COLUMBIA | SIPA
Center for Environmental Economics and Policy

CEEP Working Paper Series
Working Paper Number 25

March 2023

ERA5-Land and GMFD Uncover The Effect of Daily Temperature Extremes On US Agricultural Yields

Dylan Hogan and Wolfram Schlenker

https://ceep.columbia.edu/sites/default/files/content/papers/n25.pdf

ERA5-Land and GMFD Uncover The Effect of Daily Temperature Extremes On US Agricultural Yields

Dylan Hogan* Wolfram Schlenker†

November 2022

Abstract

Global agricultural commodity markets are highly integrated among major producers. Prices are driven by aggregate supply rather than what happens in individual countries in isolation. Estimating the effects of climate change hence requires a globally representative weather data set. Recently, two widely used data sets that provide daily or even hourly values, GMFD and ERA5-Land, became available. We formally test whether these global data sets are as good as more fine-scaled country-specific data in explaining yields and whether they give similar response functions. While GMFD and ERA5-Land have lower predictive skill for US corn and soybeans yields than the more fine-scaled PRISM data, they still correctly uncover the underlying non-linear temperature relationship. Predicted crop yield losses are of similar magnitude and precision using the daily data, but start to diverge from estimates using average temperature under increasing warming. All specifications using daily temperature extremes under any of the three weather data sets outperform models that use average temperatures. Correctly capturing the effect of daily extremes has a larger effect than the choice of weather data.

Key words: Climate Data, Daily Temperature, Panel Studies, Agriculture
JEL Codes: Q1, Q54

Acknowledgments: This work was supported by USDA National Institute of Food and Agriculture, Grant No. 2022-67023-36400.

*Columbia University, email: dth2133@columbia.edu
†Columbia University and NBER, email: wolfram.schlenker@columbia.edu

Assessing the effect of climate change on global food systems requires a global analysis of weather's impact on agricultural productivity. Recent studies using fine-scaled weather data for individual countries or regions have shown that temperature extremes, especially extreme heat, are a main driver of agricultural yields (12; 16; 3; 2). Incorporating the full temperature distribution between the daily minimum and maximum provides much better predictions of heat-related yield losses (14; 6). Averaging over time (monthly rather than daily data) or space (larger grids) can mask this nonlinear relationship (1). However, until recently, most global data sets have only provided monthly data that can mask daily extremes (e.g., CRU, University of Delaware) and global studies were forced to rely on this more aggregated monthly data (8).

Recently, two new daily data have become available and are used extensively. They are the Global Meteorological Forcing Dataset (GMFD), which includes daily minimum and maximum temperature measurements on a 0.25 degree grid. The other is ERA5-Land, which provides some of the most detailed temperature data both temporally (hourly) as well as spatially (0.1 degree) for the entire world. Given the global coverage over all agricultural area, these two data sets have the advantage of offering a standardized weather product, which is crucial for a unified global analysis that drives prices, comparative advantages, production, and trade (5; 7; 9). However, these global daily weather data sets have not been systematically assessed in how well they explain outcomes of interest compared to more detailed fine-scaled weather data sets that are available for individual countries. If the daily data include measurement error, the resulting estimates would suffer from attenuation bias, which is especially important for studies of non-linear effects. For example, if the daily maximum is not correctly captured due to measurement error in a data set, too much (or too little, depending on the sign of the error) of the temperature exposure is counted as harmful yield-decreasing heat rather than beneficial yield-increasing moderate temperatures, severely biasing both coefficients, given that the sign of the coefficients switches between these categories. The role of temperature extremes can only be uncovered if their exposure is correctly captured in the underlying weather data set.

To better understand whether global climate data sets can uncover the effect of extreme heat on crop yields, we estimate statistical yield-weather relationships using (1) a modified version of the fine-scaled PRISM data set, which only provides weather measurements in the United States, and (2) the same measurements constructed from the more aggregated, but globally available, data sets, GMFD and ERA5-Land. We compare model performance across data sets for the area where they overlap, i.e., the United States. We show that all three data sets give comparable response functions and climate change predictions. Our paper focuses on agriculture, but similar response functions have been estimated for other sectors (e.g., energy (11) or mortality (4)). While agriculture accounts for only 4% of global economic activity, it comprises more than a quarter of GDP in some of the least developed countries, which are among the most exposed to extreme weather and the least equipped to invest in adaptation. Accounting for the sensitivity of yields in these locations is vital for evaluating the extent of the inequalities associated with climate change impacts.

Our statistical analysis links yields to each weather data set. Yields data are publicly available at the county level in the United States: 73% of all counties report corn yields, while 63% of all counties report soybean yields in at least some of the years of our 70-year sample period 1950-2019. We aggregate all weather data sets to the county level, weighted by the share of cropland in a cell (as measured by a satellite-derived cropland mask that is at the 30 $\times$ 30 meter scale). For the fine-scaled PRISM data set (1/24 degree grid), we link each grid cell to the county in which the centroid is located. For the coarser ERA5-Land and GMFD data set (0.1 degree and 0.25 degree, respectively), we weight by the area of each grid cell that falls within a county. Given the importance of non-linearity, we first derive all non-linear transformations before aggregating the data to the county level. We construct measures of degree days, which represent the total heat experienced by crops above a threshold in a day, as well as the length of time crops are exposed to each one-degree Celsius temperature interval in each day, summed across the growing season April - September.

For each data set, we estimate nonlinear relationships between temperatures and corn or soybean yields, controlling for precipitation, a quadratic time trend (to capture technological change) as well as county fixed effect to capture all time-invariant confounding factors (e.g., soil quality). Earlier research has shown that different temperature ranges can have opposing effects, where moderate temperatures are yield enhancing while very hot temperatures greatly reduce yields. We consider three functional form assumptions to capture possible nonlinear relationships. First, we estimate piecewise linear regressions that separate the impacts of moderate days and extreme days around a critical temperature threshold. Since this specification requires a break-point between beneficial temperature days and costly temperature days, we apply a data-driven cross validation approach to inform this decision separately for each crop and weather data set. Second, we estimate a flexible 8^{th} order Chebychev polynomial, which smoothly characterizes the temperature yield relationship and does not require a threshold assumption. Third, we consider a semi-parametric specification, which estimates a separate yield effect for each three-degree temperature range.

1 Empirical Results

We find that relationships estimated on county-level yield data in the United States are quantitatively similar for all three functional form assumptions.

We asses model performance using a cross-validation procedure that quantifies the ability of the response functions to predict crop yields outside of the estimation sample. We show that PRISM models outperform the ERA5-Land and GMFD models in predicting crop yield impacts out of sample, particularly for soybean yields, with ERA5-Land and GMFD models performing about the same despite different temporal and spatial resolutions. While there are differences in out-of sample prediction accuracy, we also show that spatially-uniform warming scenarios applied to the response functions project climate impacts of similar magnitude and precision.

All specifications and weather data uncover an asymmetric relationship where yields are increasing in temperature for moderate temperature ranges, but sharply decrease in temperature at the upper end.

Figure 1 displays the response functions and 95% confidence intervals for the temperature sensitivity of crop yields for all three weather data sets. The figure has six plots, with rows examining different crops and columns different functional forms for the temperature-yield relationship. In particular, we estimate responses for two crops, corn (top row) and soybeans (bottom row) with three functional form assumptions, a piecewise linear specification following the agronomic concept of degree days (left column), an 8^{th} order polynomial in temperature (middle column), and a step function that estimates a separate temperature effect for each three-degree bin (right column). Within each plot, response functions estimated from PRISM, ERA5-Land, and GMFD data are indicated by color. The red lines show responses based on PRISM, which provides spatially fine-scale daily data available only for the United States. The blue lines show responses based on ERA5-Land, using the data for the US in this globally representative hourly data set at the 0.1 degree resolution. The green lines show responses based on GMFD, again focusing on the US portion of the data for this globally representative daily data at the 0.25 degree resolution.

The dependent variable in all cases is the log of crop yield (i.e., bushels per acre). The response functions in the figure are normalized at zero at the PRISM exposure-weighted average temperature.[1] All three climate data sets estimate similar response functions across

[1]Since our model includes county fixed effects, the results should be interpreted in relative terms, i.e., the difference in height (y-value) for two different temperatures (x-values) rather than in absolute levels so the

crops and specifications, with modest yield increases associated with going from cold to moderate temperatures and sharp yield decreases once temperatures pass a threshold. Notably, the cross validation procedure we use to choose the critical temperature threshold for the piecewise linear specification yields different breakpoints for each climate data set. For corn, the threshold varies from 27°C (ERA5-Land) to 30°C (GMFD). For soybeans, the threshold varies from 28°C (ERA5-Land) to 30°C (PRISM).

Table 1 shows the effect of changes in the temperature exposure on annual crop yields. Specifically, it gives the predicted yield change (in percent) for replacing a full day at a moderate 20°C, i.e., a 24-hour exposure to 20°C (recall that we are counting partial days) with either (i) a full day at 40°C or (ii) a full day at the model-specific yield-maximizing temperature. The effect of extreme heat on crop yields is comparable across climate data sets. Substituting a full day at a moderate 20°C with a full day at 40°C decreases corn yields between 4.1% and 4.2% (SE 1.0%-1.3%) for all three data sets. However, models disagree on the benefits of being at the yield-maximizing temperature. Substituting a full day at a moderate 20°C for a full day at the yield-maximizing temperature increases yields by 1.1-1.2% (SE=0.3%) based on PRISM and GMFD relationships, but only 0.4% (SE=0.2%) according to ERA5-Land. The pattern is similar for soybeans, with an additional full day at 40°C reducing yields between 3.6% to 4.4% and a full day at a yield-maximizing temperature increasing yields between 0.3% and 1.1%.

All regression models include two additional sets of controls: a quadratic function of total growing season precipitation and state-specific quadratic time trends. The precipitation control features a similar inverted-U shape with statistically significant linear and quadratic terms for all crops and climate data sets, indicating that moderate precipitation levels are optimal for yields. State-specific quadratic time trends control for technological progress common to counties within a state. The temperature sensitivity of crop yields is robust to changes in the time trend, e.g., by including more flexible controls like year fixed effects to account for common year-specific shocks.

The fine-scaled PRISM data has the best out-of-sample model performance

We compare regression models from each data set using a cross validation procedure that calculates the root-mean squared prediction error (RMS) of out-of-sample predictions. In particular, we measure the reduction in RMS relative to a baseline model that includes county fixed effects and quadratic state-level time trends but excludes the weather variables.

normalization is inconsequential.

We thereby measure how much of the error can be explained by the inclusion of various weather variables. The cross validation procedure consists of 1,000 repetitions in which we randomly use 85% of the years in our panel data set in the estimation of the regression coefficients and then predict crop yields for the remaining 15%. We sample years rather than observations to avoid significant spatial correlation in yields across counties within a year.

Figure 2 shows comparisons of RMS reductions across climate data sets for corn and soybeans, with PRISM indicated by the red bar, ERA5-land indicated by the blue bar, and GMFD indicated by the green bar. On average across our three main specifications (piecewise linear, 8^{th} order polynomial, and 3-degree bins), PRISM improves upon the baseline model of corn yields by about 15.6%, which is a 32.0% larger reduction in RMS than ERA5-Land (11.8%) and 28.5% larger reduction than GMFD (12.17%). For soybean yields, PRISM reduces RMS by 14.2% on average, which is 29.8% larger than ERA5-Land (11.0%) and 36.5% larger than GMFD (10.4%). Note that for both crops, Welch tests find a statistically significant differences in out-of-sample performance between PRISM and either of the two global data sets, however, a statistically insignificant difference between the two global data sets ERA5-land and GMFD. Consistent with prior work (12; 15), the three main specifications, which leverage daily temperature extremes, outperform models that include quadratic functions of the average temperatures over the growing season. This is even true across models: the RMS reduction is larger using any of the weather data sets or non-linear specifications than when using average temperatures. Accounting for daily extremes consistently improves model fit.

Predicted yield declines from uniform warming scenarios are similar in terms of magnitude and precision across all climate data sets for the models using daily temperature extremes, but diverge for the model using average temperatures.

Figure 3 summarizes projected yield losses, in percent, from uniform warming scenarios between 1°C and 4°C, including point estimates and 95% confidence intervals.[2] The vertical axis is the percent reduction in aggregate US crop yields relative to average yields between

[2]Many studies have used general circulation models (GCMs) to project agricultural losses of spatially heterogeneous climate change scenarios, such as those from CMIP5 or CMIP6. However, biases in GCM simulations can be large relative to historical observations from a particular data set (10). These biases propagate through nonlinear response functions to projections of climate impacts. The CMIP ensembles account for this error by bias-correcting GCM output to a single historical data set, e.g., GMFD for CMIP5 and ERA5 for CMIP6. For our purposes, uniform warming scenarios allow us to compare across data sets and to highlight for what range of warming the models start to diverge.

1960 and 1989 after a uniform shift in the temperature distribution, and the horizontal axis indicate the four model specification we used: the first three use daily data, while the fourth (average temperature) relies on the seasonal average. The top and bottom frames show projected impacts for corn and soybean yields, respectively.

While some counties experience yield benefits from moderate warming scenarios, negative aggregate impacts across all specifications, crops, and warming scenarios are driven by increases in the number of extremely hot days. Disagreement between climate data sets regarding the threshold between beneficial and harmful temperature days drives modest differences in impacts for a given main specification and crop. Across the first three specifications and all climate data sets, corn losses from a 2°C warming range from 13-16% and soybeans losses range from 9-11%. A 4°C uniform warming is projected to reduce yields between 31-35% for corn and 25-29% for soybeans. Consistent with earlier studies focusing on the functional form of temperature in climate analyses (12; 15), all climate data sets agree that excluding daily variation in temperature by focusing on seasonal averages results in lower projected impact as they fail to capture the asymmetric relationship where being above the optimal temperature is much worse than being below the yield-maximizing temperature.

2 Discussion

Several studies have highlighted the importance of extreme temperatures on agricultural yields using fine-scaled data. Recently, two global data sets have become available that are temporally fine-scaled (providing daily minimum and maximum temperature or even hourly temperature values), but spatially more aggregated. The effect of temperature extremes might be masked if the weather data is aggregated, both temporally (a monthly average hides peak exposure within a month) as well as spatially (aggregating over larger areas can hide that some part of it experienced extremes).

While these global data sets have been used, there has not been a systematic assessment of whether they uncover the non-linearities and crucial effects of temperature extremes. Our results suggest that novel climate data sets with daily observations are useful additions for estimating global crop yield response functions, particularly if the objective is to simulate impacts of future climate scenarios. All three data sets correctly capture the harmful effects of temperature extremes when linked to US corn and soybean yields. When we aggregate the data to seasonal averages, the common approach undertaken in previous global assessments that relied on monthly weather measures that cannot adequately capture daily temperature extremes, the out-of-sample predictive performance is greatly reduced. In fact, any of the three data sets using any of the three specifications using daily data perform better than any specification using monthly averages. The provision of daily global temperature data is hence an important step forward to accurately simulate the effects of climate change on agricultural yields and prices.

Differences in predictive skill may be driven by measurement error induced by a combination of the spatial and temporal resolution of the data sets as well as the interpolation methods used to fill gaps in the weather record. Appendix Figure B1 further examines the role of spatial and temporal aggregation. A "resampled PRISM model" estimated after spatially aggregating PRISM to the ERA5-Land resolution before constructing our county panel reduces the gap in RMS between the PRISM and ERA data sets by about 55%. In other words, half of the lead PRISM has over ER5 in predicting yields is caused by the fact that it is spatially more disaggregated. The remaining gap in RMS is likely due to interpolation algorithms (or "reanalysis") applied to the underlying weather observations. Reanalysis incorporates outputs of numerical weather prediction models to create a globally consistent record, which, in the process, may mask the most extreme weather events that are critical for accurately predicting yields.

On the other hand, temporally aggregating ERA5-Land from hourly temperature readings to daily minimum and maximum temperature before calculating degree days using a spline interpolation between minimum and maximum temperature negligibly reduces RMS, suggesting that the sub-daily hourly records do not provide important additional information. Taken together, a global data set on minimum and maximum temperature correctly identifies the importance of temperature extremes and provides yield predictions that closely mirror more fine-scaled weather data.

One qualification of our finding is that the global weather data sets obviously incorporate the weather station network of a country, which is dense in the United States. Other countries have fewer stations and the global weather data sets accordingly might be worse at capturing extreme temperature extremes.

3 Conclusion

We find that two global data sets, ERA5-Land and GMFD, that provide daily temperatures uncover a nonlinear relationships between extreme heat and US crop yields that closely align with the results that are obtained using the high resolution PRISM data set. In particular, we find similar effects of extreme temperatures on both corn and soybean yields across three specifications and the three data sets. On the other hand, the benefits of yield-maximizing temperatures vary somewhat between models: the estimated relationship is lower under ERA5-Land than the other two climate data sets. Predictive power, in terms of RMS, is lower for the global data sets than for PRISM, with GMFD and ERA5-Land performing about the same. Despite these differences, projections of climate impacts from spatially uniform warming are similar across models that include daily temperature extremes. Finally, we validate the benefits of daily weather observations in the global data sets. Specifying temperature as averages over the growing season rather than accounting for daily extremes universally lowers predictive power and simulated climate impacts. The appropriate functional form is more important than the choice of weather data set.

References

[1] Maximilian Auffhammer, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, 7(2):181–198, 2013.

[2] Marshall Burke and Kyle Emerick. Adaptation to Climate Change: Evidence from US Agriculture. *American Economic Journal: Applied Economics*, 8(3):106–140, August 2016.

[3] Ethan E. Butler and Peter Huybers. Adaptation of US maize to temperature variations. *Nature Climate Change*, 3:68–72, January 2013.

[4] Tamma Carleton, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang. Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*. *The Quarterly Journal of Economics*, 137:2037–2105, April 2022.

[5] Arnaud Costinot, Dave Donaldson, and Cory Smith. Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World. *Journal of Political Economy*, 124(1):205–248, 2016. _eprint: https://doi.org/10.1086/684719.

[6] Anthony Louis D'Agostino and Wolfram Schlenker. Recent weather fluctuations and agricultural yields: implications for climate change. *Agricultural Economics*, 47(S1):159–171, 2016.

[7] Christophe Gouel and David Laborde. The crucial role of domestic and international market-mediated adaptation to climate change. *Journal of Environmental Economics and Management*, 106:102408, 2021.

[8] David B. Lobell, Wolfram Schlenker, and Justin Costa-Roberts. Climate Trends and Global Crop Production Since 1980. *Science*, 333(6042):616–620, 2011. _eprint: https://www.science.org/doi/pdf/10.1126/science.1204531.

[9] Ishan Nath. Climate Change, The Food Problem, and the Challenge of Adaptation through Sectoral Reallocation. Technical Report 21-29, Center for Economic Studies, U.S. Census Bureau, September 2021. Publication Title: Working Papers.

[10] Julian Ramirez-Villegas, Andrew J Challinor, Philip K Thornton, and Andy Jarvis. Implications of regional improvement in global climate models for agricultural impact research. *Environmental Research Letters*, 8(2):024018, 2013.

[11] Ashwin Rode, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E. Kopp, Kelly E. McCusker, Ishan Nath, James Rising, and Jiacan Yuan. Estimating a social cost of carbon for global energy consumption. *Nature*, 598(7880):308–314, October 2021.

[12] Wolfram Schlenker and Michael J. Roberts. Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598, 2009.

[13] R. L. Snyder. Hand calculating degree days. *Agricultural and Forest Meterology*, 35(1-4):353–358, October 1985.

[14] Jesse Tack, Andrew Barkley, and Lawton Lanier Nalley. Effect of warming temperatures on US wheat yields. *Proceedings of the National Academy of Sciences*, 112(22):6931–6936, 2015.

[15] Jesse Tack, Andrew Barkley, and Lawton Lanier Nalley. Effect of warming temperatures on us wheat yields. *Proceedings of the National Academy of Sciences*, 112(22):6931–6936, 2015.

[16] Jarrod R. Welch, Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures. *Proceedings of the National Academy of Sciences*, 107(33):14562–14567, 2010.

Table 1: Crop Yield Sensitivity to Various Temperatures

Yield relative to a full day at 20°C day	PRISM	ERA5-Land	GMFD
Corn yields			
40°C	-4.15	-4.09	-4.24
	(1.02)	(1.25)	(1.21)
Yield-maximizing temperature	1.06	0.39	1.19
	(0.28)	(0.20)	(0.31)
Soybean yields			
40°C	-4.29	-4.33	-3.56
	(0.44)	(0.42)	(0.23)
Yield-maximizing temperature	1.02	0.26	1.13
	(0.22)	(0.19)	(0.20)

Notes: Table provides the predicted change in yields, in percent, associated with replacing a day (24-hour exposure) at a moderate 20°C with (1) a day at 40°C day and (2) a day at the yield-maximizing temperature, which varies based on the climate data set. Standard errors are provided in parentheses.

Figure 1: Comparing the Yield-Temperature Response across the Three Weather Data Sets

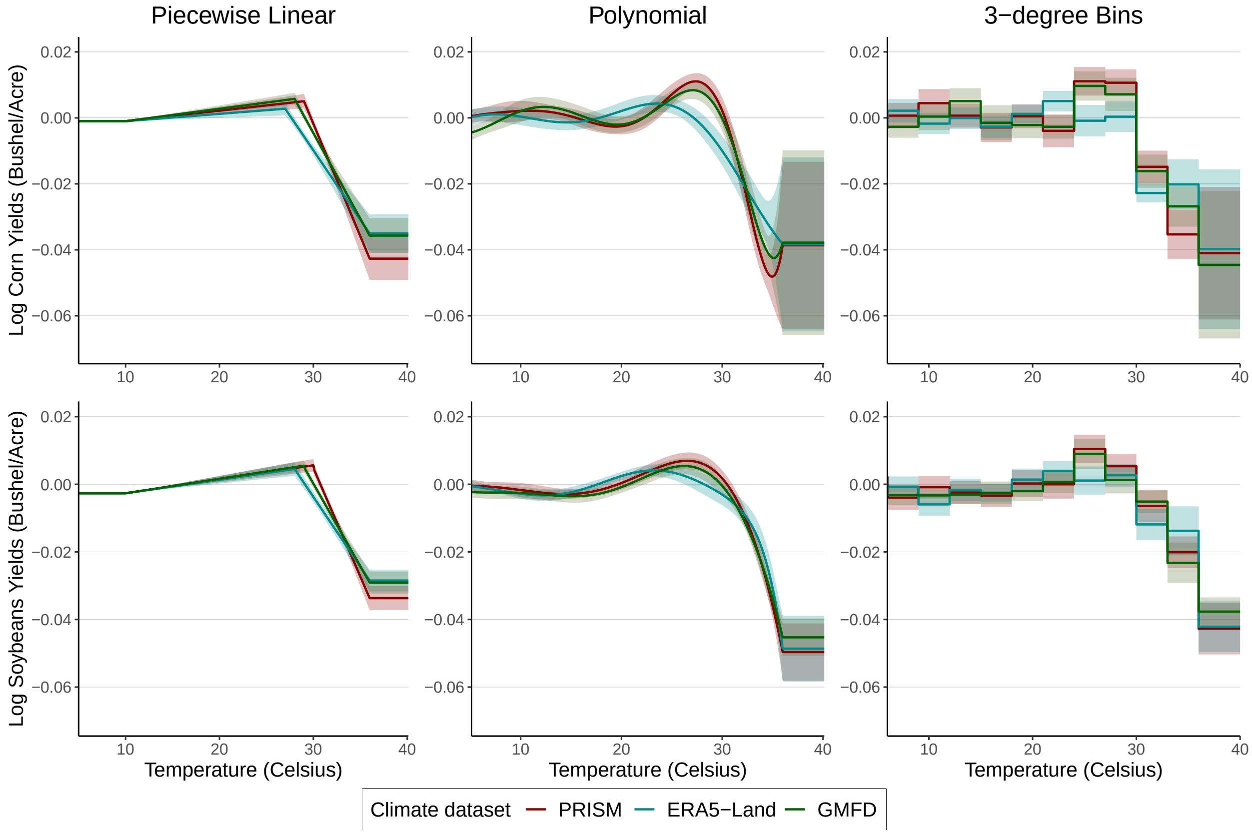

Notes: Each graph estimates the relationship between yields and temperature using both the fine-scaled PRISM data set (shown in red) as well as the more aggregate but globally available ERA5-Land (shown in blue) and GMFD (shown in green). The 95% confidence bands are added. The top row provides results for corn yields, while the bottom row gives the results for soybean yields. The left column estimates a piecewise linear function, the middle column an 8th order polynomial in temperature, and the right column uses temperature bins.

Figure 2: Out-of-sample Model Predictions Across Specifications and Weather Data Sets

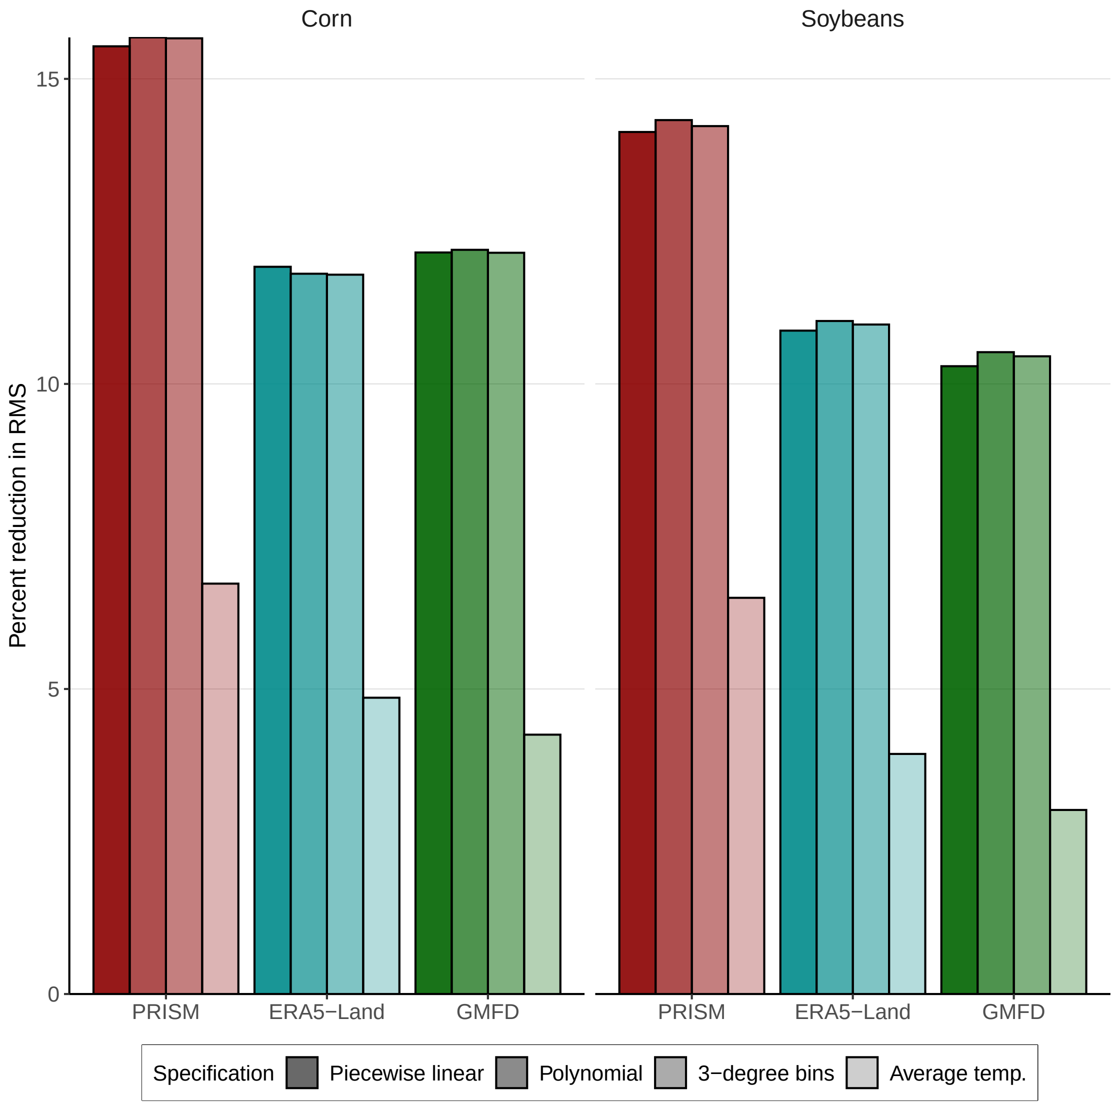

Notes: Figure compares out-of-sample prediction for piecewise linear regression models estimated using weather observations from PRISM, ERA5-Land, and GMFD data sets. The vertical axis shows the percent reduction in root-mean-squared error (RMS) relative to a baseline model that excludes temperature and precipitation variables. For each data set, piecewise linear response functions are estimated 1,000 times, each time randomly sampling 85% of the years from the full panel. RMS is calculated based on each piecewise linear model's prediction of the remaining 15% of years. Note that years are sampled rather than observations because there is significant spatial correlation in yields across counties within a year, whereas year-to-year weather fluctuations are random. Results are provided for corn and soybeans yields, and climate data sets are indicated by color.

Figure 3: Simulating Yield Losses For a Range of Uniform Warming Scenarios

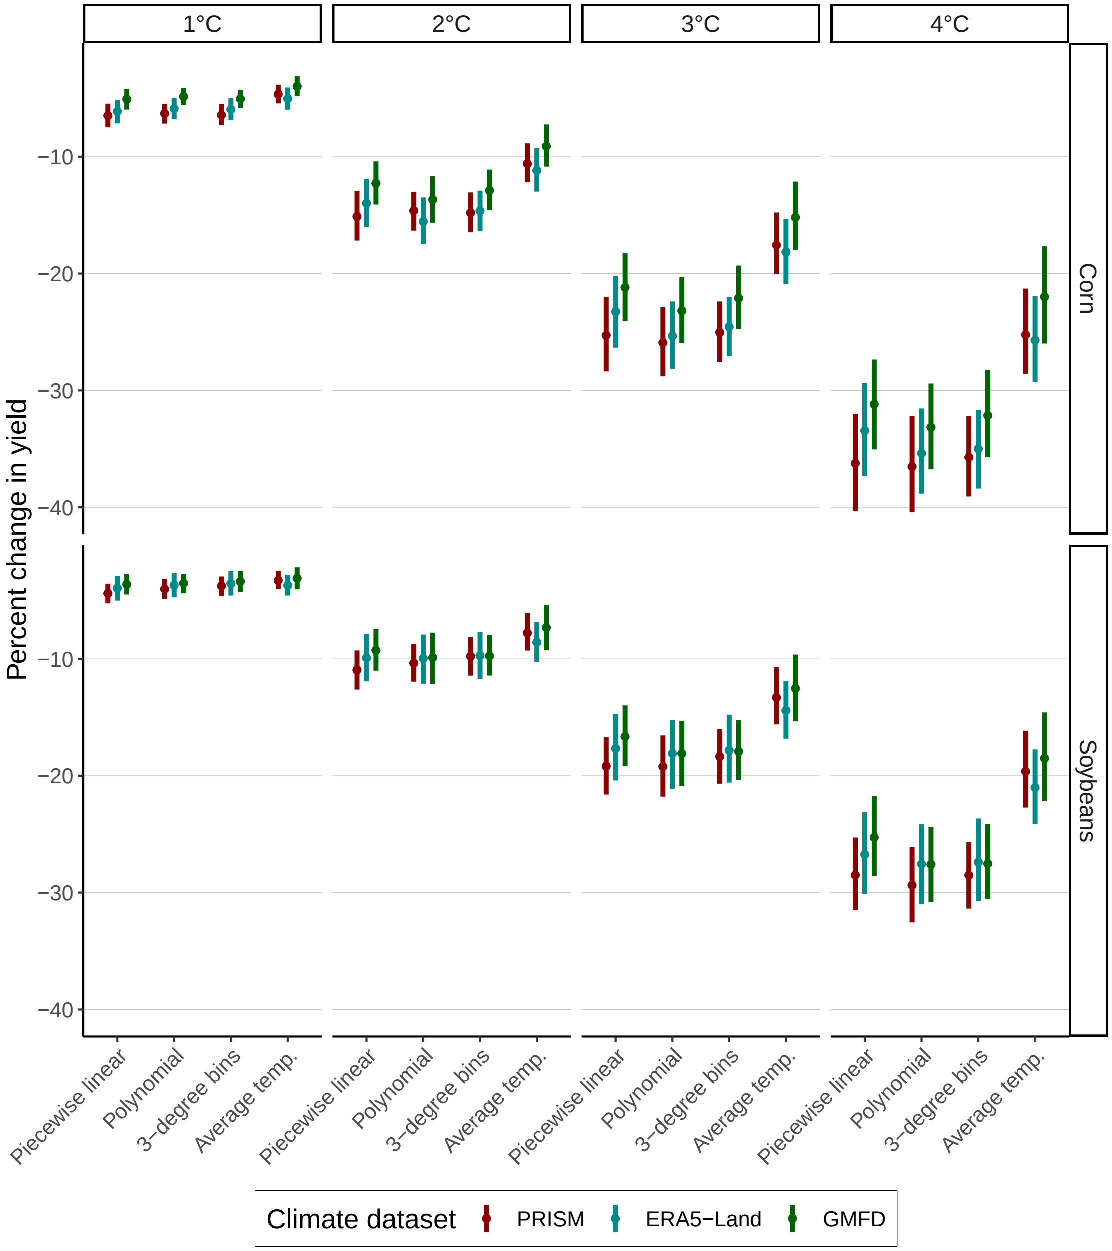

Notes: The four panels in each row provide projected climate change impacts on crop yields under uniform warming scenarios between 1°C and 4°C. Each panel provides the point estimates (circle) and 95% confidence bands (vertical bars) for twelve estimates when each of the four specifications are paired with the three weather data sets. The four specifications are listed on the horizontal axis: piecewise linear, eighth-order Chebychev polynomial, 3-degree bins, and a quadratic in average temperature. Colors indicate the climate data set used to estimate the response functions (PRISM in red, ERA5-Land in blue and GMFD in green). The top panel shows projected impacts on log corn yields, while the bottom panel displays the results on log soybean yields. The vertical axis is the predicted decline in overall US yields in percent.

A Methods

Weather data. The three climate data sets used for this analysis are (1) a modified version of the the Parameter-elevation Regressions on Independent Slopes Model (PRISM) data set from the Northwest Alliance for Computational Science and Engineering based at Oregon State University; (2) the fifth generation of the European Reanalysis (ERA5-Land) data set from the European Centre for Medium-Range Weather Forecasts; and, (3) the Global Meteorological Forcing Dataset (GMFD) from the Terrestrial Hydrology Research Group at Princeton University.

The following paragraphs outline the data construction methods used to obtain county-level temperature and precipitation records from each data source. For each grid cell, we approximate the distribution of temperatures within each day using a sinusoidal curve fit between minimum and maximum temperature measurements (13), except for ERA5-Land, which provides hourly measures (see below). We calculate two measures of daily temperature exposure. The first is the amount of time a pixel is exposed to each 1°C interval each day. The second follows the agronomic concept of degree days, which measures for how long and by how much temperatures exceed a threshold. For example, for a threshold of 30°C, a hypothetical day of constant 32°C temperature contributes 2 degree days, as would two days at 31°C, while a day of constant 28°C temperature contributes 0 degree days. Note that in order to preserve nonlinearities in the temperature record, it is important to construct daily temperature measurements at each weather grid cell before aggregating measurements to the county level. All daily grid cell-level temperature and precipitation measurements are combined with a high resolution cropland raster that allows us to weight cells based upon the share of cropped area.[3] Weather data are linked to counties in the contiguous US in two ways: for the fine-scaled PRISM data we link a cell to county if it's centroid falls within the county. For the spatially coarser GMFD and ERA5-Land data sets, we derive the fraction of a grid cell that overlaps with the county boundary as described next in more detail:

PRISM provides daily minimum temperature, maximum temperature, and total precipitation data on a 4 km × 4 km grid across the contiguous United States. To stay consistent with the other weather sets, we focus on the period 1950-2019.[4] We link a PRISM cell to the county in which its centroid is located.

ERA5-Land provides average hourly temperature and total precipitation measurements on a global 0.1 degree (≈11.1 km) grid for the period 1950-2019. The process of aggregating these data to the county level is similar to the process applied to PRISM with two exceptions. First, hourly temperature data allow us to calculate daily temperature exposure directly

[3]We obtain high resolution cropland information from the USDA National Agricultural Statistics Service Cropland Data Layer (CDL). CDL provides crop-specific land cover masks for the continental US. Data are available annually from 2008 to 2021. We take an average of all available cropland rasters at the native 30 meter resolution before aggregating to the appropriate climate grid cell size.

[4]The raw PRISM data is modified to maintain a constant set of stations over time, thereby ruling out that changes in temperatures are due to a change in station coverage. To obtain a balanced panel of weather observations, we fill in missing weather measurements with the distance-weighted average of the cumulative density function of surrounding stations. For example, if the 10 closest weather stations are on average at their 70th percentile, the station's missing value is set to the 70th percentile of its own measurements.

instead of interpolating the within-day distribution using a sinusoidal curve.[5] Second, we aggregate weather observations to counties weighted by cropped area as well as the share of each grid cell that intersects the county. Thus, a cell that only partially overlaps a county's boundary is weighted less than one that lies entirely within the county.

GMFD provides daily minimum temperature, maximum temperature, and total precipitation data on a global 0.25 $\times$ 0.25 degree ($\approx$28 km) grid for the period 1950-2010. GMFD observations are aggregated to the county level using a combination of the methods described above. We use a sinusoidal interpolation of within-day observations to measure temperature exposure, and we include area weights in the spatial aggregation to account for the low resolution of the data.

Yield data. Yield data are collected from the U.S. Department of Agriculture's National Agricultural Statistical Service. We use corn and soybean yields from the years 1950 to 2019 for the main analysis; however, 9 years are dropped for regressions with GMFD, which only provides a weather record through 2010. When merged with the weather data, our panel consists of 128,169 observations for corn yields and 102,674 observations for soy yields, with approximately 14,000 observations dropped for regressions on GMFD weather variables.

Regression models. We model the relationship between weather and yields following (12), which assume that temperature effects are additively separable over the growing season: an additional growing degree day experienced just after planting has the same effect on yields as an additional growing degree day experienced right before harvest.[6] In particular, we model plant growth $g(h)$ as a nonlinear function of heat h. Thus, log yield y_{it} in county i and year t is

$$y_{it} = \int_{\underline{h}}^{\overline{h}} g(h)\phi_{it}(h)dh + z_{it}\delta + c_i + \epsilon_{it}. \tag{A1}$$

$\phi_{it}(h)$ is the time distribution of heat over the growing season for county i and year t; $\underline{h}$ and $\overline{h}$ are the lower and upper bounds of temperatures observed over the growing season; z_{it} is a vector of additional time-varying controls, including a quadratic function of total growing season precipitation and quadratic state-specific time trends,which account for technological progress common to all counties within a state; c_i are county fixed effects, which flexibly control for time-invariant characteristics of counties that confound the weather-yield relationship, such as soil quality. Standard errors are clustered at the state level.

We focus on the months April through September to capture the main growing season for both corn and soybeans. Using the temperature constructions described above for each

[5] Hourly data should produce temperature variables with less measurement error than those derived from PRISM. In Appendix B, we assess the degree to which this impacts model performance by applying the sinusoidal fit to the minimum and maximum daily temperatures from ERA5-Land and comparing out-of-sample predictions. We find that models estimated with interpolated ERA5-Land data perform negligibly worse than those based on the raw hourly data.

[6] Note that this assumption diverges from some crop simulation models, which account for different temperature effects over the life cycle of a plant.

1°C temperature interval, we approximate A1 with

$$y_{it} = \sum_{h=-4}^{41} g(h+0.5)\left[\Phi_{it}(h+1) - \Phi_{it}(h)\right] + z_{it}\delta + c_i + \epsilon_{it} \tag{A2}$$

where $\Phi_{it}(h)$ is the cumulative distribution function of heat in county i and year t. We model $g(h)$ with three functional form assumptions: a piecewise linear function estimated via degree days at a crop-specific threshold temperature, an eighth order Chebychev polynomial, and a non-parametric specification in which a separate yield effect is estimated for each 3°C temperature bin up to 36°C.

B Contrasting Spatial and Temporal Resolution

There are two differences between the ERA5-Land and PRISM model. The former is spatially more aggregated (11km grid versus 4km grid size), yet temporarily less aggregated (24 hourly temperature observations rather than providing the daily maximum and minimum temperature).

To further examine the effect of spatial and temporal aggregation, we aggregate the raw temperature data from the finer scale to the more aggregate scale and then derive again our non-linear temperature transformation. Specifically, in the case of the spatially disaggregated PRISM data, we first aggregate minimum and maximum temperature on the 4km PRISM grid to the same 11km resolution of the ERA-5grid and then re-derive the temperature measures (degree-days, etc). Similarly, for the temporally more disaggregated ERA-Land weather data, we take the minimum and maximum of the 24 hourly observations and then use the sine-interpolation between minimum and maximum temperature.

The results are shown in Figure B1. Note that aggregating the PRISM data to the same spatial resolution as the EAR5 grid (second bar) is in between the height of the PRISM grid (first bar) and ERA-5 bar (third bar), i.e., half of the difference in the better out-of-sample prediction of PRSIM is attributable to a finer scale that better captures local extremes. On the other hand, the finer temporal resolution of the ERA-5 has no benefit - when we aggregate the hourly data to the daily minimum and maximum (4th bar), it is indistinguishable from using the hourly data (3rd bar).

Figure B1: Out-of-Sample Model Prediction Accuracy: Spatial and Temporal Resolution

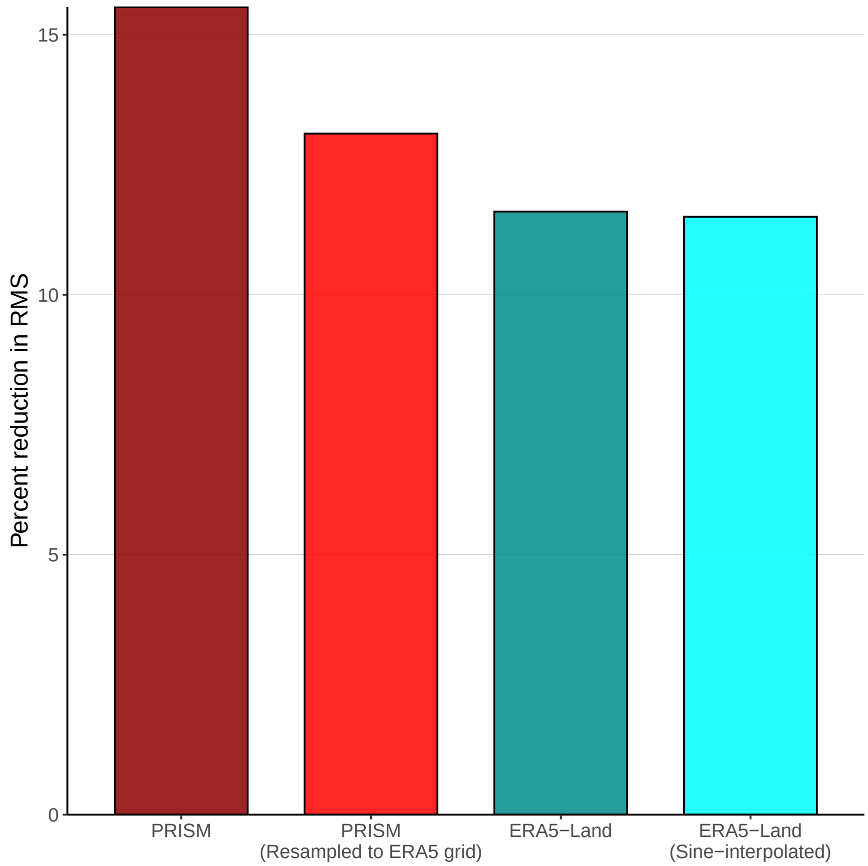

Notes: Figure compares out-of-sample corn yield predictions for piecewise linear regression models estimated using weather observations from PRISM and ERA5-Land data sets. Dark red and cyan bars show the percent reduction in root-mean-squared error (RMS) relative to a baseline model as shown in Figure 2. The light red bar shows results for a piecewise linear model estimated on PRISM data aggregated from a 1/24 degree grid to a 0.1 degree grid. The light cyan bar shows a model estimated on degree days calculated using a sinusoidal interpolation of daily maximum and minimum temperature values, rather than from hourly observations in the raw data. For each data set, piecewise linear response functions are estimated 1,000 times, each time randomly sampling 85% of the years from the full panel. RMS is calculated based on each piecewise linear model's prediction of the remaining 15% of years.