COLUMBIA | SIPA
Center for Environmental Economics and Policy

CEEP Working Paper Series
Working Paper Number 7

August 2019

Ground-Level Ozone and Corn Yields in the United States

Christopher D. A. Boone, Wolfram Schlenker, Juha Siikamäki

https://ceep.columbia.edu/sites/default/files/content/papers/n7.pdf

Ground-Level Ozone and Corn Yields in the United States*

Christopher D. A. Boone[†] Wolfram Schlenker[‡] Juha Siikamäki[§]

August 15, 2019

Abstract

We provide new empirical evidence of a nonlinear effect of ozone on corn yields. Our county-level panel analysis links observed historical yields to measures of air pollution constructed from fine-scaled hourly pollution monitor data. We find a critical threshold of 65 ppb for hourly daytime ozone concentration, considerably higher than the 40 ppb threshold derived from controlled experiments and used as the basis for air quality standards in Europe. Exposure to concentrations above this threshold has large negative impacts on yields. Our estimates indicate that a substantial fraction (between 32% and 44%) of the growth in U.S. corn yields observed between 1990 and 2014 can be attributed to reductions in peak ozone levels. A back-of-the envelope calculation reveals that the decline in peak ozone levels in the U.S. has reduced global food prices of the four basic staple crops by roughly 10% and increased consumer surplus of these commodities by $38 billion annually.

*We are grateful for helpful comments from participants at the 2012 Northeast Workshop on Energy Policy and Environmental Economics at Cornell and the AERE sessions of the 2013 ASSA meetings in San Diego, including especially Jordan Matsudaira and Barrett Kirwan, who served as discussants.

[†] Cornell University. Email: cb238@cornell.edu

[‡] Columbia University and NBER. Email: wolfram.schlenker@columbia.edu

[§] International Union for Conservation of Nature. Email: juha.siikamaki@iucn.org

1 Introduction

Corn yields in the United Sates have been trending upward at a rapid pace since 1950, experiencing a remarkable productivity gain that outpaced most other sectors of the economy. Some commonly cited reasons for this gain include the expanded use of fertilizer, irrigation, and pesticides, as well as the introduction of new crop varieties like hybrid corn (Jorgenson and Gollop, 1992; Alston et al., 1998). The United States currently produces roughly 40% of the world's corn, which is responsible for a substantial portion of global caloric consumption. Four commodities (corn, wheat, rice and soybeans) account for 75% of the calories that humans consume (Cassman, 1999).[1] Among those four commodities, U.S. corn production on average accounted for 11% of the caloric production in 1990-2014. Given its market share, any effect on U.S. corn production influences global commodity markets of all four commodities, where prices of these substitutes move closely together.

There was a general downward trend in real agricultural commodity prices over the 20th century as production increases outpaced increases in demand. This trend reversed as prices increased during the first decade of the 21st century, but since 2012 prices have resumed a downward trend. The demand for calories has increased as emerging economies switch to a meat-based diet that uses corn as feedstock. At the same time, corn-based biofuel mandates have put additional pressure on food supplies (Hill et al., 2006). The combined increase in calorie demand can only be met if yields continue to grow rapidly. Understanding the driving forces behind the strong past growth in average yields is an important step in modeling the future of commodity prices. While climate change has the potential to significantly alter yields in the future (Schlenker and Roberts, 2009), current empirical estimates are usually calculated relative to existing trends, which are taken as exogenous. A better understanding of the causes of yield growth is crucial to modeling future crop prices and food security.

In this paper we focus on the contribution of air pollution reduction to agricultural yield growth in the United States using a county-level panel of corn yields for the last 25 years (1990-2014). The counties in our sample account for 91% of U.S. corn production. We combine the data on yields with a variety of measures of local ozone exposure, which we construct using hourly readings from all ozone monitors maintained by the EPA. We focus in particular on identifying nonlinearities in the relationship between ozone and yields.

We make three contributions to the literature. First, we find strong and robust evidence that the effect of ozone exposure is large in magnitude and may result in substantial produc-

[1] On a calorie-weighted basis, the U.S. share of total production of corn, rice, soybeans and wheat was 23%, and the majority of that share comes from corn (Roberts and Schlenker, 2013).

tion losses. Our flexible nonlinear specification allows us to identify a critical threshold of 65 parts per billion (ppb): below this threshold, there is no statistically significant effect of ozone exposure on yields; above this threshold, however, there is a negative and significant effect on yields, and the magnitude of the effect is linearly increasing in ozone concentrations above 65 ppb. We compare our measure to others that are used as the basis for pollution standards. Our measure performs better at explaining corn yields than those used as the basis for current air quality standards in the U.S. and in Europe.[2] It also outperforms the W126 measure which was recently proposed by the U.S. Environmental Protection Agency (EPA) as the basis of a revised air quality standard.

Second, in order to get a better sense of the magnitude of the effects, we use our empirical estimates to investigate the impact of peak ozone concentrations on crop production. We compare production in our sample counties under observed ozone levels to a counterfactual where ozone concentrations are truncated at 65 ppb, i.e., all hourly ozone readings above 65 ppb are set to 65 ppb. The estimated crop losses due to ozone pollution above this threshold are as much as 20% of potential production in the early years of our sample. The estimated losses in recent years are smaller, resulting from the decline in exposure to ozone concentrations above 65 ppb. Peak ozone levels have been trending downward in recent years, and our estimates of production shortfalls due to ozone therefore decline from almost 20% at the beginning of our sample period to almost zero recently. At the same time, corn yields have been growing at a rate of 1.5% per year, for a total increase of 43% between 1990 and 2014. Our estimates imply that the reduction in peak ozone levels is responsible for between 32% and 44% of the observed growth in corn yields over this period. Forty-four percent of the observed yield growth represents a 19% overall increase in yields and corresponds to the caloric equivalent of feeding 170 million people on a 2000 calories/day diet.

Third, we derive the effect of the reduction of peak ozone levels on food prices. Elimination of peak ozone in the United States has decreased global food prices of the four basic staple crops by roughly 10% and increased consumer surplus of these commodities by 38 billion USD annually. While consumers of basic calories gain, producers suffer. U.S. corn producers have seen increased yields, yet the reduction in prices offset these gains. Producers of the four basic commodities outside the U.S. saw a reduction of surplus through lower prices.

Understanding the relationship between ozone exposure and crop yields is important

[2] The secondary standard for U.S. ozone pollution is based on the daily maximum 8-hour average, and the secondary standard in Europe is based on linear exposure above a lower threshold of 40 ppb that has been identified in laboratory studies.

for the design of regulatory standards. Regulatory agencies generally do not place limits on instantaneous concentrations but instead impose limits on some aggregate statistic calculated using the entire set of observed concentrations. The United States sets two regulatory standards for air pollutants: a primary standard designed to protect human health, and a secondary standard designed to protect "human welfare," which includes agricultural yields.[3] For ozone, the current secondary standard sets a limit based on the average concentration observed over a consecutive 8-hour period.[4] As part of a recent review process, the EPA proposed to revise the secondary standard and instead base it on a measure of cumulative monthly ozone exposure, where cumulative monthly exposure is calculated as a nonlinear weighted sum of hourly ozone concentrations.[5] In Europe, the ozone standard is based on the total cumulative exposure above a threshold concentration of 40 ppb. Our results allow us to compare the performance across alternative measures of ozone exposure in explaining the variation in crop yields.

Much of the current understanding about the effects of ozone on crop yields comes from the results of controlled experiments using growth chambers such as glass houses or open-top chambers (Heagle, 1989). In these experiments, crops are grown in fully or partially enclosed environments, where the ambient ozone concentration is controlled by adding ozone or filtering the air. More recently, experiments have been run in open-air environments using free-air gas concentration enrichment technology (Morgan et al., 2006). Experiments usually differentiate between acute damage (from short-term exposure) as well as chronic damage, that is caused by continuous longer-term exposure to lower levels. We focus on acute damage from peak hourly readings.

Although such experiments offer powerful means to test for the effects of ozone in controlled settings, the differences between the controlled settings and actual agricultural production environments, including the full range of relevant environmental conditions (e.g., weather, insects, disease) and farmers' production decisions, may considerably limit the applicability of the currently available evidence to project damages from ozone pollution in real-world agricultural production. This has led to calls for comparable assessments using actual crop production and environmental data. For example, the extensive assessment of

[3]The two standards are currently the same for most pollutants.

[4]More specifically, the limit is imposed on the 4th highest daily maximum 8-hour average, averaged over three consecutive years.

[5]The agency ultimately decided to continue using the same index, based on consecutive 8-hour averages, but to make the standard more stringent. Information on the 2015 National Ambient Air Quality Standards (NAAQS) for Ozone can be found at https://www.epa.gov/ground-level-ozone-pollution/2015-national-ambient-air-quality-standards-naaqs-ozone.

ozone pollution by the Royal Society highlights the need for empirical studies to verify the effects of ozone pollution, including potential threshold effects, using observational data under real-world growing conditions (Royal Society, 2008). Similar concerns were raised by the United States Environmental Protection Agency in the recent regulatory documents concerning the National Ambient Air Quality Standard for ozone (Environmental Protection Agency, 2008).

Previous research employing data from real-world agriculture has studied the effect of ozone on corn and soybean yields in a cross-section of several hundred fields in the Eastern U.S. (Westenbarger and Frisvold, 1995). However, cross-sectional studies are potentially subject to and limited by omitted variable bias. Another study focused on soybean yields using five years of data from three states in the Midwestern U.S. (Fishman et al., 2010). The most closely related paper to ours is by McGrath et al. (2015), who also find significant negative impacts of ozone pollution on crop yields using panel data.[6] Our analysis focuses on corn, the key agricultural crop in the U.S., and encompasses almost all corn production in the United States over a time period of 25 years. Following the results of earlier studies, we allow for the existence of thresholds and nonlinear relationships. We investigate these nonlinearities in more detail than has been done in previous observational studies or has been feasible in experimental settings where there are practical limitations on how much the treatment can be varied.

One empirical challenge that we face involves transforming our pollution data so that it corresponds to the same spatial and temporal scale as our agricultural data. Simply averaging across space or time makes it harder to identify possible nonlinear relationships and can lead to significant attenuation bias, and we therefore take great care to capture the actual ozone exposure. To do so, we use hourly ozone readings to construct several measures of ozone exposure before aggregating them over the entire growing season (March to August). We also develop a new methodology to approximate missing values using the cumulative distribution function of each monitor station. We then fit a pollution surface between monitors and average the values over the agricultural area within each county to

[6] That paper was written concurrently with ours and uses the same underlying data sources for crop production and pollution. Our paper differs in a number of key respects. By accounting for nonlinearities when aggregating across space and time and by capturing the effects of ozone exposure at different concentrations, we show that the threshold estimated using actual field-based yields is substantially higher than the threshold estimated from earlier experiments; we use these estimates to evaluate the performance of different ozone measures; and we show that failure to properly account for these nonlinearities substantially biases the estimates of pollution-related crop losses as well as projections for future yields. We also employ nonlinear transformations of temperature that have been shown to influence yields, and we discuss the effects of ozone on yields trends and food prices.

obtain a more accurate measure of the actual exposure. The five sets of ozone variables that we construct are: (i) the maximum daily 8-hour average, which is the basis for the current U.S. secondary (and primary) ozone standard; (ii) the simple daytime mean exposure; (iii) the weighted sum of all hourly daytime exposure using EPA's W126 weight, the basis for the proposed revised secondary ozone standard; (iv) a measure of cumulative linear exposure above a threshold, which has been used in laboratory experiments that form the basis for some regulatory standards in Europe;[7] and (v) a semi-parametric approach that models the relationship with the help of flexible restricted cubic splines.

In our baseline specifications we control for weather since ozone formation is correlated with temperature, which itself affects yields (Schlenker and Roberts, 2009). We also include two sets of time controls: county-specific quadratic time trends to capture overall yield trends that have varied in space (Burke and Emerick, 2016), as well as year fixed effects to capture system-wide effects like movements in global commodity prices. To address concerns about omitted variables bias, we show that our results are robust to a variety of controls and across a number of specifications. The inclusion of county fixed effects as well as county-specific time trends ensures that the parameters are identified by fluctuations around baselines that are allowed to trend separately in each county. Threats to the validity of our empirical estimates must entail some unobserved factor that varies systematically with ozone around a county-specific mean or trend. The most likely candidate would be weather. We demonstrate that our results are insensitive to the inclusion of a variety of different weather-related controls. The other obvious candidates are other (non-ozone) air pollutants, and we show that controlling for other pollutants leaves the estimated effect of ozone unchanged.

2 Data

We begin by describing the data on maize yields and ozone pollution that we use for our analysis. This section outlines in detail how we aggregate the hourly ozone readings at the monitor level into daily ozone measures at the county level (which are then aggregated over the growing season).

[7] While the European standards usually rely on a 40 ppb threshold, we vary it between 1 ppb and 120 ppb.

2.1 Agricultural Yields

Data on agricultural yields were obtained from the National Agricultural Statistics Service Quick Stats website for the years 1990-2014.[8] Annual county-level corn yields are defined as the ratio of total production in a county to the harvested area. We focus on dryland agriculture and therefore only use counties east of the 100 degree meridian (except Florida) that report corn yields in at least half of the years in our sample (13 out of the 25). These counties are shaded grey in Figure 1 and account for on average 91% of annual U.S. production in 1990-2014. The aggregate log yield in our sample is shown in Figure 2.

2.2 Pollution Data

Since yield data are reported on an annual level for each county, we need to aggregate the hourly ozone data to an annual county-level panel as well. Earlier studies linking corn yields to weather outcomes used a growing season of March 1st-August 31st of each year (Schlenker and Roberts, 2009). Our model uses the same definition of the growing season to derive a season-total ozone measure.

We start by downloading hourly monitor-level ozone readings for all monitors in the United States for the years 1990-2014.[9] Some ozone monitors only operate for a fraction of a season. Figure 3 shows the average number of monitors that report ozone readings on a day across the days of the year. There is a significant increase in the number of active monitors on April 1st. We therefore separately interpolate ozone data for the month of March (when fewer monitors are available) and the time period from April through August.

For each of the two time periods (March and April-August) we follow the same script. First, the data set is restricted to monitors that report values for at least 75% of the days in the time period of consideration in a year. We do this separately for each year as few monitors report consistently for the entire time period 1990-2014. It is possible that a monitor is included in only some years. The number of "good" monitors in our data set, i.e., the ones that have at least 75% of their daily values non-missing for at least one of the time periods, are shown in Figure 4.

Second, for each monitor, we aggregate the hourly readings into daily statistics, such as average daily concentration or total exposure above a threshold concentration. These statistics are defined in Table 1 below.[10]

[8] We downloaded data from the "Survey" program at http://quickstats.nass.usda.gov.

[9] http://www.epa.gov/ttn/airs/airsaqs/detaildata/downloadaqsdata.htm

[10] Some hourly observations are missing. Since hourly ozone readings are highly serially correlated, missing

Third, we fill in the missing daily values from the previous step by using the average cumulative distribution function (CDF) of surrounding non-missing stations. Specifically, we compute the empirical CDF of daily readings for each monitor; then, for each non-missing monitor on a particular day, we compute the corresponding percentile of that day's reading based on the monitor's CDF; and finally, for any monitor with a missing value on that day, we compute the weighted average of the percentiles from all the surrounding non-missing monitors, where the weights are the inverse squared distance, and we fill in the missing value using the value that corresponds to the weighted-average percentile from that (missing) monitor's own CDF.[11] This procedure produces a balanced panel of daily monitor-level ozone characteristics for each year and time period.

Fourth, we interpolate the daily monitor readings to the fine-scaled grid used by the Parameter-elevation Regressions on Independent Slopes Model (PRISM),[12] the grid underlying our daily weather data. It is an evenly spaced grid in both longitude and latitude with a grid size of $\frac{1}{24}$ degrees. Pollution readings at a grid point are the squared inverse distance weighted average of all monitor readings.[13]

Finally, we then average the daily ozone readings for all grid cells in the county in order to construct a county-average daily ozone reading, weighting by the cropland area in each grid cell, which was obtained from the 1992 Land Cover database provided to us by Shawn Buchholz at the Economic Research Service at USDA.

2.3 Weather Data

Ozone is highly collinear with solar radiation and temperature. It forms from Volatile Organic Compounds (VOC) and Nitrogen Oxides (NO_x). Omitting temperature controls from

observations are better approximated by the previous reading than the daily average. We thus weight each hourly observation by the time until the next reading. For example, if readings occur at 9 am and 11 am, the 9 am reading will get a weight of 2 (hours). By the same token, if the first readings of the day are missing, we count the hours until the first reported reading and assign it to the first observed reading. For example, if we calculate the daily average between 6 am and 8 pm, and readings occur at 8 am, 9 am, and 10 am, the 8 am reading will receive a weight of 3 (hours). Similarly, if the last readings of the day are missing, we assign them to the last observed reading. As a result, the weights of all readings on a day always sum to the number of hours of the day over which the ozone measure is calculated.

[11] For example, if a non-missing monitor 1 km away has a value that equals 80th percentile of all readings in 1990-2014 for the given time period (e.g., March or April-August), and another that is twice as far away has a value that equals the 60th percentile, the weighted average will be the 76th percentile. Supposing that the weighted average of all non-missing monitors corresponds to the 76th percentile, we then replace the missing value at the station with the value at the 76th percentile of *its own* cumulative distribution function.

[12] http://www.prism.oregonstate.edu

[13] We always add 1 m to the distance to avoid division by zero in case a monitor is located at a grid centroid.

the regression will bias the ozone coefficient as it would attribute the effects of temperature to ozone. We use the same fine-scaled daily weather data set of Schlenker and Roberts (2009) that gives daily minimum and maximum temperatures on a $\frac{1}{24}$-degree grid (2.5 minutes, or about 4.5 km by 4.5 km) in both longitude and latitude, but extend it to 2014. This allows us to derive degree days measures that capture the nonlinear effects of temperatures on yields.

3 Model

We estimate a panel regression linking log yields to pollution exposure over the growing season, which we fix to be March-August. Specifically, log corn yields in county i and year t are regressed on measures of ozone as well as other explanatory variables.

$$log(y_{it}) = \alpha o_{it} + \mathbf{W}_{it}\boldsymbol{\beta_1} + \mathbf{X}_{it}\boldsymbol{\beta_2} + f(t) + \gamma_i + \epsilon_{it} \tag{1}$$

The main coefficient of interest α captures the effect of the ozone measure o_{it} over the growing season in year t in county i on log corn yields. The various ozone measures are discussed below in Section 3.1. We control for weather variables $\mathbf{W}_{it}$ to capture the direct effect of weather on yields. Since ozone formation is correlated with warmer temperatures, controlling for weather variables is crucial to avoid omitted variable bias. If we omit weather variables, the coefficient α becomes more negative as it also captures the direct damaging effects of extreme heat, which is conducive to ozone formation. In our baseline specification the weather controls include total degree days between 10-29°C, degree days above 29°C, and a quadratic in precipitation. Similarly, $\mathbf{X}_{it}$ can include other control variables that could covary with ozone and directly influence yields, such as other pollutants.

Yields exhibit strong upward trends over time, and we hence include temporal controls $f(t)$. The baseline specification uses county-specific quadratic time trends to capture overall movements in average yields over time as well as year fixed effects that capture economy-wide shocks (e.g., fluctuations in world prices).

Finally, the inclusion of a county fixed effect γ_i ensures that our identification stems from comparing outcomes within a county across years. The error terms ϵ_{it} are clustered at the state level to account for spatial correlation. This procedure gives comparable standard errors to approaches that more specifically model the covariance of error terms between counties as a function of distance (Fisher et al., 2012), e.g., Conley's standard errors (Conley, 1999).

3.1 Ozone Measures

We construct five sets of ozone variables: (i) maximum daily 8-hour average, the basis for the current secondary ozone standard; (ii) simple daytime mean exposure; (iii) weighted sum of all hourly daytime exposure using EPA's W126 weight (Figure 5), the basis for the proposed revised secondary ozone standard; (iv) linear hourly exposure above a threshold, the variable used in aforementioned laboratory experiments that form the basis for some regulatory standards in Europe; while the European standards usually rely on a 40 ppb threshold, we vary it between 1 ppb and 120 ppb; and (v) a semi-parametric approach that models the relationship with the help of flexible restricted cubic splines with 7 knots. The exact specifications are given in Table 1. For all other pollutants, we construct the daily daytime average. Our interpolation procedure is further assessed and discussed in Section 5.6 through cross validation.

3.2 Contribution of Peak Ozone Levels on Yield Trends

We investigate the contribution of declining pollution levels to the observed growth in corn yields. We derive the contribution of peak ozone on corn yields more formally using a bootstrap technique that accounts for the uncertainty of our parameter estimates. We compare yields under observed peak ozone levels to a counterfactual where concentrations above 65 ppb are truncated at 65 ppb—that is, ozone exposure above 65 ppb is set to zero. The exact steps of the procedure are described in Table 2.

4 Empirical Results

4.1 Cross-Section

Earlier studies have linked field-level yields to ozone exposure in the cross-section (Westenbarger and Frisvold, 1995). For comparison, we use our data set of piecewise linear ozone exposures to replicate a series of cross-sectional estimates that are summarized in Figure 6. We use the same county-level data set and variables (except temporal controls), but limit it to one year at a time. The x-axis indicates the year that is used in the estimation.

The estimates vary significantly between years. More strikingly though, the top line of the graph reports the threshold that gives the highest R^2.[14] The optimal threshold varies

[14] We pick the threshold among the following possible candidates: 0, 10, 20, 30, 40, 50, 55, 60, 61, 62, ..., 89, 90, 95, 100, 110, 120.

anywhere between 10 ppb and 120 ppb; that is, it largely depends on what year we use.

4.2 Panel Regression Results

We find strong and robust evidence that the effect of ozone exposure is large in magnitude and may result in substantial production losses. The effect of ozone exposure on maize yields is approximately linearly increasing in ozone concentration above a threshold of 65 ppb. Figure 7 shows the regression results for two of our models that link log yields to ozone: the lines represent the effect on log yields from 100 hours of exposure to various concentrations. The blue line shows the results using restricted cubic splines with 7 knots (the knot locations are indicated by the dashed lines). Exposure to ozone concentrations below 65 ppb have no statistically significant effect on maize yields. On the other hand, exposure to ozone concentrations above the 65 ppb threshold have a negative and significant effect that is large in magnitude and increases linearly in concentrations above 65 ppb. One hundred hours of exposure to 90 ppb reduces annual corn yields by 18 log points (or about 19%) compared to exposure below 65 ppb.

For comparison, we also add a second model in Figure 7 that forces the effect of ozone to be linear in exposure levels above 65 ppb. It is shown as the red line. The slope is comparable to the model that uses restricted cubic splines. Since the model is estimated using fixed effects, the results are in relative terms (net of the constant, or group fixed effect). We therefore normalize both models so the y-value is zero on average below 65 ppb for easier comparison of the relative effects.

Regression results for these two models as well as the other specifications outlined in Table 1 are summarized in Table 3. Most of the variation in ozone is already absorbed by the county fixed effects. A model with only county fixed effects has a R-squared of 0.5. We hence report the R-squared once the fixed effects and temporal controls (county-specific time trends and year fixed effects) are partialled out. The R-squared is increasing from left to right (except for column 4a). Since we are keeping all other control variables the same between the columns, a higher R-squared implies greater explanatory power of the ozone measures. Controlling for weather variables has a R-squared of 0.257. If we finally include the linear ozone exposure measure above 65 ppb (column 4b), the R-squared increases modestly to 0.286. The reason is that much of the ozone variation is already accounted for: in space (dirty versus clean areas) through county fixed effects and over time through the year fixed effects and county-specific time trends. We prefer to control for the covariates to rule out confounding variation or spurious correlations between two trending variables.

Just for comparison, a model that only includes the linear ozone measure above 65 ppb but no other controls has a R-squared of 0.18 and the coefficient is -1.073, which is larger in magnitude than our preferred estimate of -0.737 in column (4b). Our preferred estimates are more conservative.

Figure 8 shows the R-squared for additional linear models above an exposure threshold. The piecewise linear model with the highest R-squared uses a threshold of 66 ppb. In-sample R-squared, however, has been criticized as a model selection criterion. For example, it can only go up as the number of control variables increases and is not the best measure to select between models.

We therefore also examine out-of-sample forecasts to help assess which concentration-response specification best predicts variation in our data (Efron and Tibshirani, 1993). It is an intuitive way to assess which model is best: Ultimately, we are interested in the model that can best predict yields. Assessing this criterion out-of-sample avoids overfitting the data in-sample.[15] We hence assess each model in the following way. We randomly draw 80% of the data, estimate the model, and predict log yields for the remaining 20% out-of-sample.[16] We calculate the squared prediction error for each of the observations in the 20% prediction sample, and derive the root mean squared error. We repeat this procedure 1000 times so the results are not driven by the particular selection of the 20% sample. Figure 9 plots the average percent reduction in root-mean-squared error compared to a model with no ozone variable.

The model that uses the 4th highest daily maximum 8-hr average, the basis of the current U.S. ambient air quality standard, has almost no effect on the RMSE, compared to a model without any ozone variable. This is perhaps not surprising as it only counts one day of the season and omits the rest. The highest reduction in prediction error is accomplished by the flexible spline in hourly ozone exposure over the growing season. Recall that this model with seven knots has six variables. It reduces the RMSE twice as much as a model that uses the average of daytime (6am-8pm) ozone readings. It is also better than the newly proposed W126 weighted sum of hourly ozone exposure.

The figure also shows that within the class of linear exposure models, a model with a threshold of 64 ppb has the best performance; that is, when the threshold is systematically varied between 60 ppb and 90 ppb, the reduction in RMSE peaks at 64 ppb. (This peak is also fairly close to the 66 ppb threshold that gave the highest in-sample R-squared in

[15] It is usually possible to get an almost perfect fit by adding enough covariates.

[16] If observations from a county only appear in the 20% prediction sample, we include a fixed effect to make sure the average error is zero, for all other case we use the fixed effect from the 80% estimation sample.

a linear model in Figure 8.) Moreover, the linear exposure model performs only slightly worse than the spline model.[17] While the spline model is very flexible and has the greatest predictive power, the linear exposure model is parsimonious and easier to interpret and use for forecasts, as it only requires the estimation of one coefficient. In the remainder we use a threshold of 65 ppb, which is in the middle of 64 ppb (best out-of-sample fit) and 66 ppb (highest R-squared).

4.3 Production Impacts of Ozone

In order to get a better sense of the magnitude of the effects of ozone on crop production, we use our empirical estimates to investigate the impact of extreme ozone concentrations (corresponding to concentrations above the threshold of 65 ppb identified above). We derive the estimated crop losses by comparing the sum of predicted production in our sample counties under observed ozone levels to a counterfactual where ozone concentrations are truncated at 65 ppb, i.e., all hourly ozone readings above 65 ppb are set to 65 ppb.[18]

Figure 10 displays the reduction in maize production from hourly ozone concentrations above 65 ppb under our two preferred model specifications: a flexible spline in hourly ozone exposure as well as a more parsimonious model of linear ozone exposure above 65 ppb. The figure shows that the estimated crop losses due to ozone pollution in the early years of our sample are as much as 20% of potential production. The spline model gives smaller estimated declines of up to 14% of potential production.[19] A 20% loss corresponds to the caloric equivalent of feeding 170 million people on a 2000 calories/day diet. The estimated losses in recent years are smaller, as exposure above 65 ppb has decreased over time; as shown in Figure 11, while average ozone concentrations show no discernible trend, there is a strong decline over time in exposure to extreme hourly concentrations. The large majority

[17] The W126 measure does not perform quite as well as the linear exposure model, but still performs better than the models that are based on average concentrations. The implication is that the W126 weighting function (shown in Figure 5) does not capture the impacts of ozone exposure as well as a simple linear model with a threshold at 65 ppb, but it is not too far off. Note that the three measures based on cumulative exposure perform better than all three measures based on average concentrations.

[18] Predicted production under observed ozone levels corresponds to $\widehat{log(y_{it})}$ from Equation (1); the counterfactual production is then calculated as $\widehat{log(y_{it})} - \hat{\alpha} o_{it}$ where o_{it} equals the cumulative exposure above 65 ppb.

[19] The estimates of crop losses using the spline model are smaller than for the linear exposure model but still quite large in magnitude. The gap between the two estimates highlights the uncertain nature of these predictions, but both models suggest the losses are substantial. Note that we selected the optimal threshold for the linear exposure model based on the threshold that resulted in the highest explanatory power (65 ppb), but the locations for the knots in the restricted cubic splines were chosen ex ante.

of the estimated yield reductions due to ozone are associated with the 65-90 ppb range. For example, when we set all readings above 90 ppb to 90 ppb, the estimated reduction is three quarters of the number we get in the baseline case.

The top left panel of Figure 12 replicates the predicted production reductions from hourly ozone concentrations above 65 ppb for all regression models of Table 3. The numbers in the legend refer to the column number of the regression model. Besides our two preferred specifications and the W126 measure, the production impacts are much smaller under the other model specifications. The other ozone measures pool ozone readings that matter for yields together with ozone readings that do not, resulting in a coefficient that is biased towards zero and an underestimate of the true effect of ozone on yields. The linear model of cumulative exposure above a threshold of 40 ppb, for example, puts too much weight on concentrations between 40 and 65 ppb, and as a result underestimates the negative effects of concentrations above 65 ppb. The same is true of the W126 measure, but the bias in not nearly as extreme.[20]

Traditional chamber studies rely on a linear ozone exposure measure above 40 ppb, suggesting that reducing any hourly ozone reading above 40 ppb will be beneficial. The top right panel of Figure 12 displays the predicted production impacts from hourly ozone levels above 65 ppb (grey dashed line) as well as 40 ppb (black solid line) for a regression model that is based on a linear exposure measure above 40 ppb. The predicted impacts of ozone pollution above 40 ppb are much larger than above 65 ppb as the same regression coefficient is multiplied by a larger cumulative exposure measure. Note, however, that this result relies on significant out-of-sample interpolation that might be questionable. The bottom left panel of Figure 12 shows the histogram of observed annual linear ozone exposure above 40 ppb in our data, which is never close to zero. Constructing a counterfactual where the variable is zero presumes that the estimated linear model will extend beyond the observed range of values in the data. For comparison, the bottom right panel shows the histogram of observed annual linear ozone exposure above 65 ppb, which has significant mass around zero. Our baseline results do not rely on out-of-sample predictions.

[20] Essentially, the W126 uses too low of a threshold, overweighting concentrations below 65 ppb. See Figure 5 for the shape of the W126 weighting function, which looks like a close approximation to a linear model of cumulative exposure above a threshold of about 50 ppb. As expected, Figure 9 shows that the W126 measure results in the same out-of-sample performance as the linear exposure model with a threshold of 50 ppb.

4.4 Contribution of Peak Ozone Levels on Yield Trends

Given the significant impacts of ozone on corn yields, we derive the fraction of the yield trend that is attributable to changes in peak ozone. Figure 12 shows that peak ozone levels as measured by our piecewise linear exposure measure above 65 ppb reduced aggregate corn production by up to 20% in the earlier years of our sample but had close to no impact by the end of the sample. At the same time, corn yields have been growing at a rate of 1.5% per year, for a total increase of 43% between 1990 and 2014 (Figure 2).

The distribution of the fraction of the observed yield trend that is due to reduction in peak ozone levels is shown in Figure 13 for our model using a linear exposure measure above 65 ppb (column 4b in Table 3). The mean is 44% with a standard deviation of 5.6% with a 95% confidence interval that stretches from 33% to 55%. That is, the model predicts that the reduction in ozone exposure above 65 ppb contributed to a 19% increase in yields (44% of the observed 43% growth). We conduct three sensitivity checks that all give comparable results. First, if we use historically observed yields in step 3c) instead of predicted yields, the predicted fraction that is explained by reduction in peak ozone remains at 44% with a 95% confidence interval of [33,55]. Second, if we use the average area in a county instead of the annual observed area in steps 3b) and 4b) to rule out shifts in planting areas, the predicted fraction has a mean of 42% with a 95% confidence interval of [31,52]. Finally, if we use the spline model in column (5) of Table 3, the predicted fraction has a mean of 32% with a 95% confidence interval of [14,50].

The estimated increase in yields due to ozone reduction is not geographically uniform. Figure 14 fits a separate trend in ozone exposure above 65 ppb for each county, which are all negative, and then multiplies the linear trends by the estimated coefficient from the panel regression (column 4b in Table 3) to obtain the predicted trend in yields due to ozone reductions. Counties on the Eastern seaboard exhibited the largest yield gains that are due to ozone reductions, up to 2.8% per year. Not surprisingly, these areas that had the highest pollution levels to begin with in 1990 (Appendix Figure A1). The figure also shows large contributions to yield growth in parts of Illinois and Indiana, which contain counties with high baseline yields.

5 Robustness Checks

We present several sensitivity checks to ensure that our results are driven by changes in ozone and not some other confounding variation. The robustness checks evaluate the main results both under alternative statistical modeling approaches and under a variety of controls for potentially confounding factors. We focus here on the parsimonious model examining daytime hourly exposure above 65 ppb, summed over the growing season, as it can be summarized by one parameter. For ease of comparison, the first column in each of the tables repeats our baseline results using a linear exposure model above 65 ppb, while additional columns present sensitivity checks. We also investigate how the optimal threshold changes in response to our sensitivity checks. For each specification we estimate the model using thresholds of 0, 10, 20, 30, 40, 50, 55, 60, 61, 62, ..., 88, 89, 90, 95, 100, 110, and 120 ppb. We report the coefficient estimates for the model containing linear exposure above 65 ppb to facilitate comparison across specifications, but report the threshold that results in the highest R-squared value in the footer of the table, i.e., the best in-sample measure. The optimal threshold based on the in-sample criterion is 66 ppb in our baseline, and our 46 sensitivity checks have a mean optimal threshold of 65 ppb with a standard deviation of 6.3 ppb, i.e., they remain close to our baseline.

5.1 Including Additional Weather Controls

Higher temperatures are conducive to ozone formation. One might worry that higher ozone levels simply approximate higher temperatures, which are themselves harmful for corn yields. Recall that degree days above 29°C are the best predictor of year-to-year variation in yields. Table 4 examines the sensitivity of our baseline results to the chosen temperature controls by including additional temperature controls: minimum, maximum, and average temperature, as well as the diurnal temperature range (maximum minus minimum temperature). The last four rows display what variables are included and up to which order. The first column replicates our baseline results, while consecutive columns add additional controls. The last column includes an additional 20 control variables: all four temperature measures up to order 5.

The first row of Table 4 gives the main coefficient of interest: hourly ozone exposure above 65 ppb, summed over the growing season. When we include the additional temperature variables, the coefficient on degree days above 29°C varies a lot (from -0.498 to -2.649) due to the correlation with the other temperatures variables; nonetheless, the coefficient on ozone exposure remains very robust (between -0.637 and -0.761), which makes it unlikely that the ozone variable is picking up temperature effects.

Table 5 includes higher order precipitation terms, which have even smaller effects on the ozone coefficient than temperature controls: coefficients vary between -0.731 and -0.761.

5.2 Interaction with Maximum Temperature and PM_{10}

In a second step we not only include additional temperature variables, but also interact them with our preferred ozone measure.[21] Since both hot temperatures as well as solar radiation increase ozone formation, we also include PM_{10}, which is correlated with haze and hence less sunlight. We standardize the annual interaction term by removing the mean and normalizing by the standard deviation of the demeaned series. The coefficient on the interaction variable gives the effect of a one standard deviation increase in the interaction variable as we move away from the mean.

The first row of Table 6 now gives the effect of ozone exposure above 65 ppb if the interaction variables are kept at their mean level. This average effect is attenuated a bit when interacting with both degree days above 29°C and maximum temperature, but the effect is relatively robust overall, varying between -0.579 and -0.784. The interaction with degree days above 29°C or maximum temperature is statistically significant, suggesting that ozone is more harmful if the crop is experiencing heat stress. On the other hand, ozone is less harmful if PM_{10} is higher, which implies increase in haze and hence less solar radiation.

5.3 Controlling for Other Pollutants

Pollution concentrations are highly correlated as many of them are by-products of the same industrial activities. To rule out the possibility that other pollutants besides ozone are causing the large decline in yields, Table 7 controls for average pollution levels of carbon monoxide (CO), nitrogen oxides (NO_x), particulate matter (PM_{10}), and sulfur dioxide (SO_2) one at a time as well as all of them together.[22]

The first row of Table 7 gives the estimated coefficient on ozone exposure above 65 ppb, which is very robust to accounting for other pollutants in the regression equation as the coefficient hardly moves at all between -0.733 and -0.742. None of the other pollutants have statistically significant effects. Some papers even found beneficial effects of SO_2 as crops are sulfur limited (Sanders and Barreca, 2012).

[21] We interact the season-total of the variables, not the daily values.

[22] We construct the daily daytime (6am-8pm) mean levels of the other pollutants the same way we constructed the average ozone variable as described above, i.e., we interpolate then to the PRISM grid and average them over the agricultural area.

5.4 Sensitivity to Temporal Controls

Our baseline results include two sets of temporal controls: first, we include county-specific quadratic time trends to capture smooth increases in average yields that are allowed to vary by county; on top of that, we include year fixed effects to allow for common shocks to the entire country, such as varying global price levels or breakthroughs in crop technologies.

Table 8 varies whether we include year fixed effects or not, and whether we include no time trends, linear county-specific time trends, or quadratic county-specific time trends. The coefficient on our ozone variable of interest in the first row varies somewhat, between -0.392 and -0.883 depending on the chosen time control, but not systematically, i.e., including year fixed effects sometimes increases and sometimes decreases the estimated coefficient.

Table 9 again varies whether we include year fixed effects and also varies the sensitivity of our results to whether we include common time trends, state-specific time trends or county-specific time trends. Switching from state-specific time trends to county-specific time trends has limited effect on the results, ruling out that our results are spuriously driven by common trends in pollution and yields among counties.[23] The more important factor is whether or not year fixed effects are included.

5.5 Subset of Counties and Measurement Error

Table 10 examines how the results vary depending on how the different counties are weighted. Our baseline results in column (1) use a pooled analysis where all observations receive the same weight. Column (2) use area-weights, where we weight by the average corn-growing area in a county.[24] The estimated coefficient decreases slightly from -0.737 to -0.717. Note that agricultural areas tend to be further away from monitors and the weighted regression might place more weight on counties with higher measurement error.

Columns (3a)-(3c) split the sample into counties based on their distance to the closest monitor.[25] Columns (3a) and (3b) estimate the model using two distinct subset of counties:

[23]For example, differences across counties in the adoption of yield-improving technologies such as fertilizer, irrigation, pesticides, or new seed technologies could lead to differential trends in yields. If these yield trends were for some reason correlated with trends in ozone levels, this could lead to bias in the estimate on the ozone variable. Controlling for county-specific time trends would substantially reduce this bias. It's therefore quite reassuring that the estimate on the ozone coefficient is robust across so many different specifications.

[24]The weights are hence constant over time and *not* subject to annual fluctuations in the growing area.

[25]For each of the PRISM grids we derive the distance to the closest monitor in a year. We then take the maximum of the minimum distance for all grids in a county to derive the largest minimum distance to a monitor in a county. Finally, we average the derived distance over all years for which we have corn yields in a county. Recall that some monitors only report for some of the years, and hence the distance to the closest

those with a distance below the median and those with a distance above the median, respectively. The coefficient for the areas further away is slightly smaller. Column (3c) therefore estimates the models for the two subsets of the data jointly by including an interaction term for the subset of counties whose distance is larger than the median distance.[26] It is positive but not significant, and the magnitude is of limited size. One possible explanation is that our pollution interpolation procedure performs worse for areas that are further away from monitors and we hence have attenuation bias.

Columns (4) go a step further and limit the analysis only to counties that have at least one reporting monitor. Column (4a) uses the same specification and spatial interpolation as column (1) but limits the data set to counties with a monitor. As a result, the number of counties in the sample decreases from 1747 to 474. Column (4b) uses the same counties but no spatial interpolation, i.e., it only averages all monitor readings in the county without any spatial interpolation. The coefficient decreases significantly in magnitude, which would be consistent with measurement error in the data set that simply averages all monitor readings.

To further asses the role of measurement error, we rely on instrumental variables regression. Column (5a) instruments the pollution variable in column (4a) with the simple monitor average in column (4b). Note how the coefficient hardly changes at all. If, on the other hand, we instrument the pollution variable in column (4b) with our spatially interpolated variable in (4a), the coefficient increases significantly in magnitude, which is again consistent with the fact that simply averaging all monitor readings has more measurement error than our spatial interpolation procedure. While our interpolation procedure might itself suffer from attenuation bias in counties that are far away from a monitor, it seems to do better than simply averaging monitor readings. In case there is remaining attenuation bias, our results will *understate* the true relationship, and are hence a conservative lower bound.

5.6 Cross-validation of Pollution Interpolation

Our pollution data relies on an interpolation procedure between monitors. This section presents checks on how well the interpolation procedure is working. We conduct cross-validation exercises where we omit one monitor at a time and interpolate the remaining data to the monitor location, which allows us to compare the interpolated values to the actual observed values.

monitor might change from year to year.

[26]It also interacts the weather variables and time controls with an indicator for whether the distance to the closest monitor is above the median, but these coefficients are omitted from the table due to space constraints.

Results of regressions where we regress the interpolated values at a monitor station on the actual reading are given in Table 11. The table has two panels: panel A only includes monitor fixed effects, while panel B also includes monitor-specific quadratic time trends and year fixed effects, analogous to the temporal controls we include in our baseline regression. Columns (1a)-(1c) use daily values in March-August, while column (2a) uses annual aggregates. Column (1a) uses all monitors and days in the Eastern United States, while column (1b) limits the data to monitors that lie within a county that has corn yields, and column (1c) additionally excludes days where the monitor reading is missing and had to be interpolated from adjacent stations.

The coefficients in all regressions are significantly less than 1, i.e., if the actual value deviates by a certain amount from the mean outcome at a monitor, the interpolated value deviates by less. The fraction is around two thirds in column (1a) of both panels where we use all daily values of all Eastern monitors. If we limit the data to monitors that lie within counties that have yield data in column (1b), the ratio remains the same. Excluding days where a monitor reading is missing in column (1c) also has close to no effect.

Finally, aggregating the data to season totals in column (2a) slightly increases the ratio in panel A when we only include monitor fixed effects, but decreases the ratio to 0.48 in panel B when we also include the temporal controls. The reason is that there are strong regional trends in season-total variation, and the temporal controls absorb these trends, amplifying measurement error.

If our interpolated values systematically under-predict the true variation in the data, this would bias the magnitude of our coefficients upward. On the other hand, if we add measurement error but the variation in the variable remains constant, attenuation bias will imply that our coefficient estimates are biased towards zero. The footers of each panel in Table 11 aims at disentangling the two influences. In panel A we regress both interpolated and observed monitor readings on monitor fixed effects and derive the standard deviation of the resulting error terms. Panel B also includes temporal controls in the regression before we obtain the residuals. Our interpolated variables have less variation than the observed readings and this effect if more pronounced if we include temporal controls in Panel B of column (2a), the variation we are using. On the other hand, the estimated coefficient on the interpolated ozone variable is higher, not lower, when we omit the temporal controls in column (3b) of Table 8. This makes it unlikely that we are overestimating the true coefficient. On the other hand, attenuation bias is a real concern in columns (4) and (5) of Table 10. Taken together, we feel that our interpolation procedure reduces some of the possible attenuation bias, but might still be a conservative lower bound on the true magnitude of the parameter.

6 Economic Effects

Corn production in the United States accounts for 11% of the global caloric production of the four basic commodities (corn, wheat, rice and soybeans), which themselves account for 75% of the calories that humans consume. The U.S. share is obtained by converting country-level production numbers from the Food and Agricultural Organisation (FAO) of the four commodities into calories (Roberts and Schlenker, 2013) using conversion ratios of calories per pound. Given the large market share of U.S. corn production, quantity changes have global price effects, especially since the demand and supply are highly inelastic.

Global production of the four basic commodities in 1990-2014 results in 5.92×10^{15} calories, or using a 2000 calories per day diet for 365 days a year, enough calories to feed 8.1 billion people. U.S. corn production accounted for 11%, or the caloric equivalent of feeding 893 million people. The substantial decline in peak ozone levels led to a 19% increase in yields (44% of the observed yield trend of 43%), or the caloric equivalent of feeding 170 million people, which is 2.1% of global production. Given the estimated demand and supply elasticities of Roberts and Schlenker (2013), this implies an 8-12% decrease in food prices, or roughly 10%.

The caloric cost of a 2000 calorie per day diet is 47 dollars per year using an average price of 4 dollars per bushel of corn, 56 pounds per bushel of corn, and 1107 calories per pound:

$$4 \frac{\text{dollars}}{\text{bushel}} \times \frac{1}{56} \frac{\text{bushel}}{\text{pound}} \times \frac{1}{1107} \frac{\text{pound}}{\text{calories}} \times 2000 \frac{\text{calories}}{\text{day}} \times 365 \frac{\text{days}}{\text{year}} = 47 \frac{\text{dollars}}{\text{year}}$$

The 10% price drop hence increased consumer surplus by $0.1 \times 8.1 \times 47 = 38$ billion dollars annually. The estimate of 4 dollars per bushel is conservative, as prices have been more than twice as high, which would increase the predicted gain in consumer surplus. Note that consumers of calories include farmers, as corn is used as feed for livestock.

Producers of corn in the U.S. see a 19% increase in yields, but a 10% decrease in price of both corn and soybeans, which are usually grown in rotation. The next effect is hence a small gain. On the other hand, foreign producers of the basic commodities lost through lower prices.

7 Conclusion

This study provides robust real-world evidence that short-term (hourly) exposure to high levels of ozone is harmful for crop yields. Using multiple empirical modeling approaches and a wide range of robustness checks, we identify a 65 ppb threshold above which hourly ozone concentrations are harmful, and the damaging effect on crop yields is well explained using a linear damage function. Laboratory studies have often relied on similar linear threshold models, although they usually adopt much lower thresholds of around 40 ppb. (Fuhrer et al., 1997) point out that exposure above 40 ppb gives a good fit for experimental data, but it is "less certain that it provides the best fit to data for [...] semi-natural communities." Our findings suggest that a higher threshold is appropriate for maize yields. Our results also provide support for potentially adopting an ozone standard based on hourly ozone readings, not 8-hour averages. The current U.S. standard is 70 ppb, which is not far from our threshold of 65 ppb. Note, however, that the 70 ppb standard is imposed on the highest daily 8-hour average, which can mask hourly spikes above 70 ppb.

Second, multiple studies have used complex simulation models to describe and predict the effects of ozone on yields at the regional or even global level (Avnery et al., 2011; Van Dingenen et al., 2009). The parameters used in these models are generally based on the results of chamber studies. One potential concern with this approach is that small errors in the estimated effects in micro-level chamber studies could add up to large errors in the macro-level analysis. The results of this paper and similar analyses could be used to validate these types of complex models.

Third, we find that falling pollution levels in recent years were a substantial driver of the growth in corn yields. The reduction in peak ozone levels accounted for 44% of the observed yield trend over the last 25 years, and increased the surplus of consumers of these commodities by $38 billion. Understanding the determinants of yield growth can help with modeling future crop production and prices. In the U.S., exposure to high concentrations of ozone has fallen to low levels in recent years, suggesting that further benefits to corn yield growth from reductions in air pollution are likely to be minimal (and therefore interventions in other areas may be required to sustain the high rate of yield growth). On the other hand, many developing countries have ozone levels that exceed those in the Eastern U.S., suggesting the potential for further yield gains through pollution reduction. However, there could be important differences across countries that cause the relationship between ozone and yields to differ from the results observed in this study. Given the relative lack of existing data from developing countries (Mauzerall and Wang, 2001), the substantial effects that we find in the U.S. suggest at the very least that it may be worth investing additional resources into data collection and analysis focused on pollution and agriculture in these regions.

References

Alston, Julian M., Philip G. Pardey, and Johannes Roseboom. 1998. Financing agricultural research: International investment patterns and policy perspectives. *World Development*, 26(6): 1057–1071.

Avnery, Shiri, Denise L. Mauzerall, Junfeng Liuc, and Larry W. Horowitz. 2011. Global crop yield reductions due to surface ozone exposure: 1. Year 2000 crop production losses and economic damage. *Atmospheric Environment*, 45(11): 2284–2296.

Burke, Marshall and Kyle Emerick. 2016. Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 8(3): 106–40, DOI: http://dx.doi.org/10.1257/pol.20130025.

Cassman, Kenneth G. 1999. Ecological intensification of cereal production systems: Yield potential, soil quality, and precision agriculture. *Proceedings of the National Academy of Sciences*, 96, p. 5952–5959.

Conley, Timothy G. 1999. GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1): 1–45.

Efron, Bradley and Robert J. Tibshirani. 1993. *An Introduction to the Bootstrap*: Chapman and Hall.

Environmental Protection Agency. 2008. National ambient air quality standards for ozone; final rule. *Federal Register*, 73(60): 16436–16514.

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. 2012. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *American Economic Review*, 102(7): 3749–3760.

Fishman, Jack, John K. Creilson, Peter A. Parker, Elizabeth A. Ainsworth, G. Geoffrey Vining, John Szarka, Fitzgerald L. Booker, and Xiaojing Xu. 2010. An investigation of widespread ozone damage to the soybean crop in the upper midwest determined from ground-based and satellite measurements. *Atmospheric Environment*, 44(18): 2248–2256.

Fuhrer, J., L. Skärby, and M.R. Ashmore. 1997. Critical levels for ozone effects on vegetation in Europe. *Environmental Pollution*, 97(1-2): 91–106.

Heagle, Allen S. 1989. Ozone and crop yield. *Annual Review of Phytopathology*, 27: 397–423.

Hill, Jason, Erik Nelson, David Tilman, Stephen Polasky, and Douglas Tiffany. 2006. Environmental, economic, and energetic costs and benefits of biodiesel and ethanol biofuels. *Proceedings of the National Academy of Sciences*, 103(30): 11206–11210.

Jorgenson, Dale W. and Frank M. Gollop. 1992. Productivity growth in U.S. agriculture: A postwar perspective. *American Journal of Agricultural Economics*, 74(3): 745–750.

Mauzerall, Denise L. and Xiaoping Wang. 2001. Protecting agricultural crops from the effects of tropospheric ozone exposure: Reconciling science and standard setting in the United States, Europe, and Asia. *Annual Review of Energy and the Environment*, 26: 237–268.

McCarthy, James E. 2010. Ozone air quality standards: EPA's proposed January 2010 revisions. Technical report, Congressional Research Service.

McGrath, Justin M., Amy M. Betzelberger, Shaowen Wang, Eric Shook, Xin-Guang Zhu, Stephen P. Long, and Elizabeth A. Ainsworth. 2015. An analysis of ozone damage to historical maize and soybean yields in the United States. *Proceedings of the National Academy of Sciences*, 112(46): 14390–14395, DOI: http://dx.doi.org/10.1073/pnas.1509777112.

Morgan, Patrick B., Timothy A. Mies, Germán A. Bollero, Randall L. Nelson, and Stephen P. Long. 2006. Season-long elevation of ozone concentration to projected 2050 levels under fully open-air conditions substantially decreases the growth and production of soybean. *New Phytologist*, 170(2): 333–342.

Roberts, Michael J. and Wolfram Schlenker. 2013. Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review*, 103(6): 2265–2295.

Royal Society. 2008. *Ground-level ozone in the 21st century: future trends, impacts and policy implications*, Science Policy Report 15/08: Royal Society.

Sanders, Nicholas J. and Alan Barreca. 2012. The external costs of pollution reduction: Agriculture and the incidence of the acid rain program. *Working Paper*.

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.

Van Dingenen, Rita, Frank J. Dentener, Frank Raes, Maarten C. Krol, Lisa Emberson, and Janusz Cofala. 2009. The global impact of ozone on agricultural crop yields under current and future air quality legislation. *Atmospheric Environment*, 43(3): 604–618.

Westenbarger, David A. and George B. Frisvold. 1995. Air pollution and farm-level crop yields: An empirical analysis of corn and soybeans. *Agriucltural and Resource Economics Review*, 24(2): 156–165.

Figure 1: Location of Ozone Monitors Used in Analysis (1990-2014)

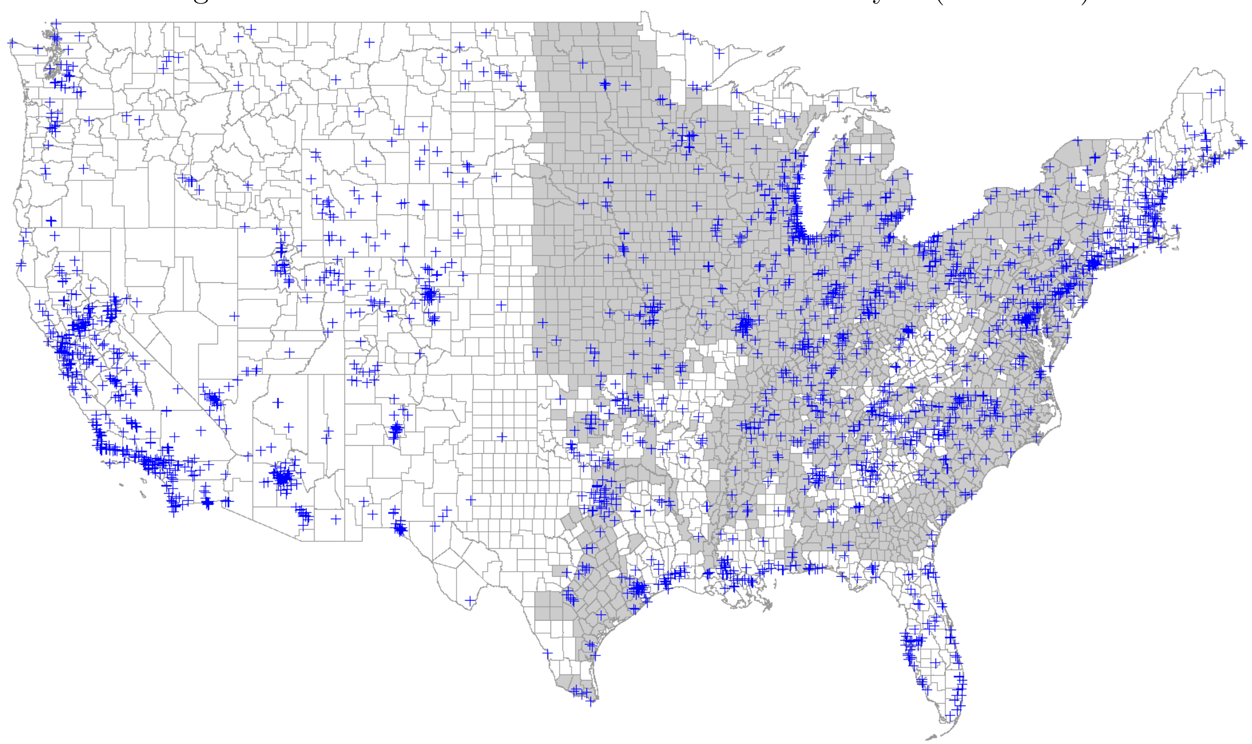

NOTE: Figure shows the locations of ozone monitors and the counties included in our sample. We use monitors if at least 75% of the observations are non-missing in March or April-August of a year. We separate the year into these two sub-periods as many ozone monitors only report for part of the year as shown in Figure 3. Counties east of the 100 degree meridian (except Florida) that report corn yields at least for 13 years in 1990-2014 (half the time) are shaded in grey.

Figure 2: Trend in Log Yields

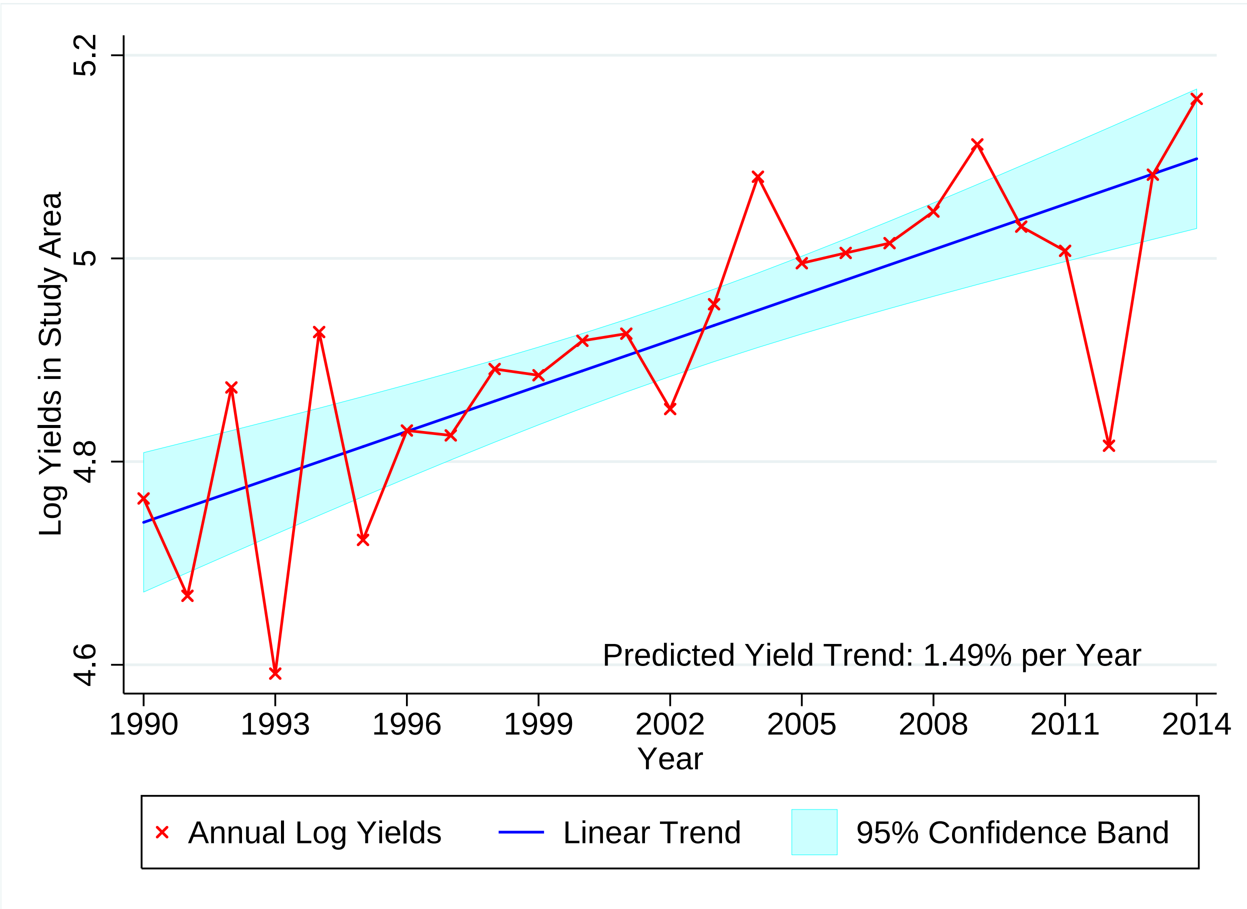

NOTE: Figure displays annual log yield averaged over the counties in our sample (shown in grey in Figure 1), weighted by harvested area. Counties included in our baseline sample, i.e., that report yields for at least half the years, accounted on average for 91% of the corn that was produced in the United States in 1990-2014. A trend is fitted to the annual aggregate data: yields on average grow 1.49% per year.

Figure 3: Seasonality of Ozone Monitors (1990-2014)

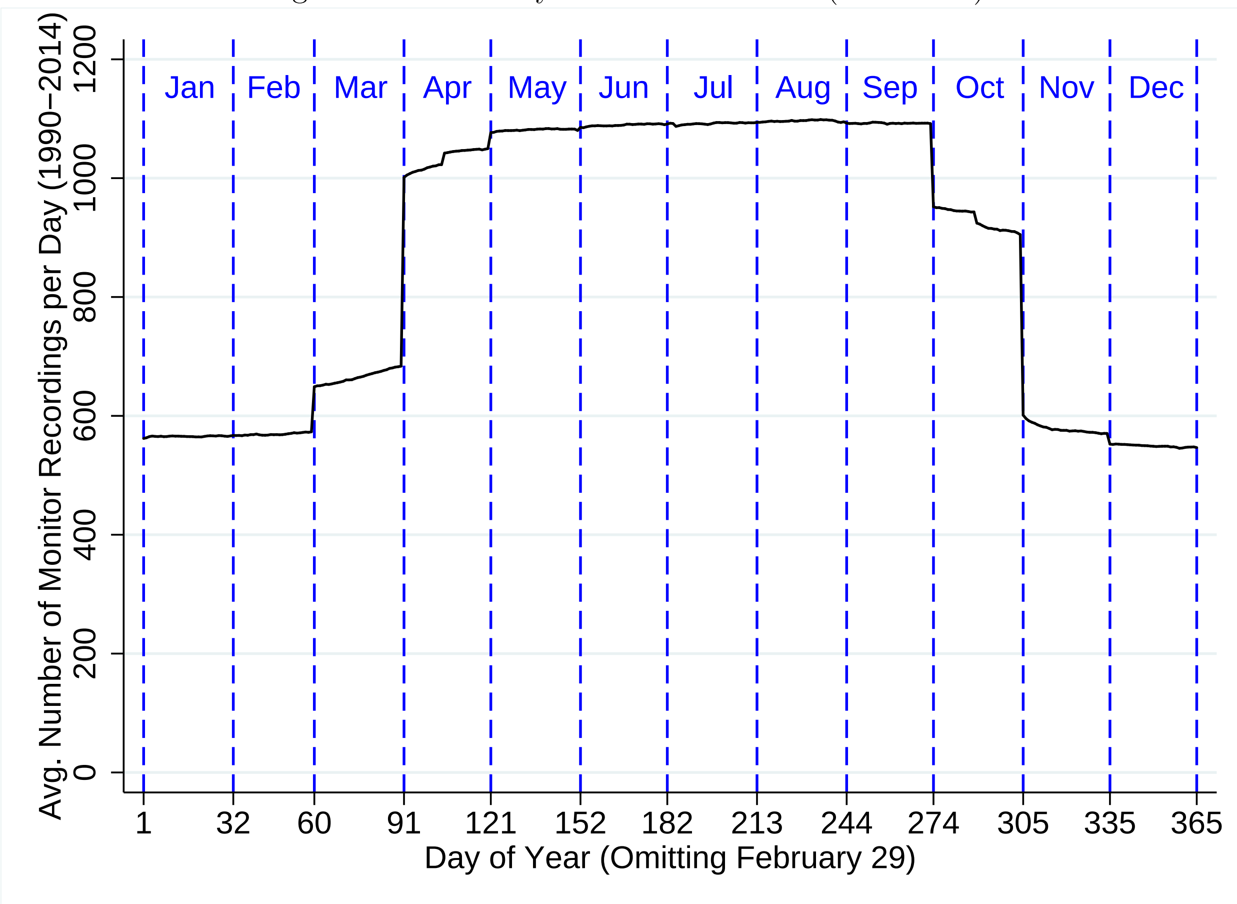

NOTE: Figure shows the number of monitors that are on average reporting on a particular day in 1990-2014. Several monitors come online in April. To make years consistent for this figure, we drop February 29th in leap years.

Figure 4: Number of Monitors Over Time

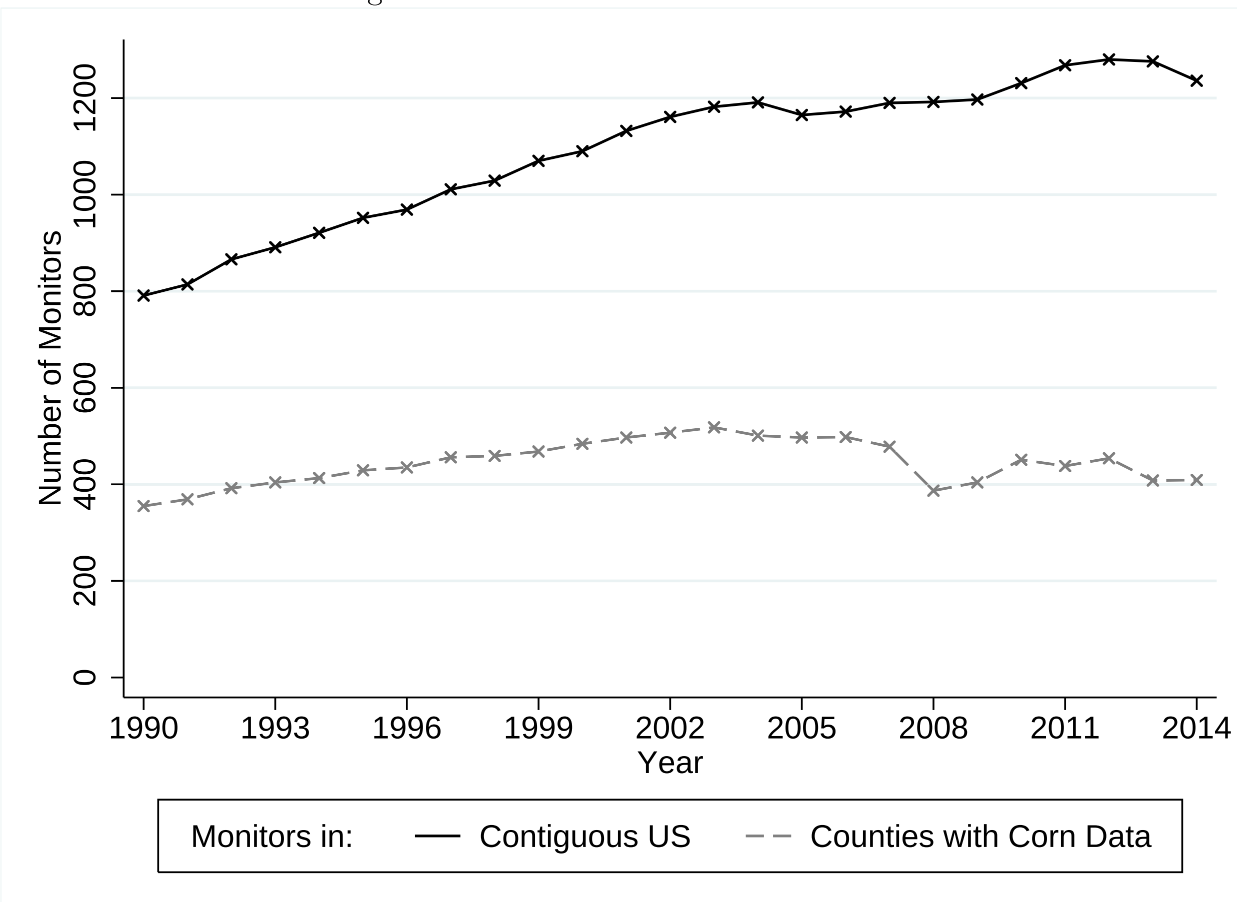

NOTE: Figure shows the number of monitors where at least 75% of the daily values in March or April-August are non-missing. We separately interpolate values in March and April-August, as there are fewer monitors reporting in March (See Figure 3). The black line shows the overall number of monitors in the database, and the grey line shows the number of monitors that fall into counties with corn yield data shown in grey in Figure 1.

Figure 5: Weighting Function W126 by EPA

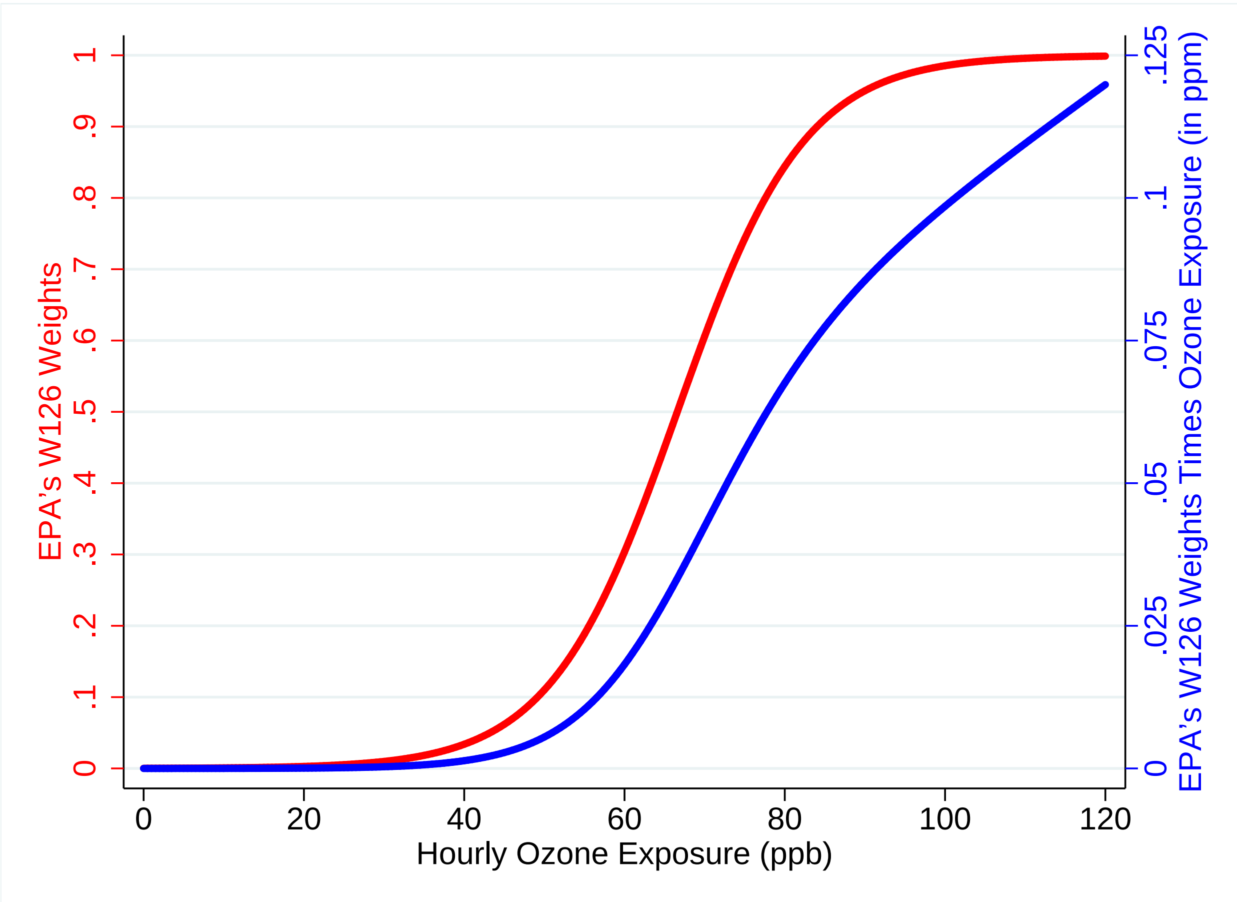

NOTE: Figure displays the weighting function W126 (left axis) as well as the product of the weight and the pollution exposure (right axis). The season-total weighted sum is obtained by multiplying each hourly ozone observation between 8am and 8pm during the growing season (March-August) by the weight and then summing all weighted observations, i.e., all blue values.

Figure 6: Cross Sectional Estimates of the Effect of Ozone on Corn Yields

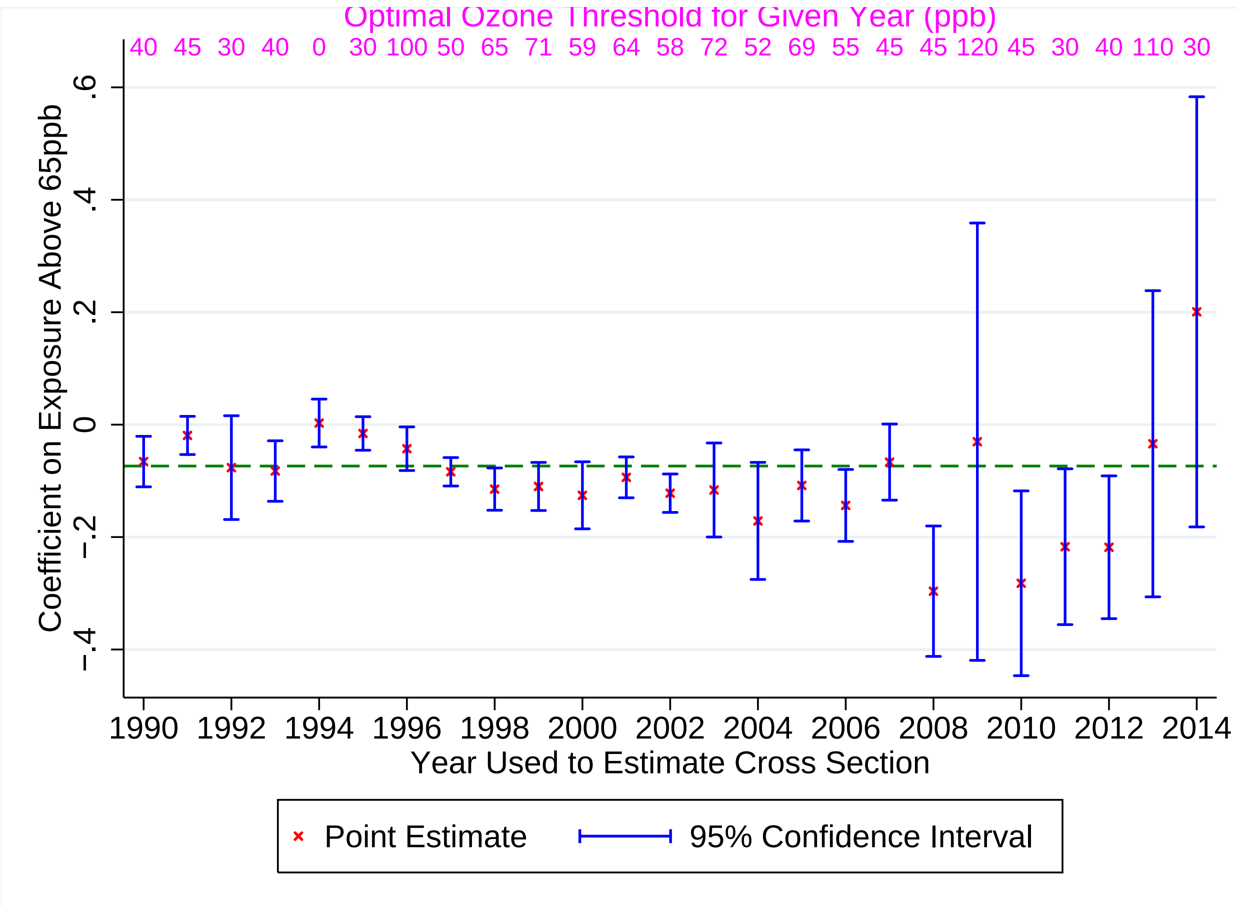

NOTE: Figure displays results of county-level cross-sectional regressions. Point estimates are shown as (x) and the 95% confidence intervals are added as (−). The x-axis indicates the year used in the estimation. We use the same ozone and weather measures as in our baseline model in column (4b) in Table 3. For comparison, the estimate of column (4b) in Table 3 is added as the green dashed line. The top of the figure gives the threshold that results in the highest R^2 or equivalently, the lowest AIC/BIC.

Figure 7: The Effect of Ozone on Corn Yields

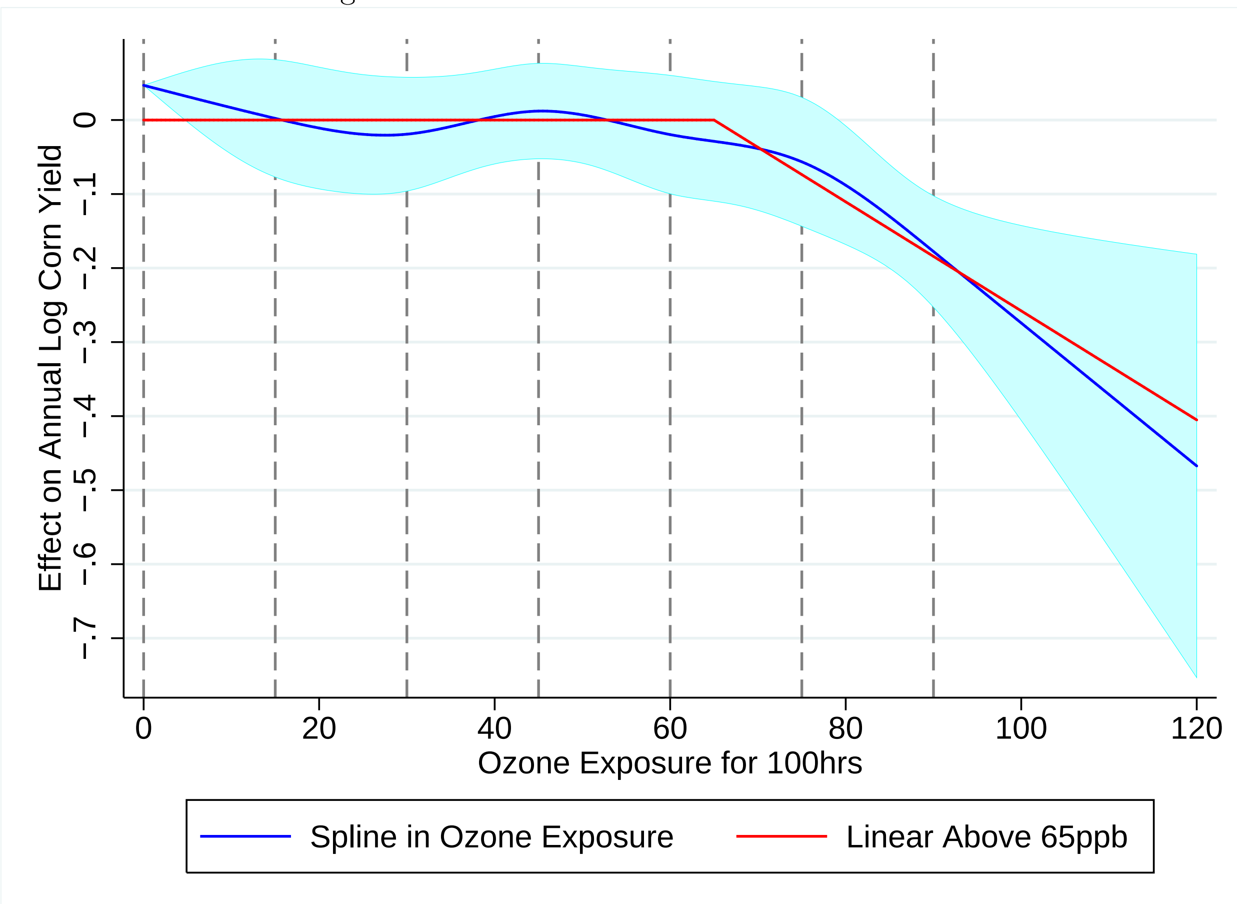

Note: Figure shows the estimated effect of ozone on log corn yields from two different models. First, the blue line shows the results of a model where log yields are allowed to flexibly depend on hourly ozone readings. The model uses restricted cubic splines in ozone with 7 knots (indicated by dashed lines at 1, 15, 30, 45, 60, 75, and 90ppb). The 95% confidence band is added as shaded area. The second model forces the effect of ozone to be linear above a threshold of 65ppb and is shown in red. The slope coefficient is estimated in a regression model. Since the regression model includes county fixed effects that allow for difference in average yields, the graph should be interpreted in relative terms, i.e., by comparing if pollution is shifted from one concentration to another. We normalized the spline plot so the estimated average effect below 65ppb is zero.

Figure 8: Highest In-sample R^2

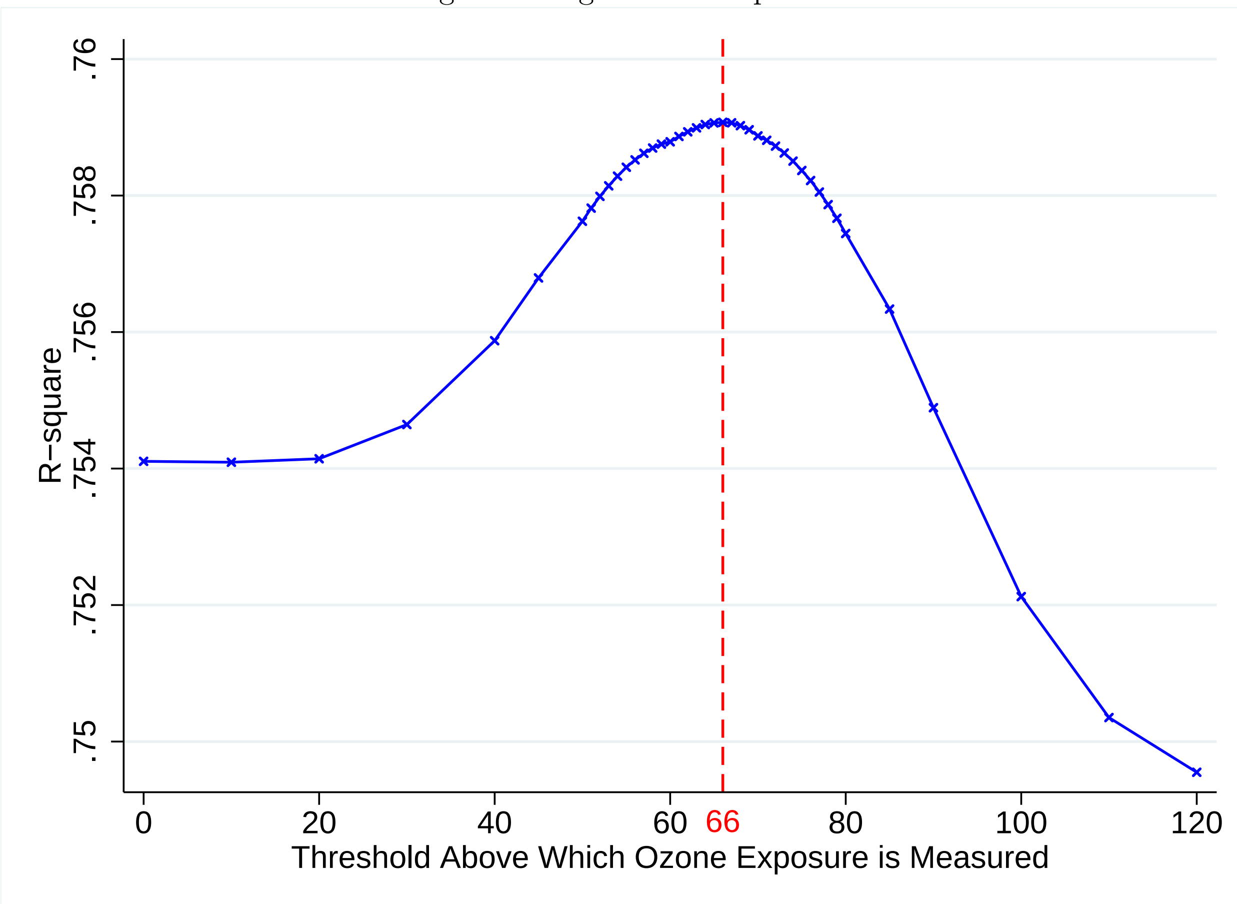

NOTE: Figure shows the R^2 for piecewise linear models. The x-axis varies the threshold above which ozone exposure is measured (being 10ppb above the threshold is forced to be 10 times as bad as being 1ppb above the threshold), while the y-axis gives the R^2 from a model that also includes the weather variables of Table 3, county-level quadratic time trends as well as year and county fixed effects. The highest R^2 is reached for a model that uses a threshold of 66ppb.

Figure 9: Out-of-Sample Prediction Error

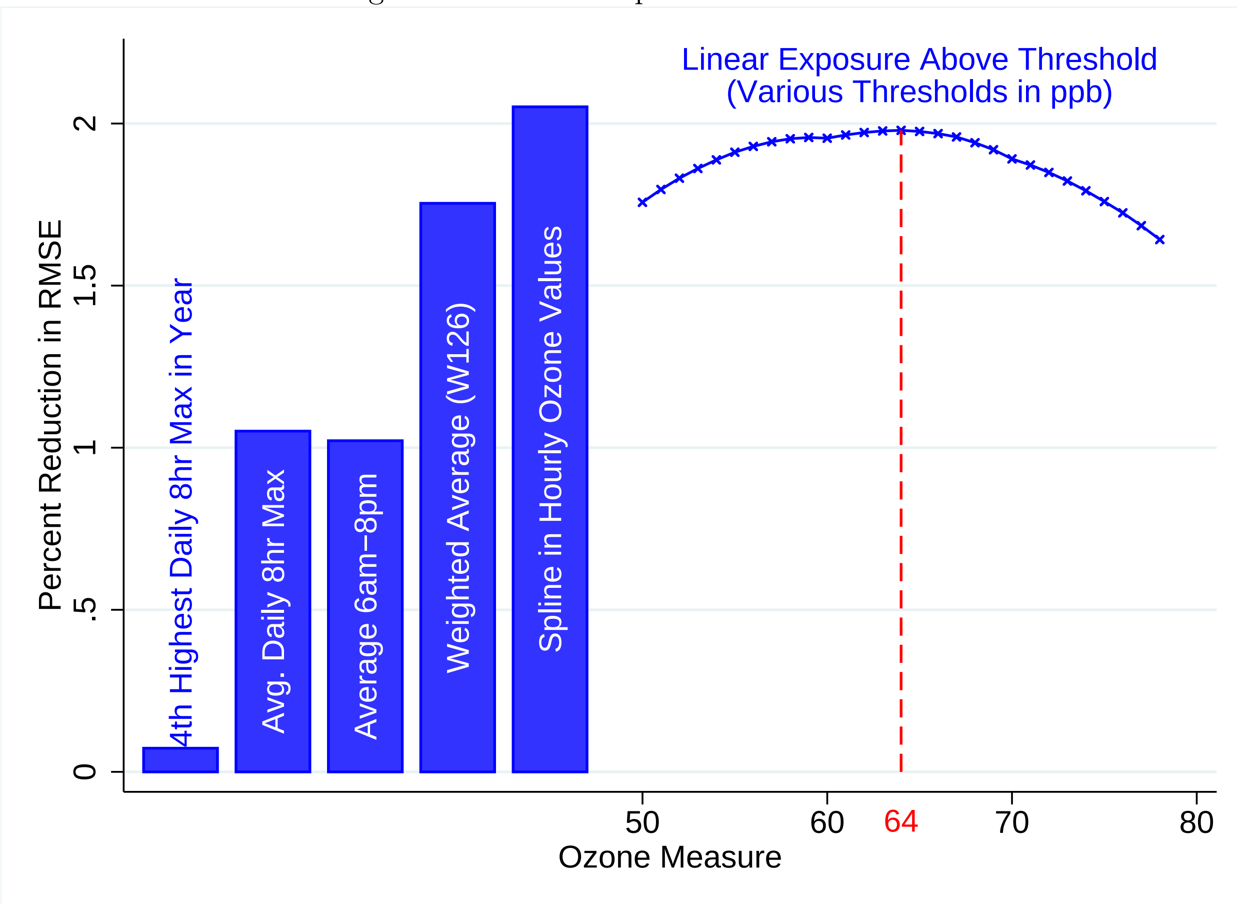

Note: Figure shows the reduction in out-of-sample prediction error compared to a model without any ozone control. All models include the four weather variables of Table 3, county-specific quadratic time trends as well as year and county fixed effects. The first five bar charts include the following ozone variables, respectively, as outlined in the specification in Table 1: (1a) The 4th highest of the daily maximum 8hr average (March-December, the remainder uses the growing season March-August); (1b) mean of all daily maximum 8-hour averages; (2) average of all hourly ozone readings 6am-8pm; (3) weighted sum of hourly ozone readings 8am-8pm (EPA's W126 weights); (5) restricted cubic spline in hourly ozone with 7 knots. The remaining line plots the reduction in out-of-sample prediction error as a function of the exposure threshold of a piecewise linear model (specification 4). For each model, we estimated the parameters 1000 times using 80% of the data and predicted yields for the remaining 20%.

Figure 10: The Effect of Observed Hourly Ozone Above 65ppb on Total Corn Production

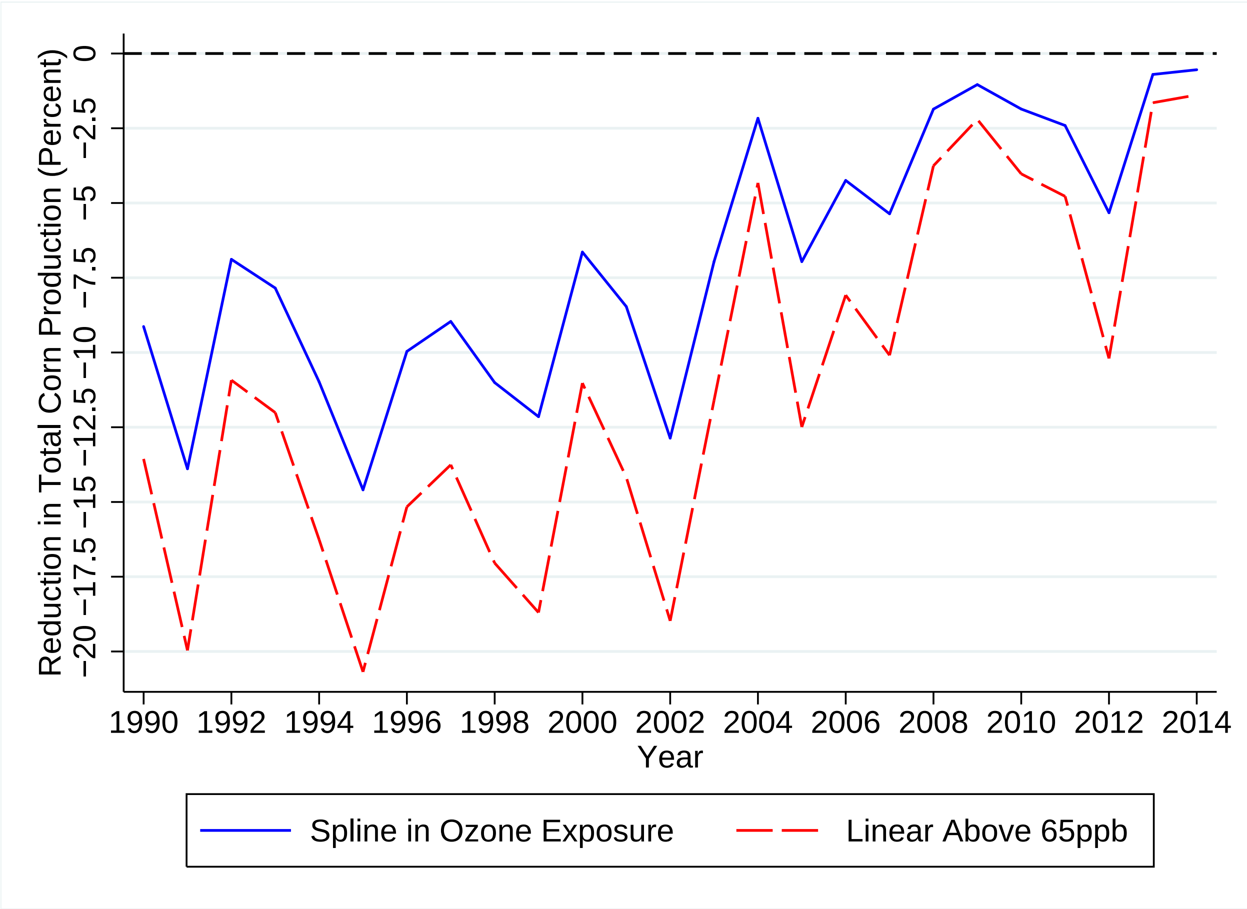

NOTE: Figure displays the percent reduction in overall corn production caused by hourly ozone levels above 65ppb. We compare predicted yields under observed ozone levels to a counterfactual where all hourly ozone readings above 65ppb are set to 65ppb.

Figure 11: Ozone Levels Over Time

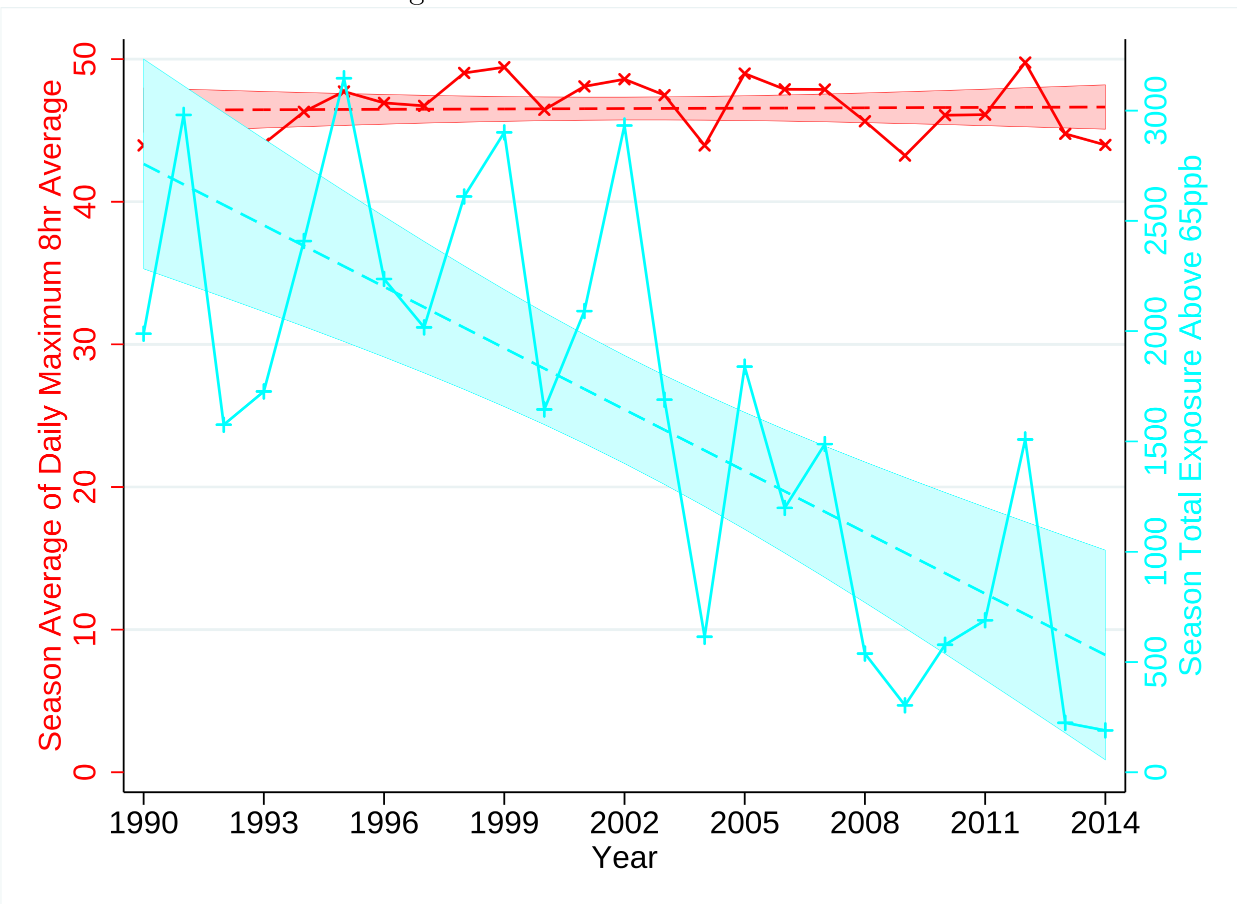

Note: Figure shows ozone levels over time. The solid red line (left y-axis) shows the average of the highest daily 8-hour average concentration over the growing season (March-August), while the solid cyan line (right y-axis) shows the cumulative exposure above 65ppb during daylight hours (6am-8pm) over the growing season. Ozone levels are weighted averages of the county-level data, where we weight each county by the average growing area in 1990-2014, i.e., weights do not change year-to-year. A linear trend (dashed lines) as well as a 95% confidence band on the trend are added.

Figure 12: Eliminating Hourly Ozone Readings above 65ppb and 40ppb under Various Model Specifications

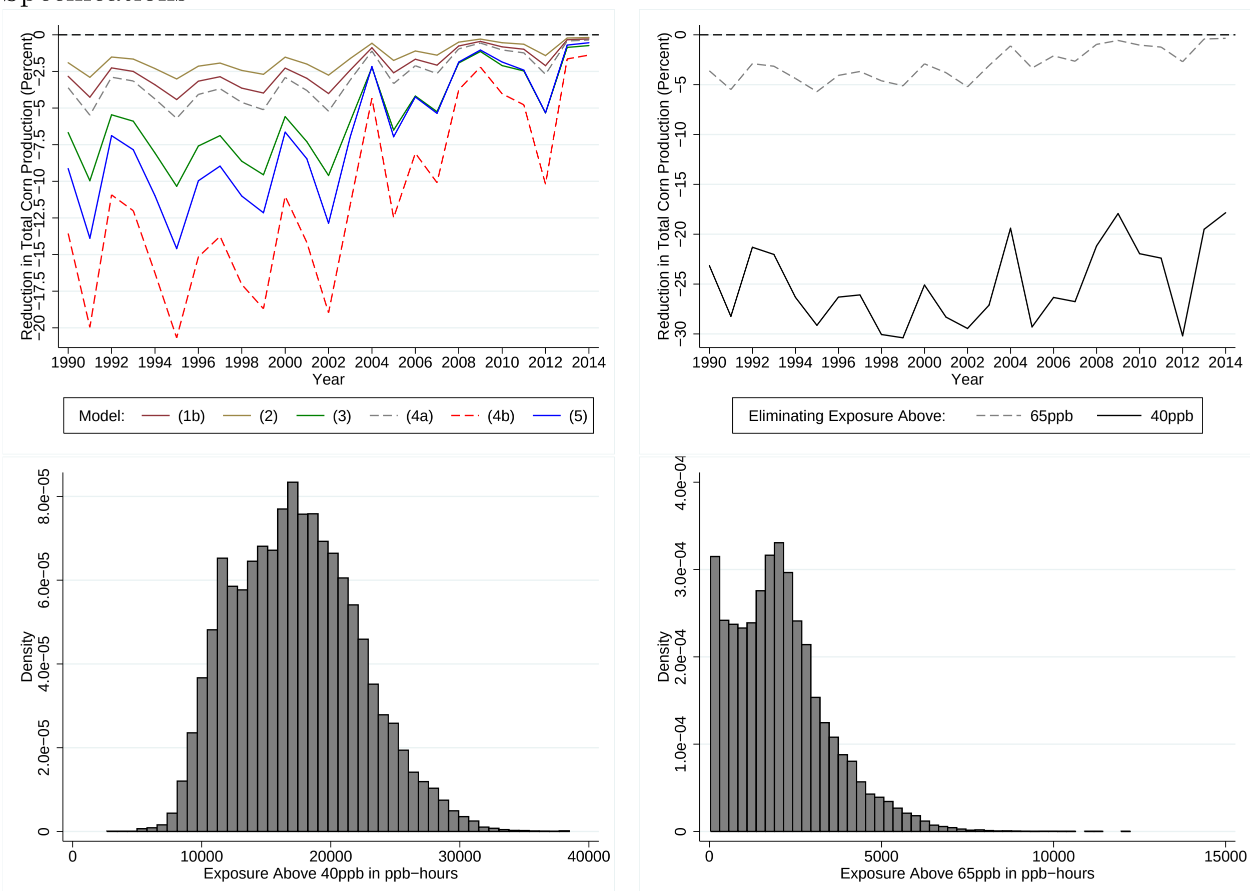

NOTE: Top panels display the percent reduction in overall maize production caused by ozone. We compare predicted yields under observed ozone levels to a counterfactual where all hourly ozone readings above a threshold are set to equal the threshold. The top left panel eliminates ozone above a threshold of 65ppb using various models; the model numbers correspond to the column headers of the regressions in Table 3. In the top right panel, the model incorporating linear exposure above a threshold of 40ppb (Model 4a in Table 3) is used to examine the scenarios where all hourly ozone readings above 65ppb are set to 65ppb (dashed line) and all hourly ozone readings above 40ppb are set to 40ppb (solid line). The bottom row shows the histogram of observed linear ozone exposure above 40ppb (bottom left panel) as well as 65ppb (bottom right panel).

Figure 13: Fraction of Yield Trend Due to Reduction in Peak Ozone

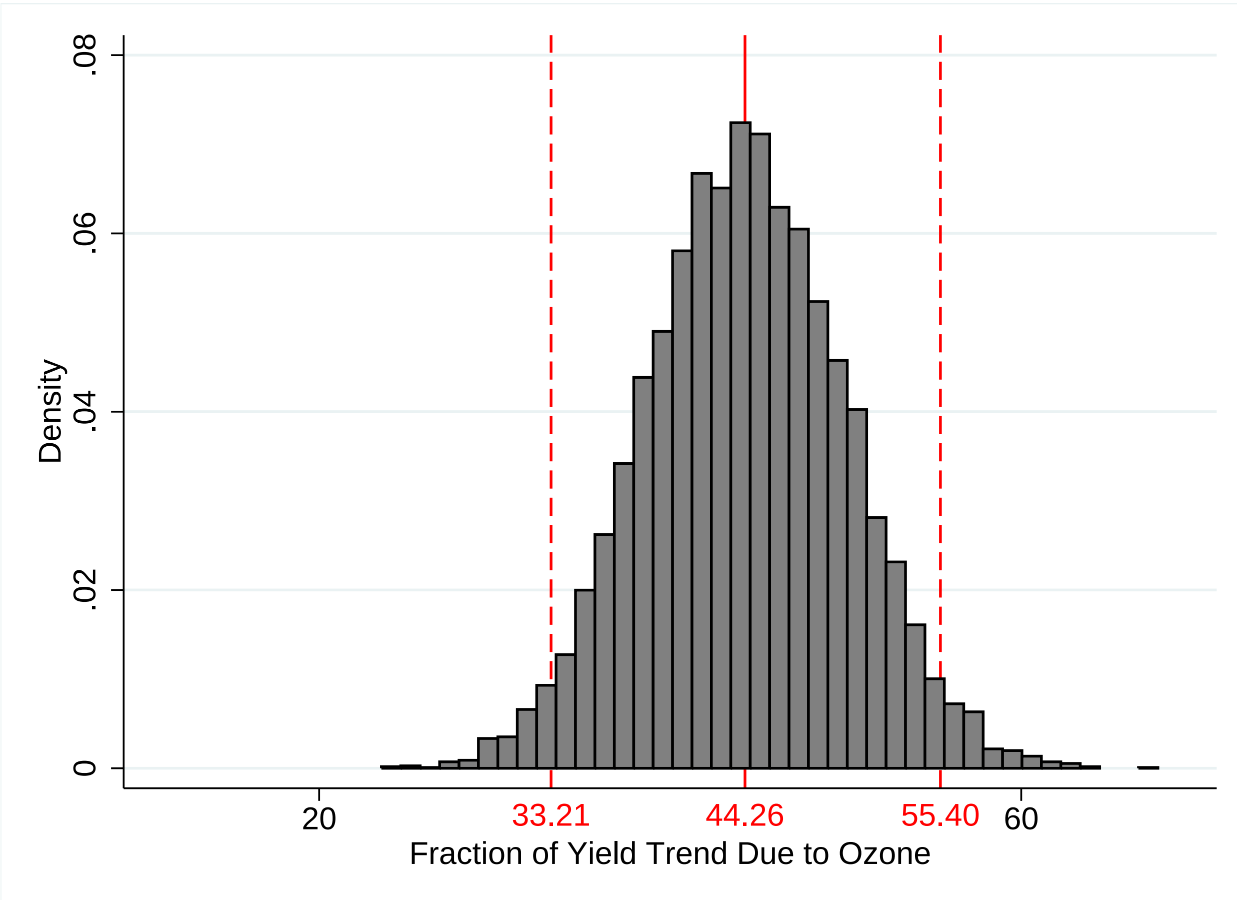

NOTE: Figure displays the distribution of the 10000 draws to evaluate the fraction of the observed yield trend (from 1990 to 2014) that is due to reduction in peak ozone levels as measured by ozone exposure above 65ppb. The mean is 44% with a standard error of 5.6%. The 95% confidence interval reaches from 33% to 55%.

Figure 14: Yield Trend Due to Trends in Peak Ozone Levels

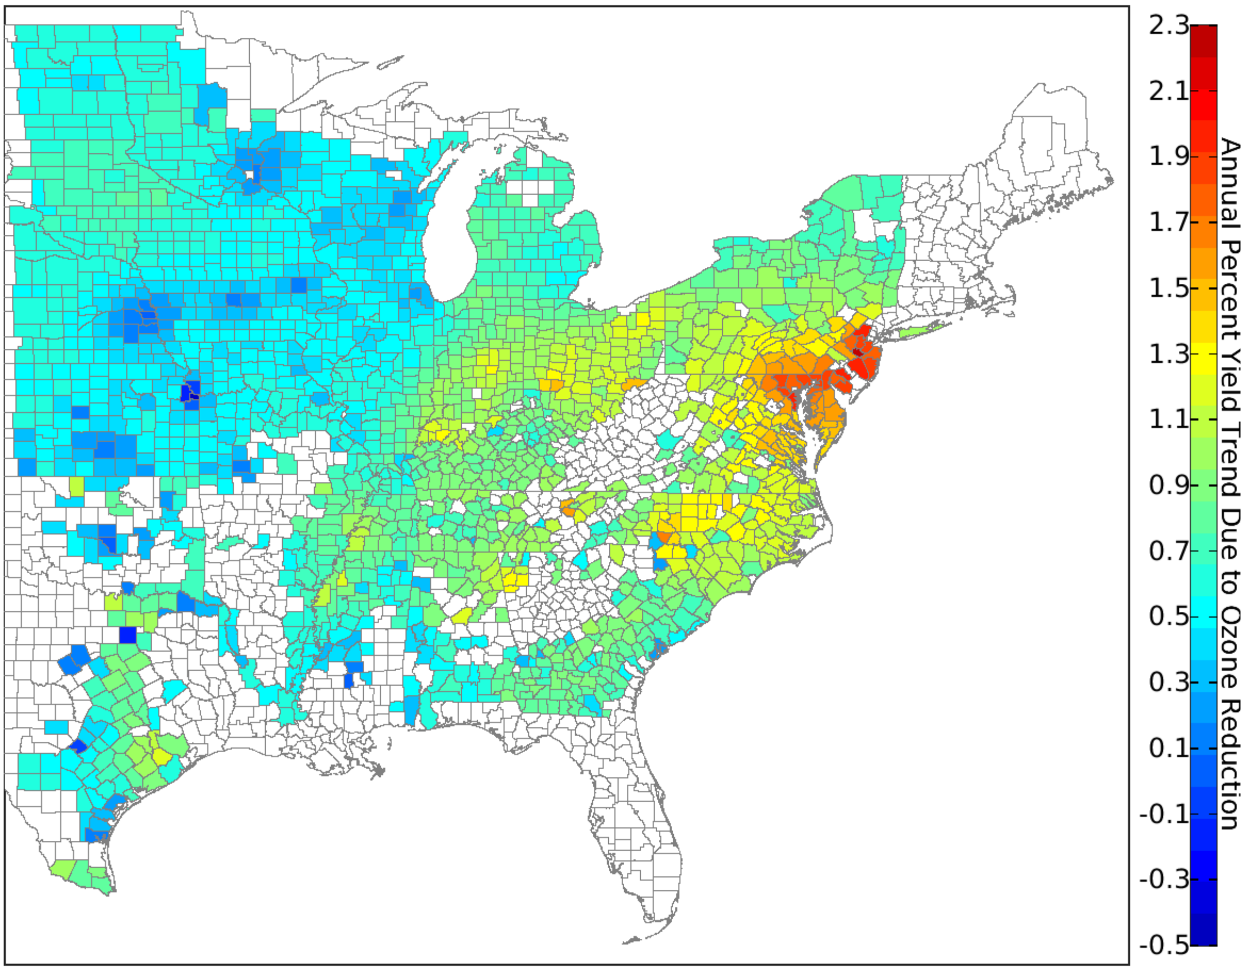

NOTE: Figure displays the annual trend in log yields due to the observed trends in season-total exposure above 65ppb. We derive this by estimating a separate trend in season-total exposure above 65ppb for each county and then multiply the trend by the estimated coefficient on the ozone measure from our panel regression. We only include counties that report yields for half of the years in our sample, i.e., have at least 13 observations in 1990-2014. The trend in overall log yields in our sample of counties is 1.49% per year (Figure 2).

Table 1: Definition of Ozone Specifications

Spec	Description
(1)	**Daily maximum 8-hour average in ppb**. We calculate average ozone levels for consecutive 8-hr intervals starting at each hour of the day, i.e., the average ozone concentration from midnight to 8am, 1am-9am, etc. We utilize the maximum daily 8-hr average in two ways:
(1a)	We pick the fourth highest for the year (March-December). We separately interpolate the ozone data for October and November-December using the above procedure. The current U.S. ambient air quality standard is based on the fourth highest of these daily maximum 8-hr averages, averaged over three consecutive years. This boils down to the pollution reading on one day of the growing season.
(1b)	We take the average of all daily maximum 8-hr averages over the growing season (March-August).
(2)	**Daily mean in ppb:** We derive the simple average of all hourly observations between 6am and 8pm of each day and then average the daily values over all days of the growing season (March-August).
(3)	**Weighted sum of hourly exposures by EPA:** EPA recently proposed a new secondary ozone standard based on a weighted sum of all hourly observations. The primary ambient air quality standard is designed to protect human health, while the secondary standard is designed to protect human welfare, i.e., visibility, damages to crop, etc. The proposed revision to the secondary standard is outlined at https://www.gpo.gov/fdsys/pkg/FR-2010-01-19/pdf/2010-340.pdf. (See also McCarthy (2010) and https://www.epa.gov/ozone-pollution/2008-national-ambient-air-quality-standards-naaqs-ozone.) The weighting function is shown in Figure 5. Specifically, we multiply all hourly ozone observations between 8am and 8pm by the appropriate weight and sum these values over all days in the growing season (March-August). We follow EPA, which uses a time period of 8am to 8pm, but results are the same if we start at 6am instead to keep the time of the day consistent with other measures. Since ozone is generally very low in the morning, this has no significant effect on the measure.
(4)	**Linear hourly exposure above a threshold in ppb-hours:** Earlier U.S. ambient air quality standard were based on hourly observations, and chamber studies have suggested that the damaging effects of ozone are linearly increasing above a threshold—for example, being twice as high above the threshold is twice as harmful. We therefore derive linear exposure measures above various thresholds b. For each hour between 6am and 8pm, we calculate the difference between the monitor reading and the threshold if the monitor reading is larger than the threshold and then sum it for these hours of the day for all days of the growing season. For example, a linear exposure measure above 65ppb sums $\max\{v-65,0\}$ for all hourly values v between 6am and 8pm for all days of the growing season (March-August).
(5)	**Spline polynomials in hourly ozone exposure:** Instead of imposing weights on various ozone readings, we estimate the effect of various hourly ozone levels on annual yields. We use flexible restricted cubic splines, which are a series of third-order polynomials that approximates the unknown function between consecutive knot locations subject to the constraints that the polynomials continuously "blend" into one another (are continuous and have a continuous derivative at the knot). For a restricted cubic spline, the function is forced to be linear below the smallest and above the largest knot. Since we are pairing annual yields with a season worth of hourly pollution readings, we need to aggregate the hourly readings. While the functional form is flexible, the assumption we impose is that the effects are additively separable. Specifically, let the effect of an hourly ozone exposure o_{hit} be given by the splines $s_1(o_{hit}),\dots,s_{n-1}(o_{hit})$, where n is the number of spline knots. In our baseline specification we use seven knots at 1ppb, 15ppb, 30ppb, 45ppb, 60ppb, 75ppb, and 90ppb. Summing the effect over all hours h of the growing season we get: $$\sum_h \sum_{k=1}^{n-1} \alpha_k s_k(o_{hit}) = \sum_{k=1}^{n-1} \alpha_k \underbrace{\sum_h s_k(o_{hit})}_{S_{kit}} = \sum_{k=1}^{n-1} \alpha_k S_{kit}$$ We are hence left with $n-1$ variables $S_{1it},\dots,S_{(n-1)it}$ which are the spline polynomials evaluated at each hourly reading and then summed over the entire growing season (March-August). The coefficients $\alpha_1,\dots,\alpha_{n-1}$ then give us the estimated effect of various ozone levels on annual log yields.

Table 2: Bootstrap Procedure: Fraction of Yield Trend Explained by Ozone Reduction

Step	**Description**
1)	Estimate our preferred model specification (columns (4b) and (5) in Table 3).
2)	Take 10,000 random draws of the joint distribution of all parameters of the model in 1. For each of the draws, evaluate steps 3-5.
3)	Calculate yield trend with observed historical peak ozone levels.
3a)	Get predicted yields in each county using the observed variables and parameters from step 2.
3b)	Derive the average yield in a year, which is the area-weighted average of all predicted yields from step 3a. In the baseline we use the observed corn acreage.
3c)	Estimate the trend in predicted annual yields from step 3b.
4)	Calculate yield trend under counterfactual where peak ozone is eliminated.
4a)	Get predicted yields in each county using the observed variables and parameters from step 2 except that *ozone above 65ppb is set to 65ppb.*
4b)	Derive the average yield in a year, which is the area-weighted average of all predicted yields from step 4a. In the baseline we use the observed corn acreage.
4c)	Estimate the trend in predicted annual yields from step 4b.
5)	Fraction explained by ozone is 1-$\frac{\text{Trend in 4c}}{\text{Trend in 3c}}$. We plot the 10,000 outcomes.

Table 3: Effect of Ozone on Log Corn Yields under Various Specifications

Model	(0)	(1a)	(1b)	(2)	(3)	(4a)	(4b)	(5)
Ozone Measure		-0.004***	-0.219***	-0.252***	-0.152***	-0.188***	-0.737***	
		(0.001)	(0.040)	(0.048)	(0.020)	(0.028)	(0.088)	
Ozone - Spline 1								-0.030
								(0.036)
Ozone - Spline 2								0.030
								(0.441)
Ozone - Spline 3								0.378
								(1.249)
Ozone - Spline 4								-1.384
								(1.714)
Ozone - Spline 5								1.760
								(1.926)
Ozone - Spline 6								-1.585
								(2.000)
DDays 10-29°C (1000)	0.296**	0.298**	0.384***	0.382***	0.359***	0.380***	0.303***	0.291***
	(0.117)	(0.117)	(0.115)	(0.114)	(0.109)	(0.112)	(0.106)	(0.097)
DDays $\geq$ 29°C (100)	-0.555***	-0.553***	-0.506***	-0.508***	-0.487***	-0.494***	-0.498***	-0.496***
	(0.060)	(0.060)	(0.054)	(0.054)	(0.050)	(0.051)	(0.051)	(0.052)
Precipitation (m)	1.119***	1.115***	1.147***	1.140***	1.095***	1.112***	1.077***	1.067***
	(0.209)	(0.207)	(0.201)	(0.201)	(0.197)	(0.200)	(0.191)	(0.191)
Precipitation (m) Squared	-0.850***	-0.848***	-0.928***	-0.923***	-0.896***	-0.909***	-0.870***	-0.861***
	(0.165)	(0.163)	(0.156)	(0.157)	(0.154)	(0.156)	(0.151)	(0.152)
R-squared	0.2573	0.2584	0.2714	0.2714	0.2818	0.2766	0.2861	0.2874
Observations	40552	40552	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25	25	25

NOTE: Table regresses log yields in counties east of 100 degree meridian (excluding Florida) on ozone and weather controls over the main growing season March-August for the years 1990-2014. All specifications control for weather (degree days 10-29°C, degree days above 29°C, and a quadratic in precipitation) and include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or temporal controls, i.e., only what is explained by the ozone or weather variables. Columns differ in how the effect of ozone is modeled. Specifications of Table 1 are indicated in the top row:
(0): No control for ozone
(1a): 4th highest of daily maximum 8-hour averages during the season in ppb
(1b): Season average of daily maximum 8-hour averages in ppb
(2): Season average of daily ozone average (6am-8pm) in ppb
(3): Weighted sum of hourly exposure 8am-8pm (EPA's W126 weights) in 10 ppm
(4a): Linear exposure above 40ppb (6am-8pm) in 1000ppb-hrs
(4b): Linear exposure above 65ppb (6am-8pm) in 10000ppb-hrs
(5): Restricted cubic spline in ozone with 7 knots for ozone readings between 6am and 8pm in ppm.
Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Effect of Ozone on Log Corn Yields - Higher Order Temperature Terms

	(1)	(2a)	(2b)	(2c)	(2d)	(2e)	(3)	(4)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.761***	-0.675***	-0.637***	-0.647***	-0.659***	-0.649***	-0.644***
	(0.088)	(0.089)	(0.099)	(0.097)	(0.099)	(0.096)	(0.100)	(0.100)
Degree Days 10-29°C (1000)	0.303***	0.479***	1.560***	-0.038	0.805	-0.058	-0.066	-0.432
	(0.106)	(0.092)	(0.531)	(0.691)	(0.932)	(1.904)	(1.809)	(1.819)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.499***	-0.637***	-1.828***	-2.421***	-2.375***	-2.610***	-2.649***
	(0.051)	(0.054)	(0.168)	(0.264)	(0.316)	(0.541)	(0.541)	(0.474)
Precipitation (m)	1.077***	1.135***	1.235***	1.328***	1.320***	1.315***	1.275***	1.255***
	(0.191)	(0.196)	(0.225)	(0.213)	(0.212)	(0.208)	(0.204)	(0.206)
Precipitation (m) Squared	-0.870***	-0.890***	-0.984***	-1.041***	-1.024***	-1.009***	-0.972***	-0.956***
	(0.151)	(0.154)	(0.175)	(0.169)	(0.171)	(0.168)	(0.162)	(0.163)
R-squared	0.2861	0.2871	0.2979	0.3068	0.3089	0.3115	0.3161	0.3171
Opt. Threshold (R^2/AIC/BIC)	66	65	67	68	68	68	68	68
Observations	40552	40552	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25	25	25
Min Temperature Polynomial	-	1	2	3	4	5	5	5
Max Temperature Polynomial	-	1	2	3	4	5	5	5
Max-Min Temp Polynomial	-	-	-	-	-	-	5	5
Avg Temperature Polynomial	-	-	-	-	-	-	-	5

NOTE: Column (1) is the same as column (4b) of Table 3, while additional columns add further weather controls. The last four rows give the highest order polynomials that are included for minimum and maximum temperature, diurnal range (maximum-minimum temperatures), and average temperature. Columns (2a)-(2e) include various polynomials of minimum and maximum temperature, while column (3) adds the diurnal range, and column (4) adds average temperature. Higher order temperatures are first calculated for each day and then averaged over the growing season. All specifications include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Effect of Ozone on Log Corn Yields - Higher Order Daily Precipitation

	(1)	(2a)	(2b)	(2c)	(2d)	(3)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.761***	-0.732***	-0.732***	-0.731***	-0.731***
	(0.088)	(0.089)	(0.087)	(0.088)	(0.087)	(0.087)
Degree Days 10-29°C (1000)	0.303***	0.479***	0.302***	0.302***	0.302***	0.303***
	(0.106)	(0.092)	(0.106)	(0.106)	(0.106)	(0.106)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.499***	-0.496***	-0.496***	-0.496***	-0.496***
	(0.051)	(0.054)	(0.051)	(0.052)	(0.052)	(0.052)
Precipitation (m)	1.077***	1.135***	1.091***	1.090***	1.102***	1.127***
	(0.191)	(0.196)	(0.192)	(0.199)	(0.208)	(0.197)
Precipitation (m) Squared	-0.870***	-0.890***	-0.857***	-0.857***	-0.858***	-0.861***
	(0.151)	(0.154)	(0.156)	(0.156)	(0.157)	(0.155)
R-squared	0.2861	0.2871	0.2862	0.2862	0.2862	0.2863
Opt. Threshold (R^2/AIC/BIC)	66	65	66	66	66	66
Observations	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25
Daily Precipitation Polynomial	-	2	3	4	5	5
Min Temperature Polynomial	-	-	-	-	-	5
Max Temperature Polynomial	-	-	-	-	-	5
Max-Min Temp Polynomial	-	-	-	-	-	5
Avg Temperature Polynomial	-	-	-	-	-	5

NOTE: Column (1) is the same as column (4b) of Table 3, while additional columns add further weather controls. The last four rows give the highest order polynomials that are included for daily precipitation, minimum and maximum temperature, diurnal range (maximum-minimum temperatures), and average temperature. Columns (2a)-(2d) include various polynomials of daily precipitation, while column (3) adds all other temperature controls. Higher order polynomials are first calculated for each day and then averaged over the growing season. All specifications include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Effect of Ozone on Log Corn Yields - Interactions with Temperature and PM_{10}

	(1)	(2a)	(2b)	(2c)	(2d)	(2e)	(2f)	(2g)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.617***	-0.620***	-0.579***	-0.784***	-0.664***	-0.668***	-0.630***
	(0.088)	(0.078)	(0.077)	(0.075)	(0.092)	(0.084)	(0.083)	(0.080)
$\times$ Degree Days $\geq$ 29°C		-0.141***		-0.086**		-0.145***		-0.108***
		(0.035)		(0.039)		(0.034)		(0.038)
$\times$ Maximum Temperature			-0.182***	-0.129***			-0.156***	-0.082**
			(0.035)	(0.046)			(0.034)	(0.041)
$\times$ Mean PM_{10}					0.101**	0.108**	0.069	0.090**
					(0.043)	(0.042)	(0.042)	(0.040)
Mean PM_{10} (ppm)					-0.024	-0.032	-0.012	-0.024
					(0.030)	(0.030)	(0.030)	(0.030)
Maximum Temperature (C)			0.020*	0.012			0.017	0.006
			(0.012)	(0.014)			(0.011)	(0.013)
Degree Days 10-29°C (1000)	0.303***	0.285**	0.376***	0.357***	0.297***	0.279***	0.365***	0.340***
	(0.106)	(0.113)	(0.122)	(0.124)	(0.101)	(0.107)	(0.114)	(0.116)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.379***	-0.491***	-0.420***	-0.497***	-0.375***	-0.492***	-0.401***
	(0.051)	(0.055)	(0.052)	(0.061)	(0.052)	(0.056)	(0.054)	(0.060)
Precipitation (m)	1.077***	1.037***	0.993***	0.990***	1.069***	1.027***	0.999***	0.996***
	(0.191)	(0.192)	(0.189)	(0.191)	(0.184)	(0.184)	(0.182)	(0.184)
Precipitation (m) Squared	-0.870***	-0.835***	-0.823***	-0.814***	-0.866***	-0.830***	-0.826***	-0.817***
	(0.151)	(0.152)	(0.148)	(0.150)	(0.147)	(0.148)	(0.146)	(0.147)
R-squared	0.2861	0.2903	0.2910	0.2922	0.2891	0.2935	0.2925	0.2942
Opt. Threshold (R^2/AIC/BIC)	66	58	55	55	59	54	54	53
Observations	40552	40552	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25	25	25

NOTE: Column (1) is the same as column (4b) of Table 3, while other columns additionally control for interactions with temperature and PM_{10}. Interaction terms are standardized by removing the mean and normalizing by the standard deviation of the demeaned series to make the coefficients easier to interpret. All specifications include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Effect of Ozone on Log Corn Yields - Controlling for Other Pollutants

	(1)	(2a)	(2b)	(2c)	(2d)	(3)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.738***	-0.733***	-0.738***	-0.742***	-0.741***
	(0.088)	(0.088)	(0.085)	(0.087)	(0.087)	(0.084)
Mean CO (10 ppm)		-0.372				-0.374
		(0.525)				(0.503)
Mean NO_x (100 ppm)			-0.185			-0.238
			(0.400)			(0.384)
Mean PM_{10} (100 ppm)				0.277		0.270
				(0.276)		(0.279)
Mean SO_2 (100 ppm)					0.635	0.682
					(0.652)	(0.638)
Degree Days 10-29°C (1000)	0.303***	0.300***	0.304***	0.296***	0.303***	0.294***
	(0.106)	(0.106)	(0.107)	(0.102)	(0.106)	(0.103)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.497***	-0.498***	-0.499***	-0.497***	-0.499***
	(0.051)	(0.051)	(0.051)	(0.052)	(0.051)	(0.052)
Precipitation (m)	1.077***	1.076***	1.080***	1.081***	1.073***	1.080***
	(0.191)	(0.192)	(0.191)	(0.188)	(0.191)	(0.186)
Precipitation (m) Squared	-0.870***	-0.869***	-0.871***	-0.868***	-0.866***	-0.865***
	(0.151)	(0.151)	(0.151)	(0.152)	(0.151)	(0.151)
R-squared	0.2861	0.2862	0.2861	0.2867	0.2863	0.2871
Opt. Threshold (R^2/AIC/BIC)	66	66	66	66	66	65
Observations	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25

NOTE: Column (1) is the same as column (4b) of Table 3, while other columns additionally control for other pollutants: carbon monoxide (CO), nitrogen oxides (NO_x), particulate matter (PM_{10}) and sulfur dioxide (SO_2). Columns (2a)-(2d) control for one pollutant at a time, while column (3) controls for all four. All specifications include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Effect of Ozone on Log Corn Yields - Various Time Controls Part I

	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.528***	-0.727***	-0.392***	-0.513***	-0.883***
	(0.088)	(0.080)	(0.086)	(0.075)	(0.069)	(0.056)
Degree Days 10-29°C (1000)	0.303***	0.327***	0.316***	0.297***	0.187	0.440***
	(0.106)	(0.078)	(0.107)	(0.082)	(0.118)	(0.084)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.549***	-0.499***	-0.568***	-0.506***	-0.499***
	(0.051)	(0.059)	(0.053)	(0.065)	(0.058)	(0.056)
Precipitation (m)	1.077***	1.051***	1.041***	1.007***	1.030***	0.965***
	(0.191)	(0.197)	(0.181)	(0.203)	(0.184)	(0.231)
Precipitation (m) Squared	-0.870***	-0.879***	-0.842***	-0.851***	-0.834***	-0.860***
	(0.151)	(0.182)	(0.147)	(0.188)	(0.148)	(0.212)
R-squared	0.2861	0.3575	0.2756	0.3372	0.2490	0.3665
Opt. Threshold (R^2/AIC/BIC)	66	71	66	78	57	73
Observations	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25
County-Level Trends	2	2	1	1	-	-
Year Fixed Effects	Yes	No	Yes	No	Yes	No

NOTE: Column (1a) is the same as column (4b) of Table 3, while other columns vary the time controls. The last two rows of the table give the highest order polynomial in the county-specific time trends and whether year fixed effects are included. Columns (a) include year fixed effects, while columns (b) do not. Columns (1)-(3) vary the highest order polynomial in the county-specific time trends. All specifications include county fixed effects. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Effect of Ozone on Log Corn Yields - Various Time Controls Part II

	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)	(5a)	(5b)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.528***	-0.674***	-0.497***	-0.699***	-0.386***	-0.513***	-0.418***	-0.513***	-0.316***
	(0.088)	(0.080)	(0.077)	(0.073)	(0.082)	(0.072)	(0.069)	(0.057)	(0.069)	(0.059)
Degree Days 10-29°C (1000)	0.303***	0.327***	0.277***	0.300***	0.300***	0.286***	0.187	0.258***	0.187	0.239***
	(0.106)	(0.078)	(0.104)	(0.081)	(0.105)	(0.081)	(0.118)	(0.082)	(0.118)	(0.083)
Degree Days $\geq$ 29°C (100)	-0.498***	-0.549***	-0.491***	-0.542***	-0.493***	-0.561***	-0.506***	-0.549***	-0.506***	-0.564***
	(0.051)	(0.059)	(0.053)	(0.061)	(0.053)	(0.065)	(0.058)	(0.063)	(0.058)	(0.066)
Precipitation (m)	1.077***	1.051***	1.089***	1.063***	1.069***	1.032***	1.030***	1.012***	1.030***	0.997***
	(0.191)	(0.197)	(0.183)	(0.188)	(0.180)	(0.195)	(0.184)	(0.192)	(0.184)	(0.197)
Precipitation (m) Squared	-0.870***	-0.879***	-0.866***	-0.877***	-0.856***	-0.864***	-0.834***	-0.850***	-0.834***	-0.845***
	(0.151)	(0.182)	(0.144)	(0.173)	(0.144)	(0.181)	(0.148)	(0.178)	(0.148)	(0.182)
R-squared	0.2861	0.3575	0.2617	0.3337	0.2623	0.3241	0.2490	0.3215	0.2490	0.3154
Opt. Threshold (R^2/AIC/BIC)	66	71	65	69	66	78	57	67	57	78
Observations	40552	40552	40552	40552	40552	40552	40552	40552	40552	40552
Counties	1747	1747	1747	1747	1747	1747	1747	1747	1747	1747
Years	25	25	25	25	25	25	25	25	25	25
Common Trends	-	-	-	-	-	-	2	2	1	1
State-level Trends	-	-	2	2	1	1	-	-	-	-
County-level Trends	2	2	-	-	-	-	-	-	-	-
Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No

NOTE: Column (1a) is the same as column (6) of Table 3, while other columns vary the time controls. The last three rows of the Table give the highest order polynomial in the state-specific or county-specific time trends and whether year fixed effects are included. Columns (a) include year fixed effects, while columns (b) do not. Columns (1a-1b) vary the highest order polynomial in the county-specific time trends, while columns (2a-3b) vary the highest order polynomial in the state-specific time trends, and columns (4a-5b) vary the highest polynomial in the common time trend. All specifications include county fixed effects. R-squared does not include the variation that is explained by the county fixed effects or temporal controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Effect of Ozone on Log Corn Yields - Measurement Error for Subsets of Counties

	Baseline	Weighted	By Distance to Monitor			Counties with Monitor		IV Regression	
	(1)	(2)	(3a)	(3b)	(3c)	(4a)	(4b)	(5a)	(5b)
O_3 Exp. $\geq$ 65ppb (10000 ppb-hrs)	-0.737***	-0.717***	-0.623***	-0.605***	-0.623***	-0.571***	-0.337***	-0.590***	-0.447***
	(0.088)	(0.128)	(0.086)	(0.124)	(0.086)	(0.083)	(0.052)	(0.075)	(0.061)
$\times$ Large Dist. to Monitor					0.019				
					(0.123)				
Degree Days 10-29°C (1000)	0.303***	0.371***	0.255***	0.270**	0.255***	0.260***	0.247**	0.261***	0.252**
	(0.106)	(0.114)	(0.085)	(0.131)	(0.085)	(0.101)	(0.107)	(0.101)	(0.105)
$\times$ Large Dist. to Monitor					0.015				
					(0.107)				
Degree Days $\geq$ 29°C (100)	-0.498***	-0.547***	-0.595***	-0.466***	-0.595***	-0.576***	-0.597***	-0.574***	-0.584***
	(0.051)	(0.076)	(0.068)	(0.050)	(0.068)	(0.072)	(0.076)	(0.072)	(0.074)
$\times$ Large Dist. to Monitor					0.129**				
					(0.058)				
Precipitation (m)	1.077***	1.234***	1.161***	1.083***	1.161***	0.808***	0.837***	0.805***	0.815***
	(0.191)	(0.257)	(0.215)	(0.199)	(0.215)	(0.225)	(0.241)	(0.221)	(0.228)
$\times$ Large Dist. to Monitor					-0.078				
					(0.263)				
Precipitation (m) Squared	-0.870***	-1.102***	-0.874***	-0.918***	-0.874***	-0.617***	-0.622***	-0.616***	-0.617***
	(0.151)	(0.210)	(0.146)	(0.165)	(0.146)	(0.173)	(0.185)	(0.171)	(0.176)
$\times$ Large Dist. to Monitor					-0.044				
					(0.186)				
R-squared	0.2861	0.2731	0.3356	0.2481	0.2914	0.3369	0.3288	0.3370	0.3265
Opt. Threshold (R^2/AIC/BIC)	66	67	69	55	66	70	69	70	68
Observations	40552	40552	20388	20164	40552	8032	8032	8032	8032
Counties	1747	1747	874	873	1747	474	474	474	474
Years	25	25	25	25	25	25	25	25	25

NOTE: Column (1) is the same as column (6) of Table 3. Column (2) weights by the average corn growing area in a county in 1990-2014. Columns (3a)-(3c) separate counties by the distance to the closest monitor: columns (3a) and (3b) estimate separate equations for counties with distances below and above the median, respectively, while column (3c) includes an interaction term whether the distance to the closest monitor is above the median for all variables and time controls. Column (4a) replicates column (1) for the subset of counties that have a monitor, while column (4b) uses the simple average of all monitor readings (no spatial interpolation) in those counties. Finally, column (5a) instruments (4a) with (4b), while column (5b) instruments (4b) with (4a). All specifications include county fixed effects, year fixed effects, and county-specific quadratic time trends. R-squared does not include the variation that is explained by the county fixed effects or time controls. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Cross Validation: Interpolation of Pollution at Monitor Sites

	(1a)	(1b)	(1c)	(2a)
	Panel A: Monitor Fixed Effects			
O_3 Exp. $\geq$ 65ppb (ppb-hrs)	0.665***	0.667***	0.665***	0.732***
	(0.012)	(0.013)	(0.013)	(0.014)
R-squared	0.6982	0.7143	0.7168	0.7359
Resid. σ - Observed Pollution	41.73	41.65	41.95	10.52
Resid. σ - Interpolated Pollution	33.22	32.85	32.96	8.98
	Panel B: Monitor F.E. + Time Controls			
O_3 Exp. $\geq$ 65ppb (ppb-hrs)	0.657***	0.658***	0.657***	0.482***
	(0.012)	(0.013)	(0.013)	(0.024)
R-squared	0.6916	0.7071	0.7095	0.4935
Resid. σ - Observed Pollution	40.95	40.81	41.08	6.79
Resid. σ - Interpolated Pollution	32.37	31.93	32.02	4.66
Observations	2674899	1849158	1817197	17341
Monitors	1395	936	936	1253
Observed Pollution - Mean	13.45	13.66	13.85	13.45
Interpolated Pollution - Mean	14.20	14.12	14.21	14.20
Temporal Aggregation	Daily	Daily	Daily	Annual

NOTE: Table regresses interpolated pollution values (interpolated to the location of the monitors without using the monitor itself) on observed values at the monitor location. Columns vary by the temporal aggregation. Columns (1a)-(1c) use daily values for March-August, while column (2a) use the annual sum of the daily values in (1a)-(1c). Columns (a) use all monitors east of the 100 degree meridian (except Florida), while column (1b) furthermore only uses monitors that lie in a county for which we have corn yields for at least half the years (13 out of 25 years). Finally, column (1c) limits the daily data further by excluding days with missing values at the monitor that had to be filled in. Panel A regresses interpolated values on observed station values, while panel B also includes monitor-specific quadratic time trends and year fixed effects. The footer of each panel reports the residual standard deviation (σ) from regressing observed and interpolated values on the controls of each panel, i.e., the remaining variation that is not absorbed by fixed effects that we are using in our identification. Robust standard errors in parentheses, adjusted for clustering at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A1 Online Appendix: Additional Figures

Figure A1: Spatial Distribution of Baseline Pollution Levels

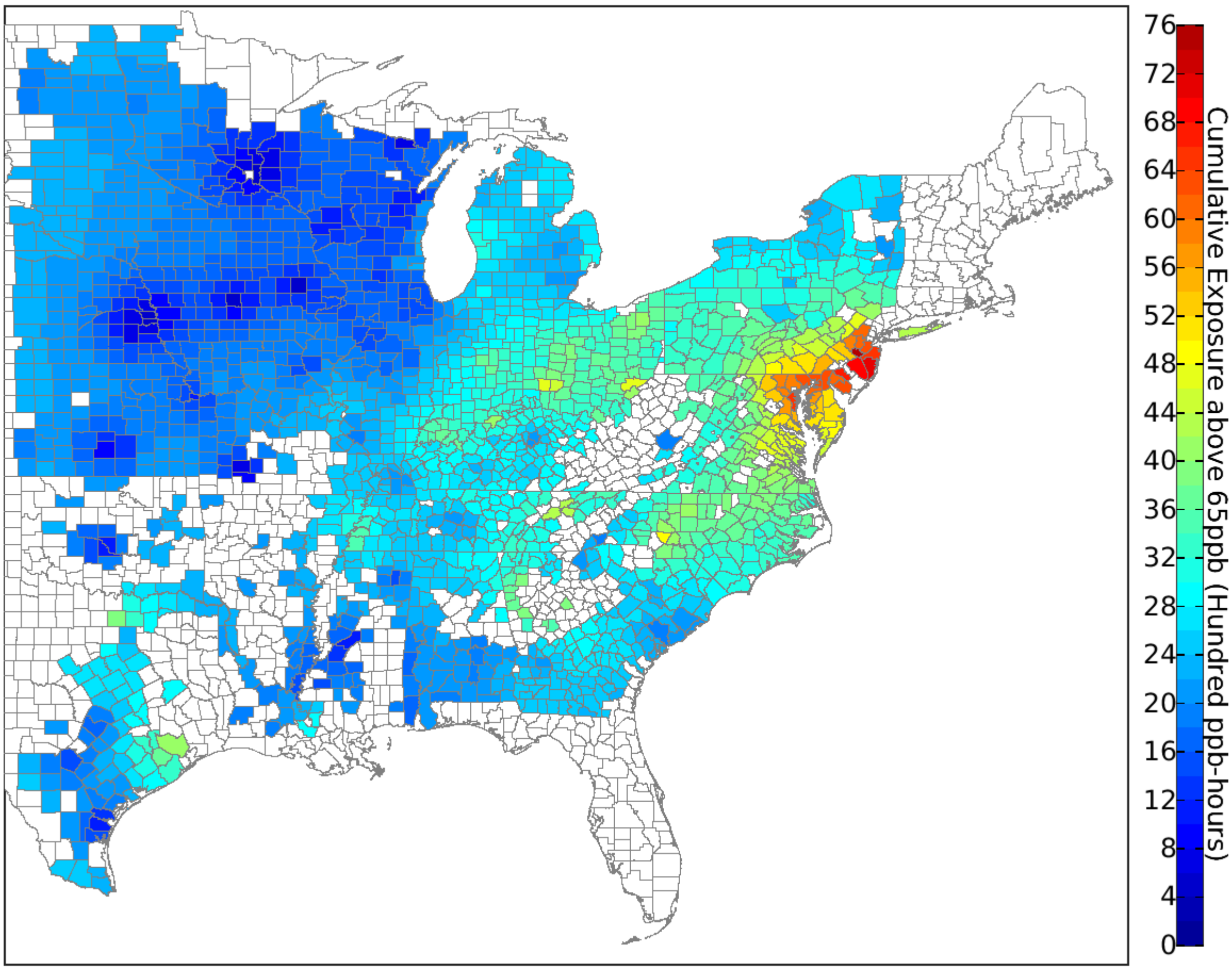

NOTE: Figure displays the spatial distribution of the average annual cumulative ozone exposure above 65ppb over the growing season (March-August) in the first five years of our panel (1990-1994).