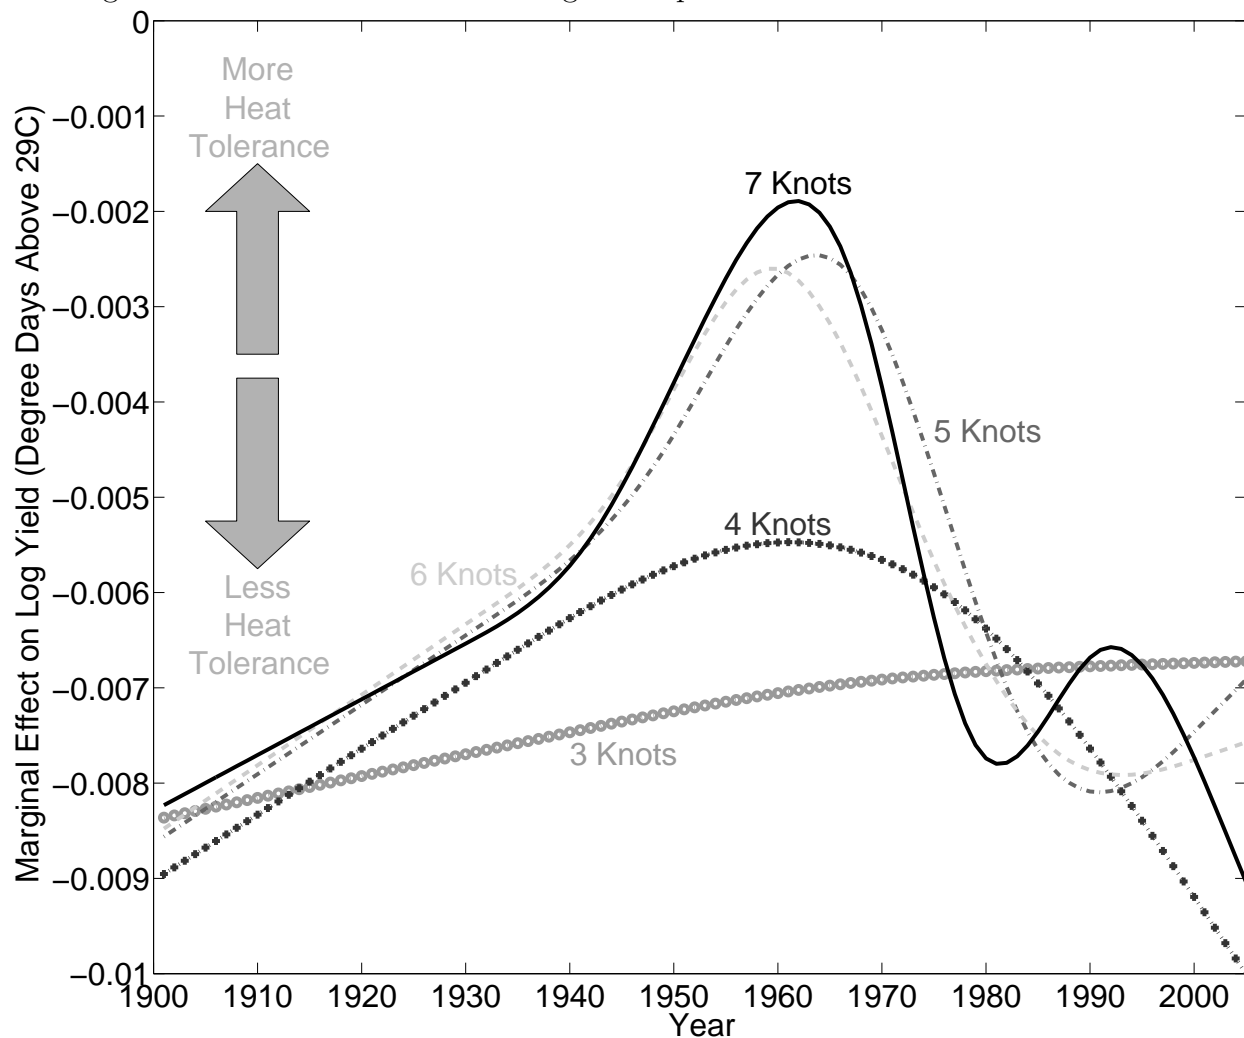
Figure 8: The Evolution of the Impact of Moderate Heat on Corn Yields



32

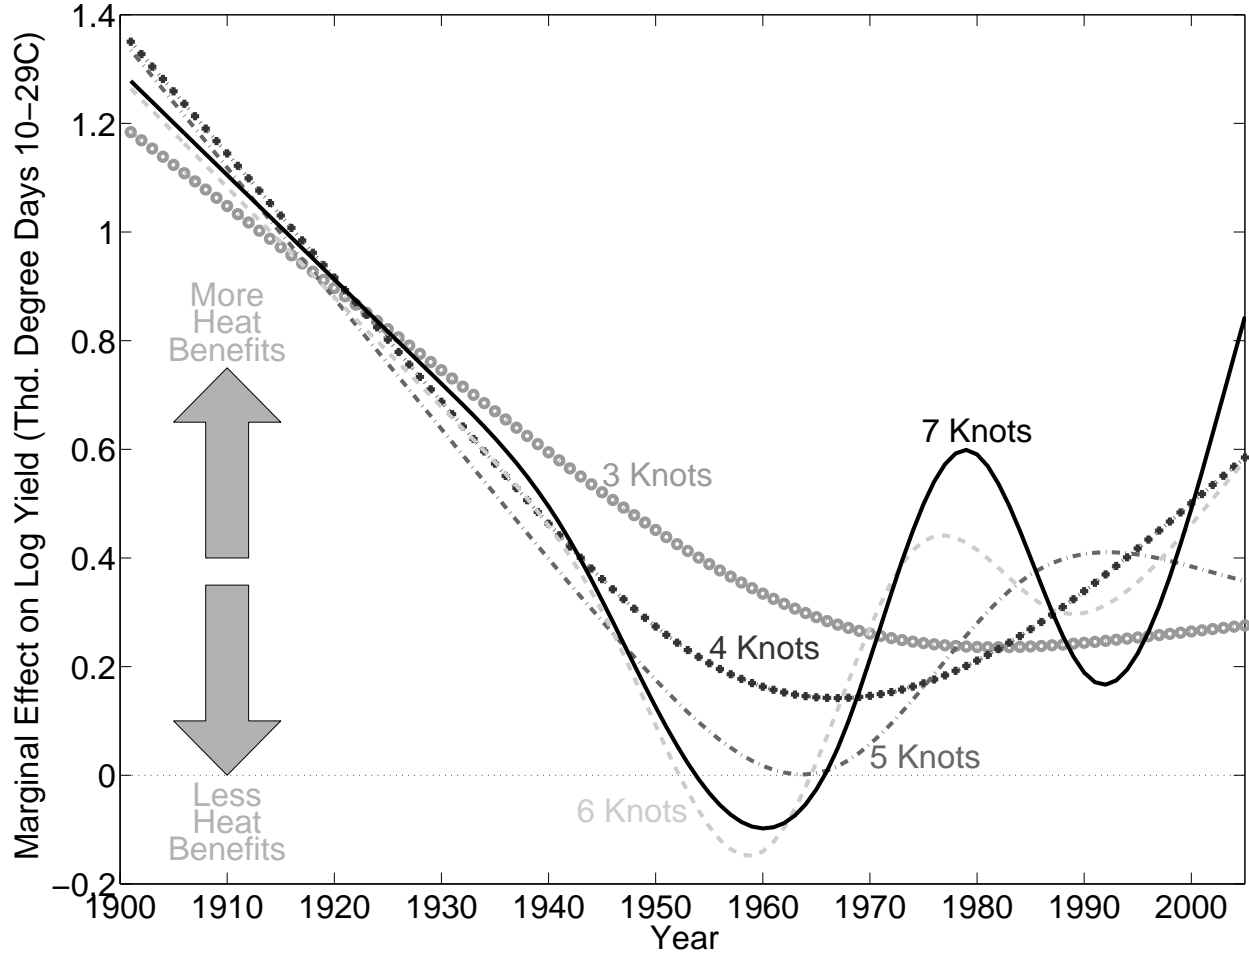
Notes: Graphs displays the effect of moderate temperatures during the growing season on log yields at fifteen periods in time. Graphs are normalized relative to the best value of each variable.

Figure 9: The Evolution of the Marginal Impact of Extreme Heat on Corn Yields



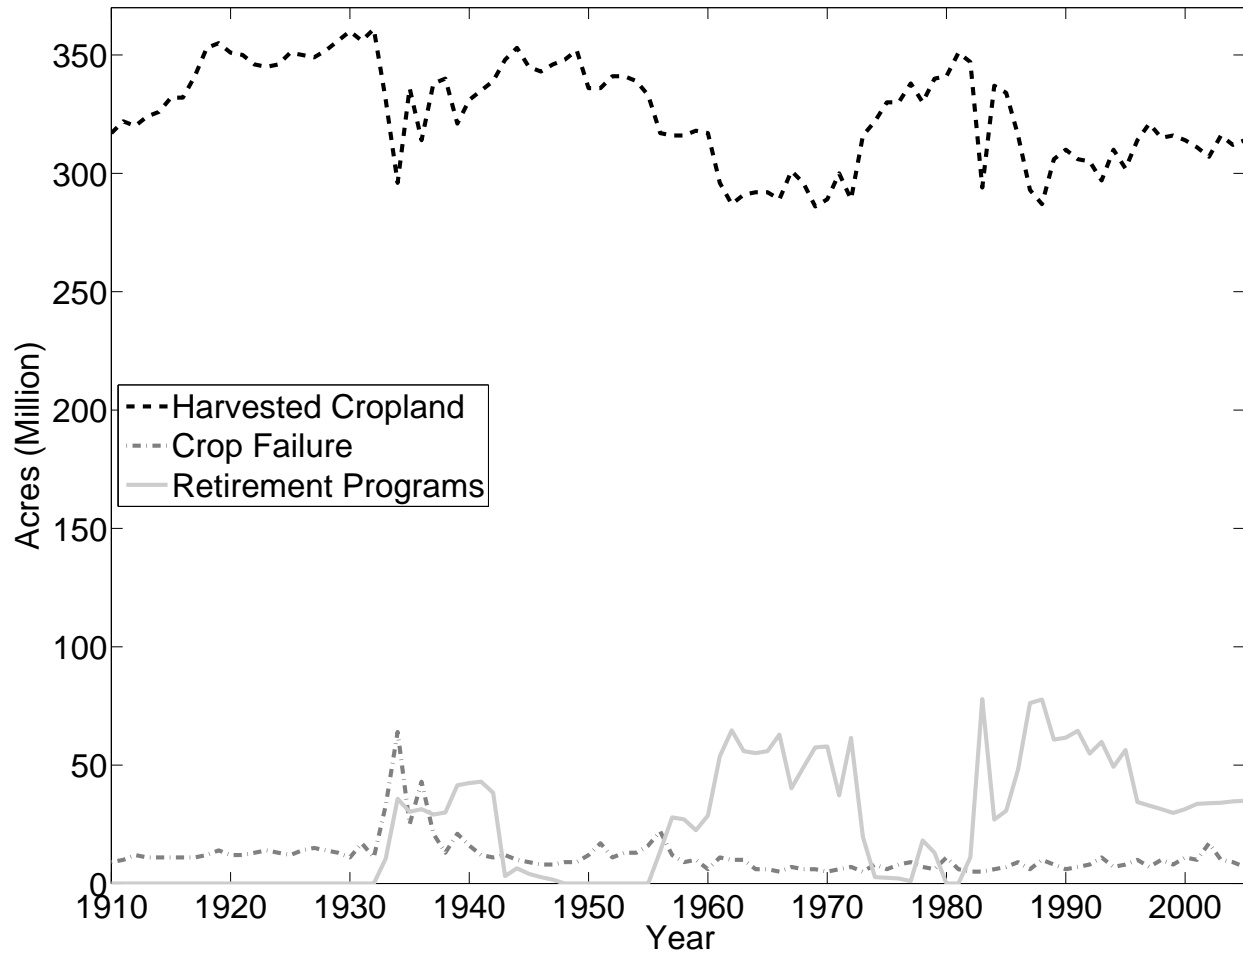
Notes: Graph displays the marginal effect of extreme temperatures (degree days above 29°C) on log yields, i.e., the slope of the regression lines in Figure 7. Cubic splines with various number of knots are used.

Figure 10: The Evolution of the Marginal Impact of Moderate Heat on Corn Yields



Notes: Graph displays the marginal effect of moderate temperatures (thousand degree days 10-29°C) on log yields, i.e., the slope of the regression lines in Figure 7. Cubic splines with various number of knots are used.

Figure 11: Cropland Failure and Retirement Programs in Indiana



Notes: Graph shows the total acreage in Indiana that is harvested, the planted acreage that is *not* harvested due to crop failure, and the acreage set aside in retirement programs.

Table 1: Analysis of Variance for Log Yield

	<i>d.f.</i>	<i>F - stat</i>	<i>p - value</i>
Degree Days Above 29C			
All Terms	5	18.99	<0.0001
Interaction Terms with Time	4	2.82	0.0289
Degree Days 10-29C			
All Terms	5	5.07	0.0003
Interaction Terms with Time	4	3.30	0.0138
Precipitation			
All Terms	11	4.09	0.0001
Interaction Terms with Time	7	2.00	0.0624
Time Trend			
All Terms	19	83.80	<0.0001
Time Trend Only	4	9.08	<0.0001

Notes: Table reports F-tests for the joint significance of key explanatory variables and their interactions with time. Our baseline model uses restricted cubic regression splines with 5 knots, which will result in four factors (variables) in the regression equation. The weather variables **Degree Days Above 29C** and **Degree Days 10-29C** consist of the weather variable (1 degree of freedom - d.f.) as well as the interactions with the four time factors (4 d.f.). The weather variable **Precipitation** consists of four factors in the amount of precipitation (4 d.f.) as well as the interaction of the linear time and precipitation term (1 d.f.) and the interaction of precipitation with the three higher order precipitation factors and vice versa (3 d.f. each). Finally, the **Time Trend** consists of 4 factors, and it is interacted with the 15 terms outlined for the three weather variables.

We used both the STATA routine *mkspline* as well as the R-package *Design*. The point estimates are identical, but the clustering option is implemented differently in both languages. We report the results from STATA, which tend to be more conservative (with the exception of precipitation).

APPENDIX

A1 Data Appendix

This appendix outlines in further detail how we construct our data set.

A1.1 Yield Data

Yield data was obtained from the National Agricultural Statistics Service (accessed March 2009). Yearly state-level yields in Indiana are available from 1866 onwards.¹¹ County-level yields in Indiana are available starting in 1929.¹² We follow the definition of the Department of Agriculture and calculate yields as the ratio of total production divided by area harvested.

The traditional definition of yields might overstate actual yields if some fields are not harvested. In a sensitivity check we define yields as total production divided by all acres planted. Unfortunately, area planted is only available from 1926 onwards for state totals and from 1972 for individual counties and hence significantly reduces our sample period. The left panel of Figure A1 displays the fraction of the planted area that was harvested in Indiana over time. While there is an upward trend, especially during the 30s, the right panel shows that the year-to-year variation in yields is similar for each definition of yields.

A1.2 Weather Data

Degree days were constructed from daily weather data. We obtained daily observations from the National Climatic Data Center Cooperative station network.¹³ The data include *daily* minimum and maximum temperature as well as precipitation. While the NCDC data has great *temporal* coverage, we combine it with the PRISM weather data set that provides better *spatial* coverage.¹⁴ The latter gives monthly minimum and maximum temperatures on a 2.5x2.5 mile grid for the United States from 1895 onward.

To construct a consistent set of weather data, we followed the following procedure for each 25-year period starting in 1901, 1910, 1920, 1930, ..., 1980.

- (i) For each of our three weather variables (minimum and maximum temperature as well as precipitation) we determine the set of stations with a consistent record, which we

¹¹http://www.nass.usda.gov/QuickStats/Create_Federal_All.jsp

¹²http://www.nass.usda.gov/QuickStats/Create_County_All.jsp

¹³<http://ols.nndc.noaa.gov/plolstore/plsql/olstore.prodspecific?prodnum=C00447-CDR-S0001>

¹⁴<http://www.prism.oregonstate.edu/>

chose to be stations that moved at most by 2.5 miles during the time period and had at most three missing values in at least 90 percent of the months.

- (ii) We fill the missing observations at stations with consistent records obtained in step (i) by regressing daily values at each station on daily values at the seven closest stations including half-month fixed effects. We use a linear regression for minimum and maximum temperature and a Tobit regression for precipitation, which has several observations at the truncation value of zero. Intuitively, the regression estimates are used to fill the missing values with a weighted average of surrounding stations with non-missing observations to give us a complete weather record at the stations with consistent weather records.
- (iii) We calculated monthly averages for the stations with consistent records in step (i)
- (iv) We regress the monthly values at each PRISM grid on the monthly averages at the seven closest weather stations from step (iii) including month fixed effects, again using a linear regression for minimum and maximum temperature and a Tobit model for precipitation. The R-squares are generally in excess of 0.999, suggesting that the PRISM data set is a weighted average of individual stations and we uncovered the weights.
- (v) We apply the regression results from step (iv) to the daily weather station data from step (ii) to derive daily weather measures at each 2.5x2.5 mile PRISM grid cell.
- (vi) We fit a sinusoidal curve between the minimum and maximum temperature of each day to calculate degree days accounting for the within day distribution of temperatures (Snyder 1985). We evaluate degree days for each bound between -5°C and $+50^{\circ}\text{C}$ using 1° steps at each 2.5x2.5 mile PRISM grid.

Once we have the daily observations on the PRISM grid, we aggregate them spatially

- (vii) We obtained the fraction of each PRISM grid cell that is cropland from a one-time Landsat satellite scan in 1992. County-level weather variables are the cropland-weighted average of all PRISM grid cells in a county.
- (viii) State-level weather data are the weighted average of all county-level measures in step (vii), where the weights are the amount of harvested corn area reported in the yield data. Since harvested corn area is not reported on a county-level before 1929, we use

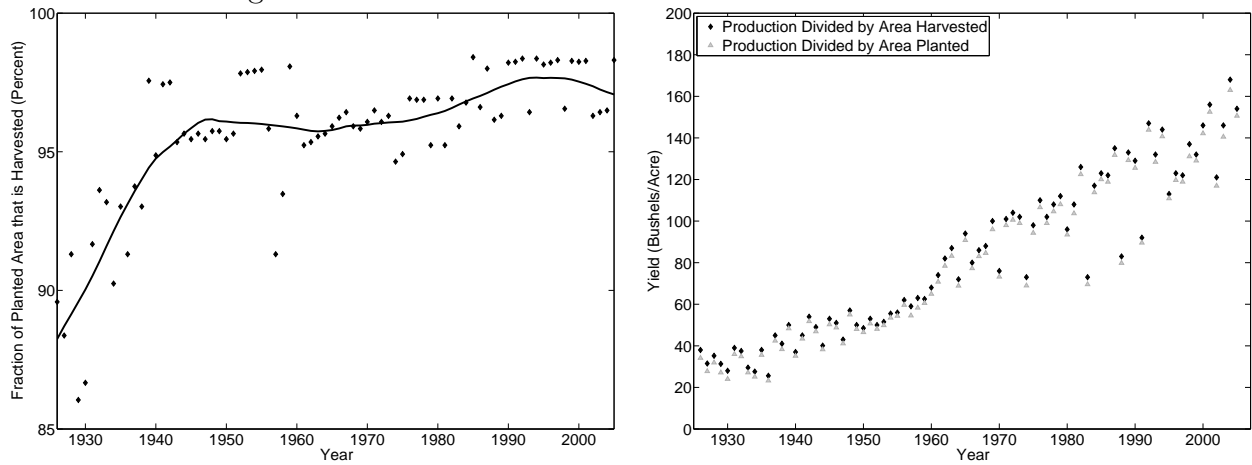
the average harvested corn area in each county in the years 1929-2005 as weights for years prior to 1929.

Finally, we aggregate the data temporally

- (ix) We define the growing season as the months March through August and add degree days as well as precipitation for all days in these months. Since total precipitation over the growing season is insensitivity to the within-day and between-day distribution, we use the monthly totals in the PRISM data set. For possibly daily interactions between precipitation and temperature we use the interpolated daily precipitation data.


Since it was impossible to get a sufficiently large set of weather stations which had consistent nonmissing records for the entire sample period 1901-2005, we instead derived the measure for 25-year intervals, starting in 1901, 1910, 1920, up to 1980. The results of interpolation series for extreme heat in the state of Indiana (degree days above 29°C) are displayed in colors in Figure A2. They appear to overlap tightly. One might still wonder whether the state results hide the fact that there are substantial errors in the county level data that get averaged out. To examine this further, we take the difference of all overlapping series in the county data. The mean absolute difference is 2.2 degree days above 29°C and the root mean squared prediction error is 3.1 degree days above 29°C, suggesting that the overlapping fit is reasonably close. Our weather data uses the average of all overlapping series.

Figure A1: Fraction of Corn Area Planted that is Harvested



Notes: Left panel shows the ratio of the corn area harvested to the area planted in Indiana 1926-2005 as black diamonds as well as a locally weighted regression with a bandwidth of one decade as grey solid line. The right panel shows yields under the two different definitions. Production divided by area harvested is shown as black diamonds, and production divided by area planted as grey triangle.

Figure A2: Interpolation Accuracy (1901-2005)



Notes: Graph shows degree days above 29°C in Indiana for each overlapping 25-year interpolation period starting in 1901-1925, 1910-1935, ..., until 1980-2005.