AGRICULTURAL ECONOMICS

Agricultural Economics 47 (2016) supplement 159–171

Recent weather fluctuations and agricultural yields: implications for climate change

Anthony Louis D'Agostino, Wolfram Schlenker*

School of International and Public Affairs, Columbia University, 420 West 118th Street, New York, NY 10027

Received 17 February 2016; accepted 3 September 2016

Abstract

We summarize recent statistical analyses that link agricultural yields to weather fluctuations. Similar to other sectors, high temperatures play a crucial role in predicting outcomes. Climate change is predicted to significantly increase high temperatures and thereby reduce yields. How good are such models at predicting future outcomes? We show that a statistical model estimated using historic US data on corn and soybean yields from 1950 to 2011 is very capable of predicting aggregate US yields for the years 2012–2015, where 2012 was much hotter than normal and is expected to become the new normal under climate change. We conclude by discussing recent studies on the implication of predicted yield declines with a special focus on adaptation and commodity prices.

JEL classifications: Q10, Q54

Keywords: Agriculture; Climate change; Climate impacts

It was only 10,000 years ago, a small span in our history, that humans transitioned from hunter-gatherers to settle and become farmers (Balter 2013). Agrarian settlements enabled a more reliable and abundant food supply and employed the predominant share of the labor force in agriculture for many centuries to come. The start of the industrial revolution led to a shift of labor away from agriculture, but it was not until the 21st century that a larger fraction of the global population began living in urban areas rather than rural areas (United Nations Population Fund 2007). Consequently, throughout almost all of human history, weather has played a crucial role in shaping livelihoods given its importance in agricultural production. Negative weather shocks and long-run climatic variability, such as the Little Ice Age, have had significant effects on civilizational upheaval (Büntgen et al., 2011) and depopulation (Zhang et al., 2007), especially in subsistence economies. More recently, this weather-food production channel has been applied to contemporary episodes of unrest such as in Syria (Kelley et al., 2015), and to civil conflict more broadly (Burke et al., 2009; Hsiang et al., 2011).

Investments in input technology associated with the Green Revolution vastly improved food security around the globe, leading to inflation-adjusted commodity prices trending downward until recently. These recent technological advancements have not insulated agriculture from negative effects. To the contrary, several studies find that modern varieties achieve higher average yields at the cost of larger sensitivity to fluctuations in high temperatures (Tack et al., 2015). Climate change poses the threat of reversing gains in average yields while also increasing the volatility of global food production. Such concerns have prompted agronomists and statistical researchers to simulate the effects of climate change on agricultural yields.

The Agricultural Model Intercomparison and Improvement Project (AgMIP), an international community of crop and agricultural trade modelers, is one such effort. Similar to the climate modeling community's Coupled Model Intercomparison Project (CMIP), models are evaluated by using common inputs (e.g., reference scenarios). The multimodel mean has been found to perform better than any individual model (Asseng et al., 2015). At the same time, reduced-form statistical analyses use the fact that weather anomalies are random and plausibly exogenous and hence ideal right-hand side variables in a regression equation. Such analyses have uncovered the very potent effect of high temperature exposure, not just in agriculture (Schlenker and Roberts 2009), but also for energy use (Miller et al., 2008), mortality rates (Deschênes and Greenstone 2011), labor supply (Graff Zivin and Neidell 2014), and economic growth (Burke et al., 2015).

*Corresponding Author. Tel.: (212) 854-1806; fax: (212) 854-5765. *E-mail address*: wolfram.schlenker@columbia.edu (W. Schlenker).

 DOI: 10.1111/agec.12315

In this article, we synthesize the current state of climate impact analyses of agricultural yields and describe the various tools being used to generate forecasts of how agriculture may be affected by global climate change. We place special emphasis on results derived from statistical methods for the United States, which produces 23% of the global calories consumed directly or indirectly (feedstock) from the four staple commodities corn, wheat, rice, and soybeans. These four staples account for 75% of the calories humans consume and hence any effect on US production has global repercussions given its market size. While it is a daunting task to simulate outcomes far into the future, we show that a statistical model for the United States that was estimated using historic data from 1950 to 2011 is very capable of predicting future outcomes in 2012–2015, including the heat wave of 2012 that is predicted to become the new norm under climate change.

1. Estimating crop yield responses

Agronomists have studied how various factors impact crop yields for a long time. A majority of agronomic crop models emphasize the detrimental effects of drought conditions (Passioura 1994). Some models allow temperature exposure exceeding a crop's cycle-specific thresholds to exert large, negative effects on yields (Paulsen 1994). Many of these factors are correlated, and it is difficult to identify them in isolation. Both temperature and precipitation influence the water balance of a plant. Drought conditions imply a lack of precipitation relative to the given temperature. Heat reduces soil moisture through two channels: first, on the supply side, evaporation directly dries out the soil and transpiration depletes soil moisture through root water uptake that gets lost by the plant. Second, on the demand side, higher temperatures increase the water demand of a plant to keep up photosynthesis (Lobell et al., 2013). Precipitation, on the other hand, only affects the supply of water by replenishing soil moisture.

Agronomic models possess several strengths: they allow for more complex interactions between various inputs (soil quality, nutrients, water availability, temperature, and precipitation) under controlled growing conditions. If the goal of the analysis is to predict yields in a particular plot, these intricate interactions are of great importance. On the other hand, the number of parameters is often so large that they cannot be jointly estimated and instead need to be calibrated. Moreover, real-world field conditions might differ from experimental setups and hence be constrained by other inputs. Passioura (1994) mentions that a "difficulty is that much of the literature deals with short-term responses to water status that may be transient in that they are eventually overridden by other changes [. . .]."

The advantage of statistical models, on the other hand, is that they often offer more degrees of freedom as they can pool observations from several locations over various years, which allows them to disentangle factors that are closely correlated. While statistical analysis includes fewer variables than crop models, they rely on the fact that weather anomalies are random and hence should be uncorrelated to other factors. The omission of other factors therefore will not bias the coefficient on the variable of interest. This implies that statistical models are better at predicting yields over larger geographic areas where other factors average out rather than an individual field. For example, Lobell and Burke (2010) compare statistical models to agronomic models of CERES-Maize and find the former outperform especially at spatial scales of higher aggregation. When examining the effects of climate change on food security and food prices, a focus on aggregate outcomes is sufficient. The next four subsections will introduce results from statistical models that show that degree days are good explanatory variables, illustrate how degree days can be calculated from daily minimum and maximum temperature, discuss limitations of statistical models, and assess their predictive power.

Both agronomic models and statistical models use data from various geographic scales, ranging from field trials to more aggregate self-reported county or country averages. The majority of agronomic studies are calibrated at "representative" sites around the world and then extrapolated to other comparable areas. The advantage of using field trials is that they often have more detailed data available for a myriad of outcomes. Statistical analysis, on the other hand, generally use more aggregate data, which reflect real-world conditions that might be very different from controlled growing conditions in labs or trail sites. There are some statistical analysis that use data from field trials (Lobell et al., 2011; Tack et al., 2015; Welch et al., 2010).

1.1. Evidence from statistical models

There is a growing literature of statistical studies linking agricultural outcomes to weather anomalies.[1] Panel models employing spatial fixed effects rely on variation in anomalies, while cross-sectional studies utilize variation in climate. From a statistical perspective, panel variation benefits from exogenous and random (except for some large-scale phenomena that make weather predictable, like El Niño) weather anomalies. A location's climatology, on the other hand, is correlated with other factors that also affect agricultural output, such as soil quality. If these other confounders are not correctly accounted for, a weather variable of interest would covary with the error term, resulting in a biased estimate. For a more detailed discussion of the advantages and challenges of using weather anomalies, see Schlenker (2016).

Statistical studies rely on historical data to estimate a response function and use it to predict outcomes under both current and future weather. Recent panel studies have shown the importance of weather extremes. Schlenker and Roberts (2009)

[1] Anomalies are computed as the deviation of an observed temperature from that location's climatology, which is typically constructed as the average over a 30 year or longer period.

use data from 1950 to 2005 to estimate a fixed effects panel model of county-level, maize, soybean, and cotton yields data in the eastern United States. The optimal temperature is around 29°C for corn, 30°C for soybeans, and 32°C for cotton. For all three crops, the relationship is highly asymmetric: being below or above the optimal temperature is suboptimal in an approximately linear fashion, but being above the threshold is roughly 10 times as bad as being the same amount below the threshold. In other words, the slope of the decline above the optimum is roughly 10 times as steep as the slope of the incline below it. This piecewise linear response function is well captured by degree days, which sum up how much temperatures exceed a threshold for a chosen period of time.[2] Climate change is predicted to reduce yields as the gains from shifting below-optimal temperatures toward optimal temperatures is more than offset by the losses from shifting optimal temperatures to warmer-than-optimal temperatures. The dominating effect of temperatures above the threshold has been confirmed in other agricultural studies and other sectors as mentioned in the previous section.

Similar relationships are observed in other parts of the world. Lobell et al. (2011) analyze field-level data from sub-Saharan African maize trials conducted by the International Wheat and Maize Improvement Center (CIMMYT). Measuring temperature exposure using a cumulated degree day approach, their results support those from Schlenker and Roberts (2009) that exposure exceeding 30°C harms maize yields. They conclude that cooler areas may benefit from additional warming, but that much of the study region already experiences temperatures in excess of the upper threshold. Drought conditions further amplify damages.

A few technical details of recent panel studies are worth repeating. First, previous studies have sometimes used quadratic specifications to model nonlinearities. This is very restrictive as a quadratic function assumes symmetry around the optimum, yet the relationship might be highly asymmetric as outlined above. It is hence preferable to use binned indicator variables for various interval ranges of the weather outcomes, or restricted cubic splines to first examine whether the response function can be approximated by a piecewise linear function underlying the concept of degree days (Schlenker 2016).

Second, nonlinearities can only be detected if the weather data are fine-scaled enough in both time and space. Section 1.4 illustrates how averaging weather variables over time and space reduces the predictive power of the model. For example, Schlenker and Roberts (2009) construct very fine (2.5 miles × 2.5 miles) gridded daily temperature and precipitation data. Since yield outcomes are reported by county, weather data are aggregated to the county-level using a satellite scan of the cropland area. However, the sequence in which aggregation and temperature transformation is performed matters. Initially deriving the nonlinear temperature transformation (described in Section 1.2) for each grid pixel and then averaging over the county gives different results compared to an approach that applies the nonlinear temperature transformation to the county-averaged weather outcome. If spatially detailed data are available, it should be transformed and then averaged, so as not to smooth over the true observed extremes. A simple example might illustrate this point: consider a county of two pixels each experiencing temperature anomalies of equal magnitude around the threshold but of opposite signs. Taking the spatial average and then performing the nonlinear transformation would lead to the false conclusion that the entire county experienced no temperatures above the threshold, since the average does not exceed the threshold.

A similar logic applies to temporal averaging; employing a sinusoidal interpolation between the daily minimum and maximum temperature (Snyder 1985) gives a different outcome for the day than applying the transformation to the daily average. This is especially important for extremes, both hot and cold, as the daily average often does not pass thresholds that the daily minimum or maximum do pass. For example, a nonlinear data generating process with threshold behavior initiated at 30°C will generate different outcomes across the following two scenarios: one with 12 hours of 33°C exposure followed by 12 hours of 23°C, and an alternative of constant 28°C for 24 hours. Both scenarios face a daily average temperature of 28°C, but only the former would obtain a positive degree days measure above 30°C.

Data analyses which smooth out extreme exposures by either spatial or temporal aggregation will underperform in estimating the response function. This finding is confirmed in Tack et al. (2015) who use detailed data on Kansas wheat trials and pair it with daily weather station data; both extreme cold (freezes) as well as hot temperatures negatively impact yields. The latter effect dominates and wheat yields are predicted to decline under climate change. The authors show that a model that captures the within-day distribution of temperatures between the daily minimum and maximum does better in terms of R^2 and out-of-sample forecasts than a model using the daily average.

Previous analyses assumed that the effects of temperatures are additively separable throughout the growing season. Welch et al. (2010) modify this approach when estimating farm-level rice yields data generated from sites in six rice-producing Asian countries. They divide the rice season into vegetative, reproductive, and ripening phases and estimate significantly detrimental effects from higher minimum temperatures in both vegetative and ripening phases, while increases in vegetative phase maximum temperature are significantly beneficial. As an alternative to the cumulated degree day or degree-bin exposure time methods, the authors explicitly include daily minimum and maximum temperatures for the growing phases as separate variables. This allows for coefficients of different magnitudes or signs. They also consider diminishing solar radiation levels which have been attributed to rising aerosol concentrations

[2] For example, a sustained temperature of 32°C for 24 hours would give 3 degree days above a 29°C threshold. The total over a growing season is simply the sum from repeating this approach for each day of the growing season.

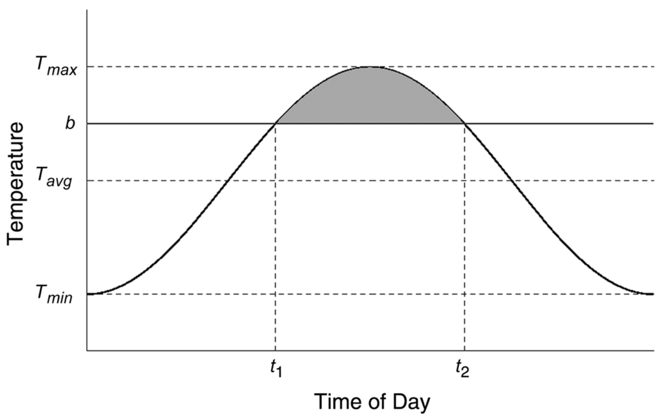

Notes: Degree days are the integral (gray area) between the temperature distribution within a day and the bound, b, above which they are measured. The daily distribution of temperatures are approximated by a sinusoidal curve between the daily minimum and maximum temperature. Temperatures exceed the threshold between times t_1 and t_2, which are formally derived in the online appendix.

Fig. 1. Deriving degree days from daily minimum and maximum temperature.

and atmospheric brown clouds (Crutzen and Ramanathan 2003; Ramanathan and Feng 2009). Aerosol concentrations are expected to decline as regional emissions standards tighten and economies transition from bio-based fuels and coal to cleaner energy systems. As a result, the negative effects of reduced radiation on the ripening phase of rice may be reversed. However, because of radiative backscattering, atmospheric aerosols are likely masking warming trends that would be larger in the absence of aerosols, suggesting that an aerosol-greenhouse gas joint management strategy must be devised to avoid rapid warming resulting from a rapid aerosol emissions reduction.

1.2. Construction of degree days

Some weather stations report hourly observations and hence have data on the within-day distribution. Even if only the daily minimum and maximum are known, a researcher can approximate the within-day distribution by a sinusoidal distribution, as done by Schlenker and Roberts (2009) and Tack et al. (2015), both following Snyder (1985). Degree days above a threshold (sometimes called cooling degree days in the energy literature), are simply the area under the temperature curve above the threshold throughout the day and as shown in Fig. 1. The mathematical derivation of degree days is given in the online appendix for both cooling and heating degree days for an arbitrary bound b.[3] Using $\bar{t} = acos(\frac{2b-T_{max}-T_{min}}{T_{max}-T_{min}})$, we get

cooling degree days

$$= \begin{cases} \dfrac{T_{max}+T_{min}}{2} - b & \text{if } b \leq T_{min} \\ \dfrac{\bar{t}}{\pi}\left[\dfrac{T_{max}+T_{min}}{2} - b\right] & \\ \quad + \dfrac{T_{max}-T_{min}}{2\pi}\sin(\bar{t}) & \text{if } T_{min} < b < T_{max} \\ 0 & \text{if } T_{max} \leq b, \end{cases}$$

heating degree days

$$= \begin{cases} 0 & \text{if } b \leq T_{min} \\ \left[1 - \dfrac{\bar{t}}{\pi}\right]\left[b - \dfrac{T_{max}+T_{min}}{2}\right] & \\ \quad + \dfrac{T_{max}-T_{min}}{2\pi}\sin(\bar{t}) & \text{if } T_{min} < b < T_{max} \\ b - \dfrac{T_{max}+T_{min}}{2} & \text{if } T_{max} \leq b. \end{cases}$$

These piecewise linear temperature variables are shown to perform exceptionally well in predicting agricultural yields in Section 1.4.

1.3. Limitations

There are several limitations of the standard panel regression method. First, the effects of carbon fertilization cannot be included because nominal spatial heterogeneity implies that any observed changes in CO_2 concentrations over the study period are absorbed by time trends or time fixed effects. The predicted effects of a change in climate conditions using a statistical panel model does not include CO_2 fertilization effects, which differ between C_3 (e.g., soybeans, rice, and wheat) and C_4 crops (e.g., corn). The former class of crops has larger yield gains from increases in CO_2.

While panel data use fixed effects to capture time-invariant omitted variables, there might still be time-varying omitted variables. Sheehy et al. (2006) emphasize the problems of omitted variable bias, noting strong colinearity between temperature and solar radiation, which biases temperature estimates if the latter is omitted. In case the correlation between temperature and solar radiation remains unchanged in the future, the coefficient would still give unbiased climate change impacts as the temperature coefficient simply picks up the joint effect of both changes in temperature and solar radiation. However, most climate models predict increases in temperature, but not necessarily of solar radiation, which depends on local pollutants and how much solar radiation they absorb. In such a case, a temperature coefficient that includes the beneficial effects of solar radiation will not give the correct climate impact as future changes in temperature and solar radiation diverge from historic patterns used to identify the model parameters.

Surface temperature maintains a high degree of spatial covariance, whereas precipitation is significantly more

[3] Tack et al. (2015) further refine this approach by linking a day's maximum temperature with the next day's minimum temperature.

heterogeneous in space. The cross-validation exercise of Schlenker and Roberts (2009), where data are interpolated to the location of a weather station and compared to the actual outcome, finds that spatial interpolation of temperature is much better than precipitation interpolation. This implies that gridded precipitation products by their nature consist of a higher noise ratio than temperature products. This is especially troublesome in a panel setup that relies on anomalies. While most datasets agree on which locations are wetter on average, there is much less agreement on whether a particular year was above or below average (Auffhammer et al., 2013). This could result in attenuation bias on the precipitation variables. Burke and Emerick (2013) find that the trend in yield growth is more sensitive to the observed precipitation trends than the coefficient that is observed in a panel setting. This might either indicate attenuation bias in the panel setting or the fact that applying water is feasible to annual weather shocks, but not sustainable in the longer term.

While degree days based on fine-scaled weather data has performed well in the United States, many developing countries have a much coarser station network that requires a lot of interpolation to fill in the gaps. In such a case, researchers must use either reanalysis products or more aggregate temperature data from spatial interpolation routines, i.e., monthly instead of daily data.[4] Schlenker and Lobell (2010) analyze yields for five key crops in sub-Saharan Africa and run four model specifications including both average growing season temperature as well as degree days, yet find that all approaches give rather comparable results. While their results vindicate the utility of even crude temperature data, weather data quality issues prohibit more extensive models which could plausibly provide policymakers richer guidance on mitigating climate losses.

Finally, some authors have criticized statistical studies for interpolating past behavior into the future, which makes the implicit assumption that technology is time-invariant. However, Burke and Emerick (2013) find that the trend in yields shows the same sensitivity to trends in hot degree days to what is observed in a panel setup that relies on weather shocks, i.e., the long-run response is comparable to the short-run response. Still, whether modern crops are just as sensitive to hot degree days as what was observed in historical data is examined in more detail in the next subsection.

1.4. Recent US yield forecasts

Previous studies have found hot degree days to be the most significant predictor of annual fluctuations in corn and soybean yields. This section tests whether models estimated solely with previous data (1950–2011) have predictive power for the past four years (2012–2015). The weather data were constructed in the same way as in Schlenker and Roberts (2009). We first derive the daily minimum and maximum temperature for each 2.5 $\times$ 2.5 mile PRISM grid cell, construct the nonlinear temperature transformation for each grid cell and day, and then sum the data over all days of the growing season and average it over the cropland area in each grid cell to obtain county aggregates. This is the basis for the county-level analysis. We pair it with annual county-level yields as reported by the National Agricultural Statistics Service (NASS) for the eastern United States, specifically all counties east of 100° longitude except Florida. The panel specification is the same as the piecewise-linear function (red lines) in fig. 1 of Schlenker and Roberts (2009).[5]

Since we are estimating a linear model, we should be able to aggregate both dependent and independent variables to higher spatial aggregation levels, i.e., annual-level yield data[6] by weighting the county data by predicted production, which is simply the actual growing area times predicted yields according to state-specific time trends.[7] The annual aggregate has only one observation per year, while the county level data has 2,276 observations per year for corn and 2,079 for soybeans. To highlight the importance of first conducting the nonlinear transformation and the summing over space/averaging over time, we also produce forecasts that first aggregate the weather data across space or time and then apply the nonlinear transformation.

Figure 2 displays the cumulative distribution of aggregate exposure to hot degree days (degree days above 29°C or 84°F) in the left graph. It is the weighted average of the county-level data, where the weights are predicted production along a trend as described in the previous paragraph. It shows the combined exposure that the growing area has experienced since April 1 of each year. The black solid line shows the historic average for the years 1950–2011, while the dashed gray lines show the distribution for individual years in 1950–2011. The colored lines show the results for the last four years 2012–2015. The blue line depicts 2012, which had a heat wave in July. As a consequence, the cumulative exposure to hot degree days rises rapidly in July when temperatures exceeded 29°C frequently. There was less heat above the threshold in August 2012 and the line tapers off. The year 1988 had the highest season total number of hot degree days as shown by the gray line eclipsing the blue line by mid-August.

The right graph of Fig. 2 shows the cumulative distribution of season-total precipitation, where the color-coding of the lines is equivalent to the left graph. Note that 2012 was not only one of the hottest years on record in terms of hot degree days,

[4] Weather data using statistical interpolation techniques are contrasted with reanalysis products that combine interpolation with physical laws on the preservation on mass and energy in more detail in Auffhammer et al. (2013).

[5] The only difference is that we cluster errors by state instead of using Conley's routine on spatial correlation. This has no effect on the point estimate but might give different standard errors. Previous research has shown that both give comparable standard errors, and clustering is much more computationally efficient.

[6] To keep the model consistent with Schlenker and Roberts (2009), we only use counties east of the 100° meridian excluding Florida in the analysis. The aggregate is the sum of those counties, which account for the dominant share of US production.

[7] We chose state-specific restricted cubic spline with the knots at 1960, 1980, and 2000. This forces linear trends above 2000 into the future.

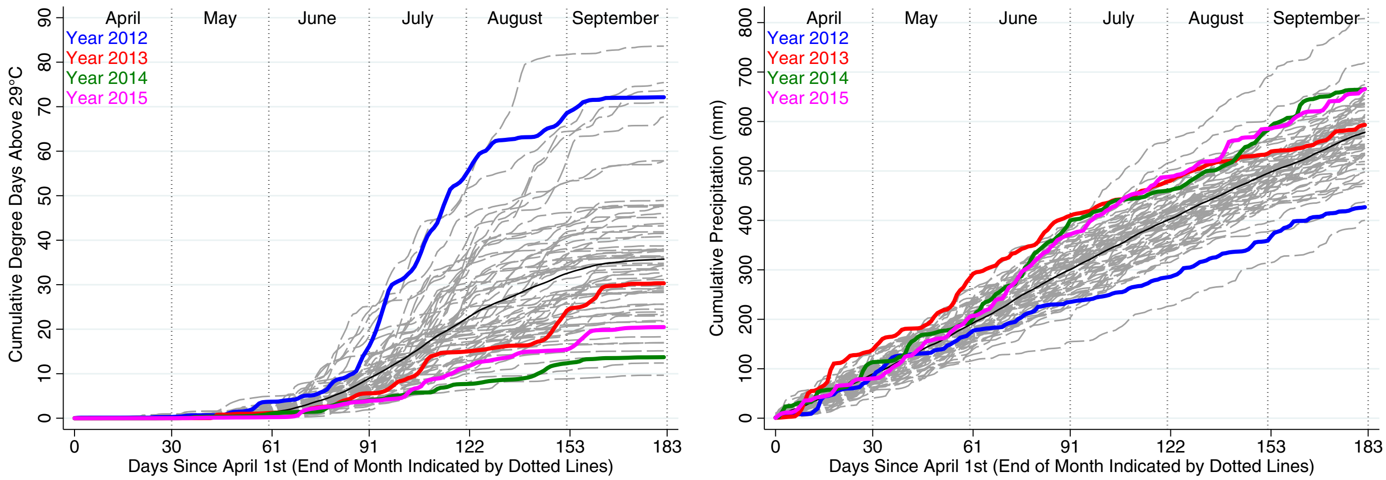

Notes: Figures display observed cumulative exposure to hot degree days (degree days above 29°C) in the left graph and precipitation in the right graph over the April–September growing season. Weather variables are weighted averages of the county-level weather outcomes, where the weights are observed growing area times predicted yields according to a trend. The black solid line is the historic average for 1950–2011, while the dashed gray lines show individual years from 1950 through 2011. The last four years are shown in color.

Fig. 2. Cumulative exposure to degree days above 29°C and precipitation.

but also one of the driest as the blue line comes in as the second-lowest by the end of September—it was second to 1988, which was the driest year on record. Recall that degree days above 29°C only measure temperatures above 29°C. A large number of hot degree days does not necessarily imply a large average temperature, since temperature fluctuations below 29°C are disregarded.

One can also easily see why 2014 and 2015 produced very strong yields. There was a limited amount of hot degree days and total rainfall was higher than usual, as shown by the lines in green and pink. The negative correlation between hot degree days and precipitation does not always hold, especially for county-level values: some counties in the very hot year 2012 experienced above average rainfall. In general, both for historic time series and for future predictions of climate change, aggregate impacts mask great heterogeneity among counties.

How well do these simple variables capture year-to-year yield variation? Figure 3 shows the results of a regression of log corn yields (left) and log soybean yields (right) on the four weather variables of Schlenker and Roberts (2009): moderate as well as hot degree days and a quadratic in season-total precipitation.[8] For corn, moderate degree days are between 10°C and 29°C, while hot degree days are above 29°C.[9] We also include quadratic time trends.[10] These four weather variable captures a very large portion of the year-to-year fluctuations in yields. The largest portion is due to the measure of hot degree days. The model predicts well out-of-sample for the years 2012–2015. While some industry representatives have suggested that modern crops are much better at withstanding heat than historic varieties used in estimating the response surface, the model prediction for 2012 does not seem to overpredict the effects of the observed heat wave for corn, and only slightly so for soybeans.

Table 1 shows predicted aggregate corn yields for the years 2012–2015, while Table 2 shows the results for soybeans. The model is estimated using data for 1950–2011, analogous to the specifications of Fig. 3. Column (1) provides the reported average yield for the counties in the sample, i.e., the eastern United States.[11] Columns (2a)–(3d) give the prediction errors in *percent* from various models, while rows vary the year of the out-of-sample forecast. Standard errors on the predictions are given in parentheses. Columns (a) do not use any weather variable, i.e., they simply predict yields along the trend. Columns (b) and (c) use one explanatory variable: average temperature over the growing season and hot degree days, respectively. Columns (d) use the four weather variables of Schlenker and Roberts (2009). Columns (2a)–(2d) run the regression and out-of-sample prediction at the county-level. County-level logyield forecasts are translated into aggregate forecasts by multiplying predicted yields[12] times the observed growing area. We then sum predicted production and area over all counties and derive predicted yields as the ratio of the two. Aggregate prediction error is simply the percent difference from the actually observed yield for the year. Standard errors are obtained from

[8] The regression in Schlenker and Roberts (2009) use monthly precipitation totals from PRISM to construct season-total precipitation measures. These numbers were not yet available for 2015 at the time this article was written. We hence use the sum of the daily precipitation measure of the interpolation routine described in the article.

[9] For soybeans, the bounds are 10–30°C and above 30°C.

[10] The county-level and state-level regressions have quadratic time trends that are allowed to differ by state.

[11] Average yields are obtained by summing all production and harvest area for counties east of the 100° meridian except Florida and then taking the ratio.

[12] Predicted yields are $e^{\log(\text{yield})} \times e^{\frac{\sigma^2}{2}}$ to account for the convexity of the exponential function.

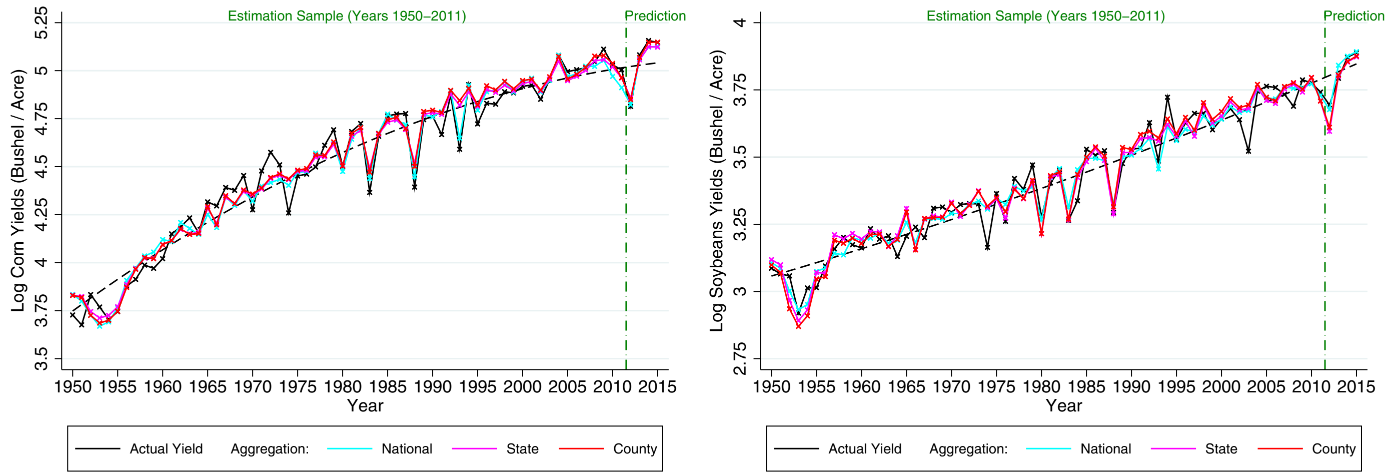

Notes: Figures display observed log corn yields (left) and observed log soybean yields (right) in black. Three statistical models were estimated using data from 1950 to 2011, and yields were predicted both in sample (1950–2011) as well as out-of-sample (2012–2015). All models included four weather variables: a quadratic in season-total (April–September) precipitation as well as two degree variables where the bounds vary by crop (10–29°C and above 29°C for corn, 10–30°C and above 30°C for soybeans). The statistical models differ by spatial aggregation. The cyan line uses annual aggregate data, i.e., 62 observations, with a quadratic time trend. The magenta and red line use state and county-level yields, respectively, each paired with a quadratic time trend by state. All weather variables are first derived for a 2.5 × 2.5 mile grid on each day and then aggregated to the eastern United States.

Fig. 3. Statistical models predicting log corn and soybean yields 1950–2015.

Table 1
Out-of-sample forecasts for corn yields 2012–2015

		Prediction error (%) under various models							
	Actual yield	County-level				Aggregate eastern US			
	(1)	(2a)	(2b)	(2c)	(2d)	(3a)	(3b)	(3c)	(3d)
Yield in 2012	123.16	27.19	14.89	4.38	3.62	23.41	18.28	4.29	1.16
			(1.41)	(0.60)	(1.91)		(3.84)	(3.73)	(4.38)
Yield in 2013	161.14	−1.91	3.13	−1.16	-1.34	−5.18	−2.98	−3.92	−1.75
			(4.19)	(1.32)	(1.27)		(3.08)	(2.46)	(2.56)
Yield in 2014	173.63	−7.85	−1.66	−0.07	−0.93	−11.58	−8.84	−3.66	−3.11
			(4.32)	(2.13)	(1.87)		(3.26)	(3.28)	(2.85)
Yield in 2015	171.87	−5.35	−7.20	−1.41	0.19	−10.30	−10.57	−5.01	−1.83
			(2.71)	(1.77)	(1.60)		(2.98)	(3.14)	(3.03)
RMSE		14.43	8.46	2.37	1.99	14.28	11.54	4.25	2.09
Weather variables		no	avg	dday	four	no	avg	dday	four

Notes: Table reports actual yields (column 1) as well as out-of sample predictions errors under various models in columns (2a)–(3d) in percent. Standard errors on the prediction errors are given in brackets and were obtained from 1,000 bootstrap runs resampled from the joint distribution of all parameters. Columns (a) include no weather variable and the model hence simply predicts yields to equal the trend. Columns (b) and (c) only include one weather variable: average temperature over the season (April–September) in columns (b) and season-total degree days above 29°C in columns (c). Columns (d) use four weather variables: season-total (April–September) degree days 10–29°C, degree days above 29°C, and a quadratic in precipitation. The statistical models also differ by spatial aggregation. Columns (2) use county-level yields with a quadratic time trend by state. Columns (3) use aggregate data by year, i.e., 62 observations, with a quadratic time trend. All weather variables are first derived for a 2.5 × 2.5 mile grid and then aggregated to the county or annual aggregate level. The models are estimated using the years 1950–2011 and predicted out-of sample for 2012–2015.

1,000 bootstrap simulations from the regression results where we repeatedly resample all parameters.

Column (2a) gives on average a forecast error of 14.4% across years for corn as shown in the second to last row of the table. Recall that the specification in (2a) simply predicts yields to equal the trend. This forecast error is reduced by not even half to 8.5% for a model using average temperature in column (2b). On the other hand, the model using only hot degree days in column (2c), which only measure temperatures above 29°C and disregards all temperature fluctuations below it, give a much smaller prediction error of 2.4%. Hot temperatures are a much better predictor of yield fluctuations than average temperatures. The model using all four weather variables reduces the prediction error slightly further to 2.0%.

Columns (3a) and (3b) replicate the analysis except that both yields and weather are first aggregated to the annual level. The model is no longer estimated using 115,205 observations of the panel, but 62 annual observations. The weather data are

Table 2
Out-of-sample forecasts for soybean yields 2012–2015

	Actual yield	Prediction error (%) under various models							
		County-level				Aggregate eastern US			
	(1)	(2a)	(2b)	(2c)	(2d)	(3a)	(3b)	(3c)	(3d)
Yield in 2012	40.20	9.28	3.23	−5.80	−7.96	10.92	8.17	1.01	−1.58
			(0.76)	(0.47)	(1.23)		(2.94)	(2.90)	(2.63)
Yield in 2013	44.40	0.22	3.73	2.68	1.19	1.73	1.89	2.15	5.12
			(3.63)	(0.52)	(0.89)		(2.53)	(2.52)	(2.12)
Yield in 2014	48.02	−6.59	−2.61	0.47	−1.57	−4.73	−4.46	−1.81	0.52
			(3.68)	(0.81)	(1.03)		(2.56)	(2.66)	(2.11)
Yield in 2015	48.82	−6.52	−7.37	−3.26	−1.21	−5.07	−6.40	−4.25	0.53
			(2.09)	(0.59)	(0.71)		(2.68)	(2.68)	(2.39)
RMSE		6.56	4.62	3.60	4.14	6.53	5.73	2.60	2.70
Weather variables		no	avg	dday	four	no	avg	dday	four

Notes: Table reports actual yields (column 1) as well as out-of sample predictions errors under various models in columns (2a)–(3d) in percent. Standard errors on the prediction errors are given in brackets and were obtained from 1,000 bootstrap runs resampled from the joint distribution of all parameters. Columns (a) include no weather variable and the model hence simply predicts yields to equal the trend. Columns (b) and (c) only include one weather variable: average temperature over the season (April–September) in columns (b) and season-total degree days above 30°C in columns (c). Columns (d) use four weather variables: season-total (April–September) degree days 10–30°C, degree days above 30°C, and a quadratic in precipitation. The statistical models also differ by spatial aggregation. Columns (2) use county-level yields with a quadratic time trend by state. Columns (3) use aggregate data by year, i.e., 62 observations, with a quadratic time trend. All weather variables are first derived for a 2.5 × 2.5 mile grid and then aggregated to the county or annual aggregate level. The models are estimated using the years 1950–2011 and predicted out-of sample for 2012–2015.

Table 3
Out-of-sample forecasts by aggregation level

	Prediction error (%) under various temp. aggregation							
	Corn				Soybeans			
	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Prediction error in 2012	1.70	−0.37	0.90	1.16	−6.84	−6.66	−3.03	−1.58
	(4.69)	(4.56)	(4.44)	(4.38)	(2.66)	(2.69)	(2.83)	(2.63)
Prediction error in 2013	4.44	-2.72	5.37	−1.75	5.85	2.93	7.34	5.12
	(3.25)	(3.08)	(3.01)	(2.56)	(2.39)	(2.30)	(2.20)	(2.12)
Prediction error in 2014	−4.17	−3.46	−3.68	−3.11	−1.22	0.20	−0.17	0.52
	(3.25)	(3.28)	(3.03)	(2.85)	(2.26)	(2.26)	(2.08)	(2.11)
Prediction error in 2015	−1.40	−5.14	−0.68	−1.83	0.36	−0.49	0.17	0.53
	(3.66)	(3.38)	(3.35)	(3.03)	(2.26)	(2.33)	(2.31)	(2.39)
RMSE	3.24	3.39	3.30	2.09	4.54	3.65	3.97	2.70
Temperature average	a/m	area	mon.	grid	a/m	area	mon.	grid

Notes: Table reports sensitivity of the aggregate (eastern United States) analysis to how degree days are constructed from temperature variables. Columns (1a)–(1d) use aggregate annual corn yields, while columns (2a)–(2d) use aggregate soybeans yields. Columns (1d) and (2d) are the same as column (3d) in Tables 1 and 2, respectively. Columns (d) use daily data for each 2.5 mile × 2.5 mile grid to construct degree days before averaging them across space and time before deriving degree days. Columns (c) average across time: they use average monthly minimum and maximum temperature instead of daily temperatures. Columns (b) average temperature across space on each day. Finally, columns (a) are the coarsest specification that average temperature across space and each month. All specification also include the same quadratic in season-total precipitation as well as quadratic time trend. The models are estimated for the years 1950–2011 (62 observations) and predicted out-of sample for 2012–2015.

averaged using production weights.[13] The results are roughly comparable, with slightly higher average prediction errors under aggregate weather measures. Recall that we are first conducting the nonlinear transformation before averaging the data. First aggregating the weather data and then taking the nonlinear transformation increases the prediction error as shown in Table 3.

[13] Production weights are constructed by multiplying actual growing area times predicted yields using state-specific time trends, as described earlier.

Table 2 replicates the same out-of-sample prediction exercise for soybeans instead of corn. Models using hot degree days in columns (c) again outperform models using average temperature in columns (b), although the reduction in error is not as big as for corn. Including all four weather variables very slightly increases the out-of-sample prediction error, but given the uncertainty with these predictions, they are not significantly different at the 95% level from one another in all but one of the eight comparisons.

Table 3 examines the role of aggregation bias for nonlinear models, which gets amplified when weather data are aggregated to the annual level. Columns (1d) and (2d) are the same as columns (2d) in Tables 1 and 2, respectively, i.e., a national-level model using four weather variables for corn and soybeans where the nonlinear transformation is first conducted for each day and grid cell and then aggregated to the annual level by summing over all days and averaging over all grids. Table 3 aggregates the weather data before taking the nonlinear transformation, where the most aggregate data are in columns (a) and the least aggregated are in columns (d). Sometimes only monthly weather data are available. Columns (c) therefore first average temperatures for each month for each grid cell before conducting the nonlinear transformation for degree days and then averaging them over all grid cells. By the same token, some weather datasets report at a much coarser grid. Columns (b) first average daily temperatures over grids in the eastern United States before conducting the nonlinear transformation for degree days and then summing over all days of the growing season. Columns (a), the model with the most aggregate weather data, first average temperatures over all grids and all days of a month before conducting the nonlinear transformation for degree days. For both corn and soybeans, conducting the nonlinear transformation before averaging over space or time (or both) in columns (d) gives the lowest prediction errors. When nonlinearities are important, averaging over space and time can dilute these nonlinearities. These should be conducted on the smallest possible grid cell and time step.

In summary, a very simple statistical model with either just one variable (measure of hot degree days over the growing season) or four weather variables (moderate and hot degree days plus a quadratic in total precipitation) give very good forecasts for the last four years that were not used in the estimation of the coefficients. This gives us some confidence that the model is adequate to simulate the effects of climate change on crop production. In the case of corn, the hot year 2012 was especially well predicted.

2. Implication of yield declines

Are predicted yield declines form statistical panel regression valid predictions of climate change? Below we discuss whether adaptation might mitigate the predicted yield declines and what effect they might have on commodity prices.

2.1. Adaptation in current growing areas

Can climate impact estimates for agriculture be reliably extrapolated from controlled field or greenhouse experiment results alone, or statistical correlations based on past behavior as outlined in the previous section? Agricultural practices are dynamic and adjustable. Economic actors will adapt to evolving conditions and if current practices lose their optimality, new practices will displace them. On the other hand, if new practices are too costly, it might not be worthwhile to engage in them. Extensive and intensive margins of adjustment will be available, and the range of options available to farmers, though situation-specific, include modifying the growing season or adjustments in crop management activities like alteration of inputs, switching of crop varieties, or crop-switching.[14] While statistical panel analyses capture within-season adaptation to weather shocks, i.e., change of some inputs like more use of irrigation water, they do not incorporate responses to permanent shifts in climate like crop switching.

Some insight into which practices may be feasible under a changed climate can be learned from the practices of farmers elsewhere who have been subject to warmer conditions in the past. This forms the basis of the Ricardian approach (Kurukulasuriya and Mendelsohn 2007; Mendelsohn et al., 1994), which employs cross-sectional techniques using farmland value in estimating warming effects on agriculture. While some such changes will doubtless occur, impediments are likely to block their complete emulation. For example, parts of the Western United States have benefited from generous federal irrigation program subsidies, which have capitalized into higher private land values. It is unlikely such subsidy levels will be extended in the future to areas not already possessing comparable irrigation systems, and even if they were, they represent government transfers and not benefits from a societal perspective. Since irrigation is observed to mitigate the effects of hot degree days, the damage function estimated from currently irrigated lands cannot be simply mapped onto areas whose future climate will approximate that of currently irrigated lands (Schlenker et al., 2005). Many adaptation measures are costly, and it is important to account for these costs or otherwise one will overestimate the potential benefits of adaptation.

Butler and Huybers (2013) run county-level degree day models of maize yields in the eastern United States and uncover large heterogeneity in hot degree days responsiveness (degree days above 29°C during the growing season, as described in the previous section). They then regress the estimated response on the location's hot degree days climatology and find that warmer counties are more heat-tolerant, suggesting that regional adaptation and use of appropriate cultivars have been effective at reducing heat-induced losses. Under a uniform 2°C warming scenario, continuation of observed adaptation would more than halve aggregate yield losses. However, they overstate the potential benefits of adaptation by assuming that it is costless; farmers can obtain a lower sensitivity to hot degree days if their area warms. If such technology was available at no cost, risk-averse farmers should already be adopting them in current climates as they reduce the variability of output at the current climate while giving the same average yield. In nonlinear panel models, the trade-off between a reduction in weather-sensitivity and average yields can be identified. Schlenker et al. (2013) find that for 2°C warming, the benefits of a lower sensitivity to hot

[14] Smit and Skinner (2002) offer a comprehensive inventory of adaptation options in agriculture.

degree days are roughly compensated by lower average yields, a result that is intuitively given by the envelope theorem, which implies that the first-order effect is given by the direct effect on yields and not the indirect effects of changes in management practices.

How closely agents can stay at the boundary of their evolving production frontiers will likely remain contested until stronger shifts in climate are realized that would allow for stronger tests of adaptation behavior. Some researchers posit that economic actors will efficiently incorporate newly revealed information about the changing climate into their decision-making, resulting in new optima minus adjustment costs. Such an assumption would project climate change costs to be substantially lower than forecasts based on extrapolation without adaptation (Kahn 2014). However, other work suggests that market distortions, discounting, and information asymmetries are some factors that may undermine optimization, or that changes in production technology are simply too costly compared to the benefits. Recent papers conclude that observed farmer adaptation thus far has been limited (Burke and Emerick 2013; Taraz 2015).

Burke and Emerick (2013) adopt a "long differences" approach on eastern US agriculture, comparing mean values for 1978–1982 agricultural outcomes against those from 1998 to 2002 to generate 20-year differences. Over this period, portions of the eastern United States experienced temperature increases on par with those anticipated over coming decades, providing some equivalence between analyzed and forecast temperature changes. The recovered sensitivities to hot degree days are not statistically different from the ones obtained in a panel regressions, suggesting that adaptation to date has not been a huge factor. While long time horizons allow for greater levels of climatic variation, agents living in the beginning state and the end state often lose comparability. This "frequency-identification trade-off" (Hsiang and Burke 2014) implies that an increasing number of confounds must be ruled out to isolate effects from the weather/climate channel.

A complementary approach involves examining the efficacy of individual actions. Taraz (2015) tests for changes in irrigation investment and crop composition among Indian farmers due to multiyear rainfall regimes, while Kala (2015) finds in her sample of Indian farmers support for ambiguity aversion—favoring planting decisions that insure against worst-outcome monsoon onset realizations—and that these decisions are consistent with a belief in climate stationarity. However, she does not control for the colinearity of temperature and precipitation and the role temperature plays in governing planting decisions via evapotranspirative effects. Colmer (2015) identifies a labor reallocation channel in Indian agriculture. In warmer and less productive years, casual laborers shift from agriculture into manufacturing. Migration is often the precursor to such occupational shifts, and in its own right has been the outcome of interest for numerous papers in development economics utilizing weather shocks as exogenous income shifters. Alternatively, Henderson et al. (2015) investigate urbanization as a more permanent response to changes in agricultural productivity. Using data from 29 sub-Saharan African countries, they estimate the effect of moisture availability, a measure of precipitation and potential evapotranspiration, on urban population and income. They find long-run drying drives urbanization, but only in locations with an existing manufacturing base.

While empirical tests of adaptation benefit from focusing on netted outcomes, translating such findings into policy requires guidance on what may work and what does not. Irrigation, for example, is considered one promising action for areas not yet at full irrigation. Evidence from more than 40 years of district-level yields data in India suggests that areas with more irrigation experience significantly lower yield losses from very high temperature, though numerous confounds including wealth, ability, and soil quality cannot be ruled out as the true causal factors. Well-designed field trials could shed light on whether such investments do lower the sensitivity to hot temperatures, but their external validity rests on how well the experiments approximate irrigation and other institutional setups. Furthermore, time horizons matter and decisions made to maximize short-run returns may be detrimental in the long run. For example, India is currently facing crises in several areas where agriculture has for years exploited highly subsidized energy to overdraw groundwater. Some of these locations are now facing salinity issues, which have led to abandonment or low crop productivity.

2.2. *Shift in growing areas*

While adapting to hot degree days in current growing areas seems to come with some challenges, another possibility might be to shift where crops are grown as climate change alters the location of the areas that offer the optimal weather. The possibility of shifts in growing areas in large parts depends on soil quality. The standard panel analysis hence runs into problems, as it is impossible to identify the spatial fixed effects for areas that currently do not grow a commodity but might in the future.

How climate change will shape the spatial pattern of agricultural production is an ongoing area of research. Reilly et al. (1994) and Rosenzweig and Parry (1994) represent some of the earliest efforts to calculate trade's role in reducing losses. Reilly et al. (1994) use a partial equilibrium model, exogenously prescribing yield effects from Rosenzweig and Parry (1994). Reilly et al. (2003) ask whether climatic change contributed to the observed westward shift in corn, soy, and wheat production, but find that mainly nonclimatic factors are responsible for it. More recently, Costinot et al. (2014) develop a "field-level" international trade model employing FAO Global Agro-Ecological Zones (GAEZ) potential yields data, both under historic and projected climates, to model each GAEZ grid cell's potential productivity of 10 key crops. This approach has the distinct advantage of using raster data with continuous coverage, which does not suffer from the problem of irregular subnational outcomes reporting that plagues many countries. In a trade-integrated global economy under climate change, climate change effects on agriculture would shave off about a quarter percentage point from global GDP.

2.3. Impacts on prices

We have thus far focused on estimating the effects of temperature on productivity per unit area (yields). From a welfare perspective, this is not sufficient. For decades, the US government has tried to limit supply to increase prices. A reduction in yields (production) might even be beneficial to farmer profits if it is offset by increasing prices, which in turn depend on how inelastic demand is. By the same token, consumer surplus is a function of commodity prices.

Real prices for commodity crops have been downward trending over the last century as production increases have outpaced demand increases. There were some spikes, e.g., the 1970s, but a general downward trend. After 2005, prices tripled due to an outward shift in demand for biofuels as well as production shortfalls. A better understanding of production and demand trends is crucial to model future prices. Prices for agricultural commodities move closely together for major producers that have access to a seaport. Landlocked countries with high transportation costs might be somewhat insulated from global movements, but they generally account for a small share of global supply and demand.

In addition to climate change, several demographic and economic challenges compound the challenge of ensuring affordable, available food. Global population continues to rise. Large segments of the world are enjoying greater purchasing power and their tastes are shifting toward more resource-intensive foods. It is unlikely that land intensification alone will be the route for satisfying food demand growth. Conversion from existing forests and peatlands to cultivable land would have the perverse effect of releasing more CO_2 and lead to a positive feedback loop of additional warming and further yield declines.[15]

If we treat climate change as a supply shock in existing areas, the split between induced additional supply and reductions in demand will be determined by the ratio of the demand and supply elasticities. A correct estimation of these elasticities is hence crucial, although empirically very challenging. Adopting a standard OLS approach will not yield consistent estimates of supply elasticities because producers endogenously determine planting area. Consider the traditional estimation setup that regresses supply, measured in aggregate calories produced in a year, on that crop's futures price. If farmers are aware of some crop-specific threat, for example a pest outbreak, they will cut back on the land area planted with that crop and switch to another crop. While demand has not changed, output has dropped, leading to higher futures prices. The outcome is a price increase in response to an output reduction—a movement along the demand curve. The classical regression setup would capture it as a supply response, implying a negative, or downward biased supply elasticity toward zero.

A unique feature of most commodity crops is that they are storable and storage can smooth production across periods. When prices are high, farmers sell their inventories and at the same time increase effort in the next period to benefit from higher prices. Starting with this storage model, Roberts and Schlenker (2013) use an instrumental variables approach to estimate the aggregate supply elasticity for four key commodities—corn, wheat, rice, and soybeans—which are all storable. They use the contemporaneous supply shock to identify the demand equation, a known instrument since Wright (1928) introduced the concept of instrumental variables. The new feature is that past supply shocks can be used to instrument the futures price and identify the supply function. They find statistically significant supply and demand elasticities, although both are fairly small at 0.11 and −0.05, respectively, suggesting that shocks to output will result in significant price changes. Two-thirds of any food that is diverted to biofuels will hence come from new production, while the remaining third will come from reductions in demand.

Has warming already affected world production and prices? Lobell et al. (2011) use the estimated elasticities to answer this question. They perform a country-level analysis by merging nationally averaged yields data for maize, rice, soybeans, and wheat with crop maps and monthly temperature/precipitation data. After linking production to weather, they predict output under the observed climate and one that takes out observed temperature and precipitation trends. They find that between 1980 and 2005, observed climate trends had already impacted crop production compared to a counterfactual without the observed trends, and that prices were 18% higher than what they would have been without trends.

3. Conclusion

We have examined recent statistical studies linking agricultural yields to weather fluctuations in a panel setting. We present new evidence that these models, estimated using historic data from 1950 to 2011 are very capable of predicting US aggregate yields out-of-sample for the years 2012–2015. The single best predictor is a measure of hot degree days that only counts for how long and by how much temperatures exceed 29°C (84°F). Aggregate production forecasts are the key variable needed to predict changes in prices.

Statistical panel models hence give useful first-order effects of predicted changes in climate, which already had a measurable effect since 1980. We discussed whether adaptation can reduce predicted losses in the future. Recent evidence finds cases of adaptation, but the effects are limited in the United States where the sensitivity to observed climate trends is roughly comparable to annual weather fluctuations, although the former should induce farmers to adapt much more. At this point, it seems

[15] Lambin and Meyfroidt (2011) recount examples of countries which increased food production without contributing to deforestation.

more likely that the main adaptation is not occurring at current growing areas, but rather through shifts in where crops are grown. Future research on how growing areas might shift is an active and important research field. If adaptation in current growing areas or movements into new growing areas cannot compensate the predicted production losses, prices will likely rise significantly to bring demand and supply back into equilibrium. This would likely benefit farmers, as the highly inelastic demand and supply elasticities imply that price increases more than offset production shortfalls. In a way, climate change would "accomplish" what government programs designed to limit supply have failed to achieve for decades. On the other hand, consumers will have to pay the higher prices and hence suffer a loss in surplus.

References

Asseng, S., Ewert, F., Martre, P., Rötter, R.P., Lobell, D.B., Cammarano, D., Kimball, B.A., Ottman, M.J., Wall, G.W., White, J.W., Reynolds, M.P., Alderman, P.D., Prasad, P.V.V., Aggarwal, P.K., Anothai, J., Basso, B., Biernath, C., Challinor, A.J., De Sanctis, G., Doltra, J., Fereres, E., Garcia-Vila, M., Gayler, S., Hoogenboom, G., Hunt, L.A., Izaurralde, R.C., Jabloun, M., Jones, C.D., Kersebaum, K.C., Koehler, A.-K., Müller, C., Naresh Kumar, S., Nendel, C., O'Leary, G., Olesen, J.E., Palosuo, T., Priesack, E., Eyshi Rezaei, E., Ruane, A.C., Semenov, M.A., Shcherbak, I., Stöckle, C., Stratonovitch, P., Streck, T., Supit, I., Tao, F., Thorburn, P.J., Waha, K., Wang, E., Wallach, D., Wolf, J., Zhaoand, Z., Zhu, Y., 2015. Rising temperatures reduce global wheat production. Nat. Clim. Chang. 5, 143–147.

Auffhammer, M., Hsiang, S.M., Schlenker, W., Sobel, A., 2013. Using weather data and climate model output in economic analyses of climate change. Rev. Environ. Econ. Policy 7(2), 181–198.

Balter, M., 2013. Farming was so nice, it was invented at least twice. Science.

Büntgen, U., Tegel, W., Nicolussi, K., McCormick, M., Frank, D., Trouet, V., Kaplan, J.O., Herzig, F., Heussner, K.-U., Wanner, H., Luterbacher, J., Esper, J., 2011. 2500 years of European climate variability and human susceptibility. Science 331(617), 578–582.

Burke, M., Emerick, K., 2013. Adaptation to climate change: Evidence from US agriculture. Working Paper.

Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. Nature 527, 235–239.

Burke, M.B., Miguel, E., Satyanath, S., Dykema, J.A., Lobell, D.B., 2009. Warming increases the risk of civil war in Africa. Proc. Natl. Acad. Sci. 106(49), 20670–20674.

Butler, E., Huybers, P., 2013. Adaptation of US maize to temperature variations. Nat. Clim. Chang. 3, 68–72.

Colmer, J., 2015. The productivity effects of labour reallocation: Evidence from India. Working Paper.

Costinot, A., Donaldson, D., Smith, C.B., 2014. Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. NBER Working Paper 20079.

Crutzen, P.J., Ramanathan, V., 2003. The parasol effect on climate. Science 302(5651), 1679–1680.

Deschênes, O., Greenstone, M., 2011. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. Am. Econ. J. Appl. Econ. 3(4), 152–185.

Graff Zivin, J., Neidell, M., 2014. Temperature and the allocation of time: Implications for climate change. J. Labor Econ. 32(1), 1–26.

Henderson, J.V., Storeygard, A., Deichmann, U., 2015. Has climate change driven urbanization in Africa? Working Paper.

Hsiang, S.M., Burke, M., 2014. Climate, conflict, and social stability: What does the evidence say? Clim. Change 123(1), 39–55.

Hsiang, S.M., Meng, K.C., Cane, M., 2011. Civil conflicts are associated with the global climate. Nature 476(7361), 438–441.

Kahn, M.E., 2014. Climate change adaptation: Lessons from urban economics. NBER Working Paper 20716, 1–31.

Kala, N., 2015. Ambiguity aversion and learning in a changing world: The potential effects of climate change from Indian agriculture. Working Paper.

Kelley, C.P., Mohtadi, S., Cane, M.A., Seager, R., Kushnir, Y., 2015. Climate change in the Fertile Crescent and implications of the recent Syrian drought. Proc. Natl. Acad. Sci. 112(11), 3241–3246.

Kurukulasuriya, P., Mendelsohn, R., 2007. Crop selection: Adapting to climate change in Africa. World Bank Policy Research Working Paper 4307.

Lambin, E.F., Meyfroidt, P., 2011. Global land use change, economic globalization, and the looming land scarcity. Proc. Natl. Acad. Sci. 108(9), 3465–3472.

Lobell, D.B., Bänziger, M., Magorokosho, C., Vivek, B., 2011a. Nonlinear heat effects on African maize as evidenced by historical yield trials. Nat. Clim. Chang. 1, 42–45.

Lobell, D.B., Burke, M.B., 2010. On the use of statistical models to predict crop yield responses to climate change. Agric. For. Meteorol. 150(11), 1443–1452.

Lobell, D.B., Hammer, G.L., McLean, G., Messina, C., Roberts, M.J., Schlenker, W., 2013. The critical role of extreme heat for maize production in the United States. Nat. Clim. Chang. 3, 497–501.

Lobell, D.B., Schlenker, W., Costa-Roberts, J., 2011b. Climate trends and global crop production since 1980. Science 333(6042), 616–620.

Mendelsohn, R., Nordhaus, W.D., Shaw, D., 1994. The impact of global warming on agriculture: A Ricardian analysis. Am. Econ. Rev. 84(4), 753–771.

Miller, N.L., Hayhoe, K., Jin, J., Auffhammer, M., 2008. Climate, extreme heat, and electricity demand in California. J. Appl. Meteorol. Climatol. 47(6), 1834–1844.

Passioura, J.B., 1994. Physiology and determination of crop yield. In: Boote, K.J., Bennett, J.M., Sinclair, T.R., Paulsen G.M. (Eds.), The Yield of Crops in Relation to Drought, Chapter 23. American Society of Agronomy, Crop Science Society of America, Soil Science Society of America, pp. 343–359.

Paulsen, G.M., 1994. Physiology and determination of crop yield. In: Boote, K.J., Bennett J.M., Sinclair, T.R., Paulsen, G.M. (Eds.), High Temperature Responses of Crop Plants, Chapter 25. American Society of Agronomy, Crop Science Society of America, Soil Science Society of America, pp. 365–389.

Ramanathan, V., Feng, Y., 2009. Air pollution, greenhouse gases and climate change: Global and regional perspectives. Atmos. Environ. 43(1), 37–50.

Reilly, J., Hohmann, N., Kane, S., 1994. Climate change and agricultural trade: Who benefits, who loses? Glob. Environ. Change 4(1), 24–36.

Reilly, J., Tubiello, F., McCarl, B., Abler, D., Darwin, R., Fuglie, K., Hollinger, S., Izaurralde, C., Jagtap, S., Jones, J., Mearns, L., Ojima, D., Paul, E., Paustian, K., Riha, S., Rosenberg, N., Rosenzweig, C., 2003. U.S. agriculture and climate change: New results. Clim. Change 57(1–2), 43–67.

Roberts, M.J., Schlenker, W., 2013. Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. Am. Econ. Rev. 103(6), 2265–2295.

Rosenzweig, C., Parry, M.L., 1994. Potential impact of climate change on world food supply. Nature 367(6459), 133–138.

Schlenker, W., 2016. Panel models in agricultural impact assessments. Working Paper.

Schlenker, W., Hanemann, W.M., Fisher, A.C., 2005. Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach. Am. Econ. Rev. 95(1), 395–406.

Schlenker, W., Lobell, D.B., 2010. Robust negative impacts of climate change on African agriculture. Environ. Res. Lett. 5(1), 1–8.

Schlenker, W., Roberts, M.J., 2009. Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. Proc. Natl. Acad. Sci. U.S.A. 106(37), 15594–15598.

Schlenker, W., Roberts, M.J., Lobell, D., 2013. US maize adaptability. Nat. Clim. Chang. 3(8), 690–691.

Sheehy, J.E., Mitchell, P.L., Ferrer, A.B., 2006. Decline in rice grain yields with temperature: Models and correlations can give different estimates. Field Crops Res. 98(2–3), 151–156.
Smit, B., Skinner, M.W., 2002. Adaptation options in agriculture to climate change: A typology. Mitig. Adapt. Strat. Glob. Change 7(1), 85–114.
Snyder, R.L., 1985. Hand calculating degree days. Agric. For. Meterol. 35(1–4), 353–358.
Tack, J., Barkley, A., Nalley, L.L., 2015. Effect of warming temperatures on US wheat yields. Proc. Natl. Acad. Sci. 112(22), 6931–6936.
Taraz, V., 2015. Adaptation to climate change: Historical evidence from the Indian monsoon. Working Paper.
United Nations Population Fund, Ed., 2007. State of World Population 2007.
Welch, J.R., Vincent, J.R., Auffhammer, M., Moya, P.F., Dobermann, A., Dawe, D., 2010. Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperature. Proc. Natl. Acad. Sci. 107(33), 14562–14567.
Wright, P.G., 1928. The Tariff on Animal and Vegetable Oils. MacMillan, New York.
Zhang, D.D., Brecke, P., Lee, H.F., He, Y.-Q., Zhang, J., 2007. Global climate change, war, and population decline in recent human history. Proc. Natl. Acad. Sci. 104(49), 19214–19219.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix