Climate Change and Agriculture Reconsidered

Anthony Fisher♣, Michael Hanemann♣, Michael J. Roberts♠, and Wolfram Schlenker♡

May 2009

Abstract

Despite the existence of a large and growing literature on the potential impact of climate change on agriculture, there still exists some disagreement about the magnitude and even the sign. Our own research suggests that the impact on U.S. agriculture is likely to be strongly negative, based on a series of studies in which we link farmland values to climate variables, and crop yields to both climate and yearly weather variables. Results are significant, robust, and consistent across data sets and methods. A recent but influential study by Deschênes and Greenstone (2007b) reports dramatically different results: based on regressions of agricultural profits and yields on weather variables, they conclude that the impact of climate change will be either insignificant or positive. In this paper we reconcile these conflicting results.

Likely explanations for the divergence between our findings and theirs are: (1) missing and almost certainly incorrect weather and climate data in their study, amplified by the use of state-by-year fixed effects that absorb most year-to-year weather variation but leave data errors intact; (2) their unusual and in our judgment incorrect treatment of climate-change predictions; (3) their use of the older Hadley II climate model for climate change predictions rather than the more recent and less optimistic Hadley III climate model used in the Fourth IPCC Report and in our studies; and (4) theoretical difficulties in their profit-based approach due to the confounding effects of storage and possibly also capital and inventory adjustments and local price movements associated with weather fluctuations. A careful account of these factors shows that the balance of evidence weighs heavily on the side of severe adverse potential impacts on U.S. agriculture by the end of the century, and probably sooner, stemming from anticipated global warming.

♣Department of Agricultural and Resource Economics, University of California at Berkeley, 207 Giannini Hall #3310, Berkeley, CA 94720. Email: fisher@are.berkeley.edu and hanemann@are.berkeley.edu.

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

♡ Department of Economics and School of International and Public Affairs, Columbia University, 420 West 118th Street, Room. 1308, MC 3323, New York, NY 10027. Email: ws2162@columbia.edu.

Agriculture is arguably the sector of the economy most directly exposed to climate and thus likely to be affected by climate change. To date, however, there exists considerable disagreement about the magnitude of potential impacts. Disagreement stems from differences in both methodology and empirical measurement. The recent paper by Deschênes and Greenstone (2007b), henceforth DG, is an important contribution to the climate impact literature and to the debate on economic methodology. DG's findings and conclusions hinge on two key factors: (1) how they navigate the distinction between climate and weather, and (2) the metrics they use for measuring both weather effects and economic impacts. In this paper we revisit DG in an attempt to reconcile differences between their work and our own.

Climate scientists emphasize the distinction between weather and climate. Weather is what occurs at a particular moment in time - typically, precipitation and temperature. Due to natural variability, weather fluctuates from one hour to another, one day to another, one month to another, and one year to another. Climate, by contrast, is the long-run pattern of weather over time. To climate scientists, therefore, climate change has a very different significance from weather change. A change in weather is inherently short-run, while climate change is a shift in the long-run pattern. Because these are different phenomena, it is not surprising that they also have different economic implications.

Differences between weather and climate have implications for the choice of an economic metric. When dealing with climate, the appropriate metric is some measure of the long-run profitability of using the land for agriculture - the equivalent of the permanent income from farming. This is an economic rather than an accounting measure, and it generally has to be inferred through some proxy. One possible measure is farmland value, which presumably reflects long run profitability, and which has been used quite extensively since Mendelsohn et al. (1994) introduced the so-called Ricardian approach for assessing the impact of climate change on agriculture. With this approach, one estimates a Ricardian or hedonic regression relating farmland value to climate variables and to other variables that control for non-climate factors which might also affect farmland value. This approach is intended to capture the long-run impact of climate on farmland value, and it allows for farm-level adaptations that might be undertaken as climate varies.

In their paper, DG question whether there are unmeasured and omitted variables that also influence farmland value and that might be correlated with the climate variables in such a way as to bias their coefficient estimates in the hedonic regression. DG's important methodological innovation is to propose an alternative approach using fixed-effects to control for time-invariant idiosyncratic features of the county (or whatever is the unit of observation) within a panel data setting. However, for this approach to work, one cannot use climate variables as regressors since, in practice, these are likely to be fixed over the duration of the panel and, hence, perfectly collinear with the county fixed-effects. For this reason, DG use annual weather rather than long-run climate, since weather does vary over the course of the panel. Also, because of their annual focus, DG use annual "profit" (reported sales minus reported costs) as their metric of value rather than farmland value; this becomes the dependent variable in their regression. They are thus measuring the effect of weather on short-run profit rather than that of climate on long-run farming profitability or land value.[1] They find no statistically significant relationship between U.S. agricultural profits, proxied by sales-minus-costs as reported in county-level data of the 1987, 1992, 1997, and 2002 agricultural censuses, and weather variables in the same years. They also find no statistically significant relationship between yields (output per acre) of the major field crops corn and soybeans and weather. They conclude that if short-run weather fluctuations have no influence on agricultural profits or output, then in the long-run, when adaptations are possible, climate change is likely to have no impact or will even prove beneficial.

With any measurement strategy there are benefits and costs, and the ultimate effectiveness of the strategy is an empirical question.[2] The benefit of DG's strategy is that it is less vulnerable to unmeasured and omitted time-invariant factors. The cost is that it may measure something different from the impact of climate on long-run profit. DG argue that their measure overstates any possible long-run adverse impact of climate because it reflects the short-run response to fluctuations in weather and therefore does not allow for longer-run adaptation, which could only be less costly. But there are also many ways farmers cope with short-run shocks that would be more costly and/or less sustainable in the long run. For example, if the short-run response to a sudden increase in temperature is to pump more groundwater, this strategy may be less sustainable and/or more costly over the long run with a permanent increase in temperature than in the short run, due to depletion of the groundwater resource. In that case, the short-run impact of a fluctuation in weather would understate the long-run impact of a permanent shift in climate.

Besides the conceptual issues associated with DG's measurement strategy, there are serious questions about how they implement it and whether it actually produces the results they

[1] Kelly et al. (2005) regress annual county-level profits in Midwestern states on both climate averages as well as yearly normalized weather deviations from averages. Since the authors include climate averages, they cannot use fixed effects to account for time-invariant factors.

[2] Smith (2007) discusses the advantages and disadvantages of using "quasi-experiments" in environmental economics as well whether recent reports on climate change warrant a change in how we evaluate environmental policies.

claim. Perhaps most importantly, there are some unusual features of the data used by DG and their representation of climate change scenarios that appear to influence their results, in each case in a direction away from finding any potential negative impact of the change.

In our own research we have considered regression models that use both cross-sectional climate variations and time-series weather variations. In Schlenker et al. (2006) we show that a better-specified hedonic model that accounts for the influence of irrigation on farmland values is robust and predicts large negative impacts from projected climate changes. In Schlenker and Roberts (2008) we find a strong relationship between corn, soybean, and cotton yields and weather, a relationship that indicates extremely warm temperatures sharply reduce yields for all three crops.

Adaptation to warmer, or even extreme, temperatures, is suggested by DG and others, and this is of course possible, especially over time with the development of new crop varieties, but it is worth noting that we find no evidence of greater heat tolerance in yield regressions in warmer regions in the South as compared to cooler regions in the North, and no evidence that relative heat tolerance has grown over time. The relationship is strong and robust and very similar whether derived from time-series variations in weather or cross-sectional variations in climate and comparable in the cross-section of farmland values. Thus, while one cannot predict whether adaptations to extreme heat may occur in the future, there appears to have been little or no adaptation in the past (at least since 1950). Climate models, in turn, project that the frequency of extremely warm temperatures will increase significantly. Holding fixed the locations where crops are grown, we predict potential losses in yields for key crops of approximately 30-40% by the end of the century under a slow- warming scenario and 60-80% under the fastest warming, "business as usual", scenario. These predictions also accord with our research that uses the hedonic approach, where potential losses in farmland value range from approximately 30% to 70% for the same scenarios over the same time period.

What explains the stark differences between our empirical findings and those of DG? With regard to DG's results on profits and yields, we present evidence showing the difference stems from several sources: (i) coding and data errors in the weather data that magnify within-state temperature fluctuations by a factor of seven; (ii) an unusual and in our judgment incorrect characterization of climate change across the units of observation; (iii) differences in underlying climate change scenarios, in particular reliance by DG on an earlier and more optimistic climate projection than that used in the Fourth IPCC assessment and in our analysis; and (iv) DG's omission of storage, and perhaps other financial or technological mechanisms, that smooth their measure of short-run profits in the presence of weather-induced output fluctuations and cause the short-run impact of weather on profit to understate the long-run impact of a permanent shift in climate.

1 Data Irregularities

To investigate differences we downloaded DG's data and STATA code from the AER website. We found several irregularities in their weather and climate data. These data irregularities explain a large portion of the differences in findings.

DG have two weather variables in their data set: the variable *dd89*, which measures growing degree days for each year and county, and *dd89_7000*, which measures the average number of degree days in each county between 1970 and 2000.[3] These two variables do not appear consistent with each other. The correlation of the county-level average of the four-year panel (*dd89*) and the 31-year average given in their data (*dd89_7000*) is only 0.39. Given the wide variation in temperatures in the cross-section, one would expect a stronger correlation between the 4-year and 31-year averages across counties. We reconstruct the same weather variables from raw data sources and find a correlation of 0.99. We also find the average of *dd89* is much lower and the standard deviation much higher than in our replication.

Second, DG's baseline climate measure (*dd89_7000*) has a value of zero degree days for 163 counties. If correct, this measure implies temperatures do not exceed 8°C (46.4°F) in those counties during the growing season (April through September). Temperatures this low would seem implausible in any state, yet many of these counties are in warm southern states (e.g., Texas).

Anomalies caused by missing or incorrect measurements, which as we shall show have an important influence on estimated impacts of climate change, are illustrated in Figures 1 and 2. We independently calculate the degree days variable *dd89_7000* used by DG and display it in the bottom panel of Figure 1.[4] Note the much smoother pattern as compared to the large discontinuous changes in the top panel. Average temperatures vary smoothly across space, where counties of the same latitude tend to have comparable average temperatures

[3] Growing degree days integrate the product of temperature and time (measured in days) above a baseline temperature and below an upper threshold. For example, our (and DG's) baseline temperature is 8°C, so one day at a temperature of 13°C would equal 5 degree days. Time at temperatures above 32°C (A common upper threshold used in DG and in our replication) is treated as if it were 32°C, 24 degrees per day at or above this temperature.

[4] We derive degree days 8-32°C by first calculating average daily temperatures in each of the 2.5x2.5 mile grids of the data in Schlenker and Roberts (2006) and average over all cells with positive agricultural area.

that increase as one moves southward. Exceptions to this rule are mountain chains like the Rockies in the West or the Appalachians in the East, where temperatures are cooler due to gains in altitude. The discontinuous pattern induces incorrect weather variation, which has an especially large influence on parameter estimates in regression models that use state-by-year fixed effects. Within-state temperature deviations in our replicated data set are roughly one seventh of DG's.

Third, DG's predicted changes under warming scenarios are discontinuous in space and range from a decrease of 880 growing degree days (equivalent to a uniform 4.8°F decrease during the growing season) to a 6572 growing degree days increase (equivalent to a uniform 35.9°F increase). This pattern is odd given that the underlying climate model does not predict cooling anywhere in the U.S. and the variance of the projected changes far exceeds that of any climate model. Predicted changes in DG's model and in our replication are shown in Figure 2. Again, compare the discontinuities in the top with the more coherent patterns in the bottom.

The large variability of DG's predicted climate changes stems from the way they combine observed weather and climate-change forecasts. The difficulty arises from the fact that general circulation models (GCMs) generate climate predictions on a coarser geographic scale than data available in historic records. DG use historic county-level data as a baseline combined with climate predictions that are uniform across each state. Thus, after climate change, Los Angeles and San Francisco, Salinas and San Joaquin Valleys, Mount Whitney and Death Valley, are all assumed to have the same climate since all are in California. Much of the within-state variation, however, is maintained in the baseline values, which are county-level averages. Such a representation of climate change therefore displays regression towards the mean, with cooler counties becoming much warmer and some very warm counties becoming cooler.

This regression-toward-the-mean effect is accentuated by apparent errors in the baseline degree-day measure. Consider for example Fresno, Kings, and Tulare counties in the southern San Joaquin Valley of California. In DG's data, Fresno is predicted to have a *decrease* of 414 degree days (equivalent uniform temperate change of -2.3°F); Kings county has an increase of 403 degree days (+2.2°F) and Tulare an increase of 4685 degree days (+25.6°F). Tulare's large increase is the result of a zero (or apparently missing) baseline. But even for Kings and Fresno counties, for which there are no missing baselines, predicted climate changes are too different for bordering counties.

This treatment of climate change is unusual. We are not aware of any other application of the Hadley GCM model that predicts decreasing average temperatures by the end of the century in any U.S. location. Rather, the standard approach is to add regional predicted *changes* from the climate models to the sub-regional baselines, thereby preserving sub-regional variation and avoiding regression toward the mean.

2 Replication and Comparison

While there are differences between DG's and our own model of yields, much of the difference in our predicted impacts stems from the data issues described above. We generally find large negative projected climate impacts from replicated profit regressions as well, though results here are somewhat mixed and less likely to be significant, for reasons we discuss in the next section. Comparisons of the original and replicated yield and profit regressions are summarized in Tables 1 and 2.

In our replications we fix the observations so they exactly match those used by DG. This excludes some agriculturally important counties, which are missing in DG. For example, 66 of Iowa's 99 counties are missing from their data set, yet Iowa is the largest producer of corn and soybeans, in turn the nation's two largest crops. On the other hand, most of Nevada's counties are included, which are highly irrigated.

Irrigation poses a problem for estimation of the effects of climate variables both in a cross-section and a panel. In a cross-section such as the Ricardian or hedonic approach the problem is that since irrigation tends to be correlated with temperature and precipitation it can bias estimates if omitted, as we discuss later in the section on robustness. In a panel, the effect of weather fluctuations depends on water availability. DG deal with this problem by estimating regressions with separate coefficients for irrigated and rainfed counties. All farms within a county that is considered irrigated are pooled in the regression equation. The problem with this approach is that water rights can (and do in California's agricultural regions) vary considerably on a sub-county level. More fundamentally, in irrigated areas the water input comes from groundwater or from precipitation falling elsewhere, and local precipitation is not a valid measure of water supply. In such cases a better measure is access to irrigation water, on an irrigation district (sub-county) level, as employed in an application to California by Schlenker et al. (2007).

The main weather variables used by DG are growing degree days, growing-degree-days squared, precipitation, and precipitation squared. All models include soil controls, county fixed effects and year fixed effects, and the regressions are area-weighted.[5] To avoid confounding our comparison with changes in specification, we use the same variables and the same weighting, as well as the same observations.[6]

Columns (1a) and (2a) of Table 1 replicate results in Table 8 of DG using their original data set, for corn and soybeans, respectively. Our replication of DG differs slightly from DG's original due to a coding error in DG that we corrected (this is detailed in the appendix). Columns (1b) and (2b) replicate DG's regression model using our reconstruction of the weather data using their specification. Columns (1c) and (2c) include one additional variable, the square root of degree days above 34°C, to account for extreme temperatures. DG include such a variable to measure the potentially harmful effect of extreme heat on profits in Table 6 of their paper, but not in their yield regression. Columns (1c) and (2c) also use a slightly different calculation for degree days that accounts for within-day temperature variation.[7] For each model, the table reports the variance explained by the weather variables, model comparison statistics, and predicted climate-change impacts under Hadley II (IS92a) and Hadley III (B2) scenarios.[8]

Our recalculated weather data differ from DG in several ways and we highlight only a few here. The greatest differences likely stem from culling of missing values and relying on weather stations with relatively complete temperature and precipitation records. In DG and in columns (1a-b) and (2a-b), degree days are calculated using the daily *average* of the minimum and maximum. That is, all time within each day is treated as if it were fixed at the average. In degree-day calculations for columns (1c) and (2c), we account for variation between the minimum and maximum in each day, which matters given the non-linear relationship between temperature and degree days and the important influence of extremely warm temperatures on impacts.

For both corn yields (1a-c) and soybean yields (2a-c), the fit improves markedly going from (a) to (c) while predicted losses under climate change simultaneously become more

[5] Yield regression use cropland-area weights, while profit regressions use farmland-area weights. In the appendix we report results from the same exercise except state-by-year fixed effects replace year fixed effects in the yield regressions. The results are similar to those reported. State-by-year fixed effects influence results for the profit regressions, so those are reported here.

[6] Alternative specifications that fit the data better are presented by Schlenker et al. (2006) and Schlenker and Roberts (2008).

[7] Construction of our weather data from raw sources is detailed in Schlenker and Roberts (2006). These data utilize the distribution of temperatures within a day between minimum and maximum instead of just the average for the day. This variation can make a difference due to the way growing degree days are defined. In Schlenker and Roberts (2008) we generalize this notion of degree days and search for the best-fitting bounds.

[8] We calculate the variance explained by weather using the ratio of the residual variance from the full specification to the residual variance from a model with all controls but no weather variables.

severe. Specifically, columns (1a) and (2a) show that DG's degree days and precipitation variables explain 12.6% (corn) and 15.3% (soybeans) of the yield variance not explained by fixed effects. Our replication using corrected data but the same specification (1b and 2b) explains about twice the variance, a strong indication of data errors in DG. Predicted losses for corn under climate change also increase from a statistically insignificant 1 percent to a strongly significant 11.5 percent (SE of 1.71 percent) under the Hadley II scenario used by DG. We report both the standard error following DG's regression as well as a standard error that adjusts for the spatial correlation of the error terms while accounting for the panel structure following Conley (1999). Predicted losses become more extreme under the Hadley III scenario. Under Hadley III losses increase to 44.5 percent (SE of 3.7 percent). While a discussion of the validity or the accuracy of either model is beyond the scope of this paper, it should be noted that earlier research against which DG compare their results relied on the Hadley III model, and we believe an appropriate comparison should leave the climate forecasts unchanged.[9]

Similar but somewhat more adverse predictions are reported for soybeans. In columns (1c) and (2c), we add extremely warm temperatures (the square root of degree days above 34C) and account for within-day temperature variation. Here the fit increases for both crops by about 25 percent relative to the (b) columns and predicted damages increase by nearly 50 percent. Using the (c) column model and Hadley III climate change scenarios predicted losses are 65.6 percent for corn and 75.7 percent for soybeans.

The second and third rows of the table present pair-wise non-nested J-tests to compare the different models. DG's model with their data is rejected when compared against either of our replications with new data, with t-statistics in excess of 15. Neither replication can be rejected when compared to DG's original model.

In columns (1a-c) of Table 2 we present the same set of comparisons for the profit regressions. A similar pattern emerges, with both the fit improving and projected damages generally increasing as we change the weather data, account for extreme temperatures, account for within-day temperature variations, and move from Hadley II to Hadley III climate projections. The overall explanatory power of weather for profits is poor, however. DG's weather data and model explain just 0.5 percent of the residual variance after removing

[9]The Hadley II scenarios were developed for use by the IPCC's Third Assessment Report; the Hadley III scenarios were developed for the IPCC's Fourth Assessment Report. Hadley III is an update and refinement of Hadley II. The key differences are that Hadley III projects a sharper temperature increase in North America, especially in summer, and a less optimistic forecast of changes in precipitation. In subsequent research DG have switched to the Hadley III model (Deschênes and Greenstone 2007a).

fixed effects. Our new data nearly triples that performance to 1.4 percent, and it increases slightly to 1.5 percent with an accounting for extreme temperatures and within-day temperature variations. Projected damages range from -6 percent using DG's data under the Hadley II climate change scenario to -53 percent using the same model, our replicated data, and the Hadley III scenario. While these estimates are statistically significant, the standard errors are economically large, equal to over 20 percent for the Hadley III predictions when our new data are used and we account for spatial correlation.

In columns (2a-c) of Table 2 we present the same set of results in columns (1a-c) with one critical difference: we use state-by-year fixed effects in place of state fixed effects. While DG find insignificant impacts in their original paper using both year fixed effects and state-by-year fixed effects, the two diverge in our replication. We find significant damages in the former and insignificant impacts under the latter under all weather data sets and climate change scenarios, although the confidence intervals of the latter are very wide. As explained in footnotes 4 and 5 of DG, state-by-year fixed effects have the advantage of capturing regional price effects, which is especially useful if production of certain crops is concentrated geographically. For example, California produces 85 percent of the lettuce grown in the U.S. A country-wide yearly fixed effect would not capture the fact that crops specific to California might face unique price shocks. However, any crop-specific price response works as natural "insurance" for farmers that grow the crop. Prices move in the opposite direction from production shocks: If yields decline, prices increase, and vice versa.[10] Accounting for region-specific price responses should therefore make predicted impacts more *negative* as it cancels out the counterbalancing price response. It is counter-intuitive that predicted changes in profits are negative and significant in a regression using year fixed effects, yet turn insignificant when one includes state-by-year fixed effects to capture region-specific price responses. What other effects apart from regional price effects might explain why the results become less damaging and insignificant with the use of state-by-year fixed effects?

A concern with the use of state-by-year fixed effects is that they absorb a significant amount of weather variance. After removing county and state-by-year fixed effects, remaining weather variance pertains only to yearly within-state deviations from county means, as for example the amount by which northern Iowa is warmer than normal in a given year compared to how much southern Iowa is warmer than normal in the same year. Generally, whenever northern Iowa is warmer than normal, so is southern Iowa, because temperatures

[10] This is especially true for speciality crops like lettuce where world trade is limited in volume and market quantity is given by production in a confined regional market.

vary smoothly in space. DG report a significant amount of within-state weather variation in their Table 2. But it turns out this variation is largely an artifact of errors in their weather data, which exhibit large discontinuous shifts across neighboring counties as discussed above.

Statistics that summarize weather variation in DG's data set and our own are reported in Table 3. The table summarizes regressions of degree days against different sets of fixed effects: (1) an intercept; (2) county fixed effects; (3) county plus year fixed effects; and (4) county plus state-by-year fixed effects. For each data set and set of fixed effects the table reports the R-square, the standard deviation of the residual weather variation not absorbed by the fixed effects (in Fahrenheit equivalent),[11] and the fraction of residuals with an absolute value greater than 1 degree Fahrenheit. While the overall standard deviations of temperatures are relatively similar in DG's measure and our replication (6.82F versus 6.07F), the two diverge drastically once we include fixed effects. The residual standard deviation of DG's temperature measure is 2.69F with county fixed effects and 2.38F with county plus state-by-year fixed effects. Our measure, on the other hand, has a residual standard deviation of 1.49F with county fixed effects and just 0.35F with county plus state-by-year fixed effects. These differences suggest a noise to signal ratio of DG's weather of about 7 to 1 in their preferred fixed-effects model.[12]

3 Profit and the Role of Storage

The preceding section shows that predicted yield impacts from climate change are negative and significant if improved weather data, a measure of extreme temperatures, and more recent climate change predictions are used in the estimation. The results from the profit regressions are mixed but weigh on the negative side, and have larger standard errors that include both sizable damages or sizable benefits. In this section we consider potential problems arising from the use of current profits to measure long-run economic impacts from weather. Our premise is that this measure omits storage and perhaps other financial or technological mechanisms that likely smooth this measure of short-run profits in the presence of weather-induced fluctuations and lead the short-run impact of weather on profit to understate the long-run impact of a permanent shift in climate.

DG construct their estimate of profit by subtracting annual production expense from

[11]We divide degree days by the number of days during the growing season.

[12]Recall that our replication of their degree days measure explains roughly twice as much of the variation in year-to-year crop yields (Table 1). Thus it is unlikely that our replication smoothes temperatures too much in space as it is superior at explaining yield variation in space.

annual farm revenues. While production expense is essentially the cost associated with the crops grown in that year, the farm revenue is not necessarily the revenue from crops grown in that year - it is revenue from crops *sold* in that year. In high-yielding years farmers accumulate stocks for the major field crops in the U.S. (corn, soybeans, and wheat, among others); in low-yielding years they deplete stocks accumulated in earlier years. Storage is thus one way for farmers to smooth weather-related shocks over time. It also creates a substantial disconnect between the weather-related shock and DG's metric for the impact of that shock, sales minus reported costs.

The amount farmers choose to place into or remove from storage is part of the error in the profit regressions. This error is directly related to the yield shock, and thus correlated with weather, DG's key explanatory variable. This creates an endogeneity bias toward zero, because storage is greater and sales lower in good years with positive weather shocks, and inventories are depleted in bad years with negative weather shocks. Specifically, if q_s is the quantity sold, q_p the quantity produced, and q_n is the quantity stored, then $q_s = q_p - q_n$ and $q_s p = q_p p - q_n p$. In a regression of sales ($q_s p$) on the value of production ($q_p p$) we have

$$q_s p = \beta_0 + \beta_1 q_p p - \underbrace{q_n p + \epsilon}_{\nu}$$

Since ν encapsulates storage it is correlated with q_p (when production is low, so is storage, and vice versa), the coefficient β_1 will be biased downward from 1. We report some estimates of the magnitude of this effect below.

Other factors besides storage could cause the short-run response to weather to understate the long-run response to climate. For example, if the short-run response to a sudden increase in temperature is to pump more groundwater, this strategy may be less sustainable and/or more costly over the long-run permanent increase in temperature as compared to the short-run transitory increase, due to depletion of the groundwater resource. Another reason could be that after a bad yield shock, a livestock producer expects higher future feed prices, and therefore chooses to slaughter breeding stock in anticipation of higher future costs. Such a reduction in cattle inventories could temporarily increase sales in a way that would not be feasible in the long run (Rosen et al. 1994).

The important point is that, for many reasons, sales-minus costs do not, in general, reflect the full economic impact of the current weather shock. Or, put another way, there may be technologies like storage, irrigation, and liquidation of capital or inventories that act to smooth profits in the face of transitory shocks. Since climate change would be permanent, such short-term smoothing technologies would not be available. DG's claim that the short-run response to weather overestimates what would be the long-run response to climate change cannot therefore be generally correct.

The most plausible argument against these dynamic considerations of storage, irrigation, and livestock inventory adjustments is DG's use of year or state-by-year fixed effects to control for prices.[13] In the competitive storage model (Williams and Wright 1991, Deaton and Laroque 1992), for example, the motive to accumulate or deplete inventories comes entirely from speculation about prices. In bad-weather years prices increase, giving farmers an incentive to deplete inventories and in good-weather years prices fall, giving farmers an incentive to replenish their inventories. The essential question is whether year or state-by-year fixed effects can suitably control for these dynamic considerations, all of which would tend to bias DG's profit regression toward zero weather effects.

There are theoretical reasons why year or state-by-year fixed effects *cannot* fully account for these factors. One reason is that there are *many* crops and agricultural activities, and prices for agricultural products do not move in perfect unison, so it is unlikely that the fixed effects should capture all price-feedback effects stemming from weather shocks. A second reason is that there exist local sub-state yield and price fluctuations not accounted for by the fixed effects. Indeed Brennan et al. (1997) present evidence showing so-called "convenience yields"–a motive to store commodities even when the futures price is below the spot price–may stem from large local price variations for which futures markets do not exist. While such local price variation may be small relative to price fluctuations overall, within-state weather variations in a given year exploited in DG's regression model are surely much smaller.[14] And if it is not local price effects that explain convenience yields, it must be that "convenience" does. The price of convenience, whatever it may be, would be low in times of relative plenty and high in times of relative scarcity, leading back to the same attenuation bias even if fully controlling for prices.

To examine this issue empirically, we conducted the following exercise, the results of which are reported in Table 4. In this exercise we regress *sales* against the *value of production*, together with DG's preferable specification that includes county and state-by-year fixed effects. The idea is that, if storage variations are fully accounted for in the model, the

[13]Both the hedonic approach and the approach taken by DG assume constant prices, and thus assume no consumer-related impacts. In effect, the goal in both approaches is to obtain a first-order approximation of the economic impact by assessing the potential impacts on fundamental productivity. In the long run, we have no idea about how prices will adjust (Cline 1996).

[14]Recall that the specification that uses state-by-year fixed effects in Table 3 has a standard deviation of 0.35 degrees Fahrenheit for the remaining weather variation.

coefficient on production value should be one. Sales should increase one-for-one by the value of each extra unit that is produced. Recall that DG's profit measure relies on *sales* in a given year to measure changes in productivity. We conduct this exercise for all crops for which the Census reports production and sales measure. Table 4 includes results for the largest four crops (corn, soybeans, wheat, and cotton) and for the aggregate of the eight crops for which data are available (the four large crops plus oats, sorghum, barley, and tobacco). If this test rejects the idea that sales proxy for the value of production, then it surely fails for aggregate production as well.[15]

Note that our proxy for production value still excludes within-state price variation that we do not observe. This biases our coefficients upward, the opposite of attenuation bias, because it makes the value of production *less variable* than it is in reality. This bias therefore could cause the estimated coefficients to suggest state-by-year fixed effects account for storage even when they do not. Despite this bias, the analysis indicates fixed effects do not fully account for storage: the coefficients are uniformly and significantly less than one.[16]

State-by-year fixed effects do account for some of the tendency to store yield shocks. The coefficients on corn and soybeans are lower if we use year fixed effects instead (Results are given in the appendix). We also see a larger coefficient on cotton, for which on-farm storage can be quite costly. But even with these controls, the coefficients are sufficiently less than one that they can have a significant impact on the imputed profits, as margins (sales minus production expenses) are low in agriculture. Production expenses in the 2002 census equaled 86% of agricultural sales, which is larger than the coefficient we get in the first column of Table 4. In the profit regressions with state-by-year fixed effects, the within-state county-level year-to-year weather variations are small, so the bias from storage could severely confound this weak source of identification.

[15] To conduct this exercise we obtained access to individual farm-level data from the Agricultural Census Micro-files. These are the raw data used to construct the aggregate sales data used in DG. The survey asks farmers to report their sales of eight crops: corn, soybeans, wheat, cotton, oats, sorghum, barley, and tobacco. It also asks for the yield of theses crops and we derive the value of production by multiplying the total production of each farm by state-level prices (reported by USDA-NASS). We then aggregate all farms in a county and follow DG's specification as closely as possible: we construct sales per acre in a county (total sales divided by the total acreage) as the dependent variable and the ratio of total production divided by total acreage as the exogenous variable, plus county and state-by year fixed effects. The dependent variable in DG is profit per acre, sales minus costs. The regression results use area-weights as does DG.

[16] In a sensitivity check shown in the bottom rows of Table 4 we exclude farms that own livestock. Such farms might use crops (especially corn and soybeans) as feedstock and hence these crops would not be included in the sales figures. The results remain robust.

4 Robustness

In the first half of their paper DG argue that results from the cross-sectional hedonic approach are not robust. Their contention is that cross-sectional studies rely on climate variations that are too closely associated with unobserved factors relating to location, and thus likely to be biased. DG's standard of robustness is the consistency of parameter estimates over different subsets of the data. However, when DG assess the robustness of the hedonic model, the weather variables they use, following Mendelsohn et al. (1994), are average monthly temperature and precipitation for January, April, July and October[17]; when they demonstrate the robustness of their fixed-effects model, the weather variables they use are summed precipitation and degree days over the growing season, as developed and used by Schlenker et al. (2006) in their hedonic analysis, and shown there to be superior on both economic and econometric grounds. The difference in the representation of weather confounds DG's comparison of robustness.

For these reasons, we believe it is appropriate to repeat DG's tests of robustness, shown in their Figure 3, using the degree days representation of temperature and applying the tests to a hedonic model. This analysis is summarized in our Figure 3. The figure shows predicted impacts from climate change using the hedonic model specified in Schlenker et al. (2006), estimated using various sets of control variables, in various census years. We replicate this analysis for farms east of the 100th meridian, an approximate boundary between rainfed and irrigated agriculture in the U.S., comprising 80% of county observations. Farms west of the 100th meridian rely on heavily subsidized irrigation water that is capitalized into farmland values. Some areas east of the 100th meridian also rely on irrigation. For example, 79 percent of the corn acreage in Arkansas was irrigated in 2007. The main difference is that access to water, where it exists at all, is much less heavily subsidized. Cline (1996) emphasizes that a hedonic approach in which observations are pooled assumes that farmers can obtain irrigation at existing marginal cost, which seems unrealistic. The bias induced by pooling observations from areas with subsidized irrigation with areas not primarily dependent on (subsidized) irrigation is demonstrated in Schlenker et al. (2005). DG assert that Schlenker et al. (2005) force the coefficient on irrigated and nonirrigated areas to be the same. The assertion is incorrect - the main point of that paper is that a *different* model is needed to deal with irrigated agriculture since the water supply there does not depend on local

[17]DG write that they are replicating Schlenker et al. (2005), but this is incorrect. They critique Schlenker et al.'s (2005) replication of Mendelsohn et al. (1994).

precipitation.[18]

Panel A in Figure 3 includes only climate variables with no other controls. If the climate variables are correlated with other variables, such as soil quality, the coefficient on the climate variables will be biased. Accordingly, panel B includes controls for both soil and socio-economic variables to avoid potential omitted variable problems. Finally, panel C additionally includes state fixed effects. Columns marked with a [0] are unweighted regressions, while columns marked with a [1] are regressions weighted by the square root of acres of farmland.

All estimates indicate strong negative impacts from climate change, even when all independent variables except for the climatic variables are dropped (panel A). Based on the robustness tests suggested by DG, a hedonic model using degree days and precipitation summed over the growing season and applied to counties not primarily dependent on irrigation clearly appears to be robust. The model also passes a wide array of additional robustness tests discussed in Schlenker et al. (2006).

5 Conclusions

Agriculture is the sector that has been most extensively studied in attempts to predict the economic impact of global climate change. This is not surprising, since climate variables such as temperature and precipitation directly enter agricultural production functions. Despite the existence of a large and growing literature, economists do not appear to have reached a consensus on the potential magnitude of the impact, or even on its sign. Our own research, involving a series of studies estimating the relationships between climate variables and farmland values, and between both climate and weather variables and crop yields, strongly suggests that the impact on U.S. agriculture over a time horizon stretching toward the end of the century is likely to be negative. Results are significant, robust, and consistent across data sets and methods.

This prediction does not account for some potential mitigating factors, such as development of crops with increased resistance to extreme heat, and CO_2 fertilization. We have however noted that evidence to date for development of crops exhibiting increased resistance is lacking. With respect to CO2 fertilization, evidence from recent experiments that more realistically simulate fertilization suggests that its impact on crop yields will be much more

[18] DG may again be confusing the replication of Mendelsohn et al. (1994) with the new material in Schlenker et al. (2005). The model introduced there is estimated solely for counties for which agriculture is rainfed, i.e. not dependent on irrigation.

limited than previously believed (Long et al. 2006)). Further, other experiments suggest that at least some of the projected increase in yields may be offset by a decline in nutritional value (Jablonski et al. 2002).

A recent study by Deschênes and Greenstone (2007b) reports dramatically different results: estimated relationships between weather variables, yields, and profits, are taken to imply that the impact of long run climate change on U.S. agriculture will be either insignificant or modestly beneficial. In this paper we try to reconcile the conflicting results by revisiting the several studies.

Likely explanations for the divergence between our findings and the contrary results reported by DG include: (1) missing and almost certainly incorrect weather and climate data in DG's study, amplified by the use of state-by-year fixed effects that absorb most year-to-year weather variation; (2) their unusual and in our judgment incorrect treatment of climate-change predictions; (3) their use of the older Hadley II climate model for climate change predictions rather than the more recent and less optimistic Hadley III climate model used in our studies as well as others in the literature; and (4) theoretical difficulties in their profit-based approach due to the confounding effects of storage and possibly also capital and inventory adjustments and local price movements that are associated with weather fluctuations. A careful account of these factors shows that the balance of evidence weighs heavily on the side of severe adverse potential impacts to U.S. agriculture by the end of the century stemming from anticipated global warming. This conclusion is subject to the possible mitigating influences of new heat-resistant crop varieties and carbon fertilization, though we have suggested that there is probably more limited scope for each of these than often claimed.

Conceptually, DG are correct in noting that omitted variables can in principle cause bias in a hedonic regression and that fixed effects can control for time-invariant idiosyncratic features of the unit of observation, in this case the county. However, it is also possible that fixed effects can increase the bias due to omitted variables if *time-varying* omitted variables are more strongly correlated with the treatment than *time-invariant* omitted variables that have been removed via the fixed effects. These fixed effects increase bias stemming from both endogeneity and measurement error. We have identified some important data errors and time-varying omitted variables, like storage, that are strongly correlated with both weather (the treatment variable) and DG's dependent variable, reported sales minus reported expenditures. These data errors and omitted variables bias toward zero results obtained by regressions that use sales as a proxy for production value.

References

Brennan, Donna, Jeffrey Williams, and Brian D. Wright, "Convenience Yield without the Convenience: A Spatial-Temporal Interpretation of Storage Under Backwardation," *Economic Journal*, July 1997, *107* (443), 1009–1022.

Cline, William R., "The Impact of Global Warming of Agriculture: Comment," *American Economic Review*, December 1996, *86* (5), 1309–1311.

Conley, Timothy G., "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, September 1999, *92* (1), 1–45.

Deaton, Angus and Guy Laroque, "On the Behaviour of Commodity Prices," *Review of Economic Studies*, January 1992, *59* (1), 1–23.

Deschênes, Olivier and Michael Greenstone, "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US," *NBER working paper 13178*, June 2007.

___ and ___ , "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather," *American Economic Review*, March 2007, *97* (1), 354–385.

Jablonski, Leanne M., Xianzhong Wang, and Peter S. Curtis, "Plant Reproduction under Elevated CO2 Conditions: A Meta-Analysis of Reports on 79 Crop and Wild Species," *New Phytologist*, October 2002, *156* (1), 9–26.

Kelly, David L., Charles D. Kolstad, and Glenn T. Mitchell, "Adjustment costs from environmental change," *Journal of Environmental Economics and Management*, November 2005, *50* (3), 468–495.

Long, Stephen P., Elizabeth A. Ainsworth, Andrew D. B. Leakey, Josef Nosberger, and Donald R. Ort, "Food for Thought: Lower-Than-Expected Crop Yield Stimulation with Rising CO2 Concentrations," *Science*, 30 June 2006, *312*, 1918–1921.

Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw, "The Impact of Global Warming on Agriculture: A Ricardian Analysis," *American Economic Review*, September 1994, *84* (4), 753–771.

Rosen, Sherwin, Kevin M. Murphy, and Jose A. Scheinkman, "Cattle Cycles," *Journal of Political Economy*, June 1994, *102* (3), 468–492.

Schlenker, Wolfram and Michael J. Roberts, "Estimating the Impact of Climate Change on Crop Yields: The Importance of Nonlinear Temperature Effects," *NBER Working Paper 13799*, February 2008.

——— **and Michael Roberts**, "Nonlinear Effects of Weather on Corn Yields," *Review of Agricultural Economics*, 2006, *28* (3), 391–398.

——— **, W. Michael Hanemann, and Anthony C. Fisher**, "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach," *American Economic Review*, March 2005, *95* (1), 395–406.

——— **, ——— , and ———** , "The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions," *Review of Economics and Statistics*, February 2006, *88* (1), 113–125.

——— **, ——— , and ———** , "Water Availability, Degree Days, and the Potential Impact of Climate Change on Irrigated Agriculture in California," *Climatic Change*, March 2007, *81* (1), 19–38.

Smith, V. Kerry, "Reflections on the Literature," *Review of Environmental Economics and Policy*, Summer 2007, *1* (2), 300–318.

Williams, Jeffrey and Brian Wright, *Storage and commodity markets*, Cambridge, UK; New York: Cambridge University Press, 1991.

Figure 1: Baseline Climate in Deschenes and Greenstone

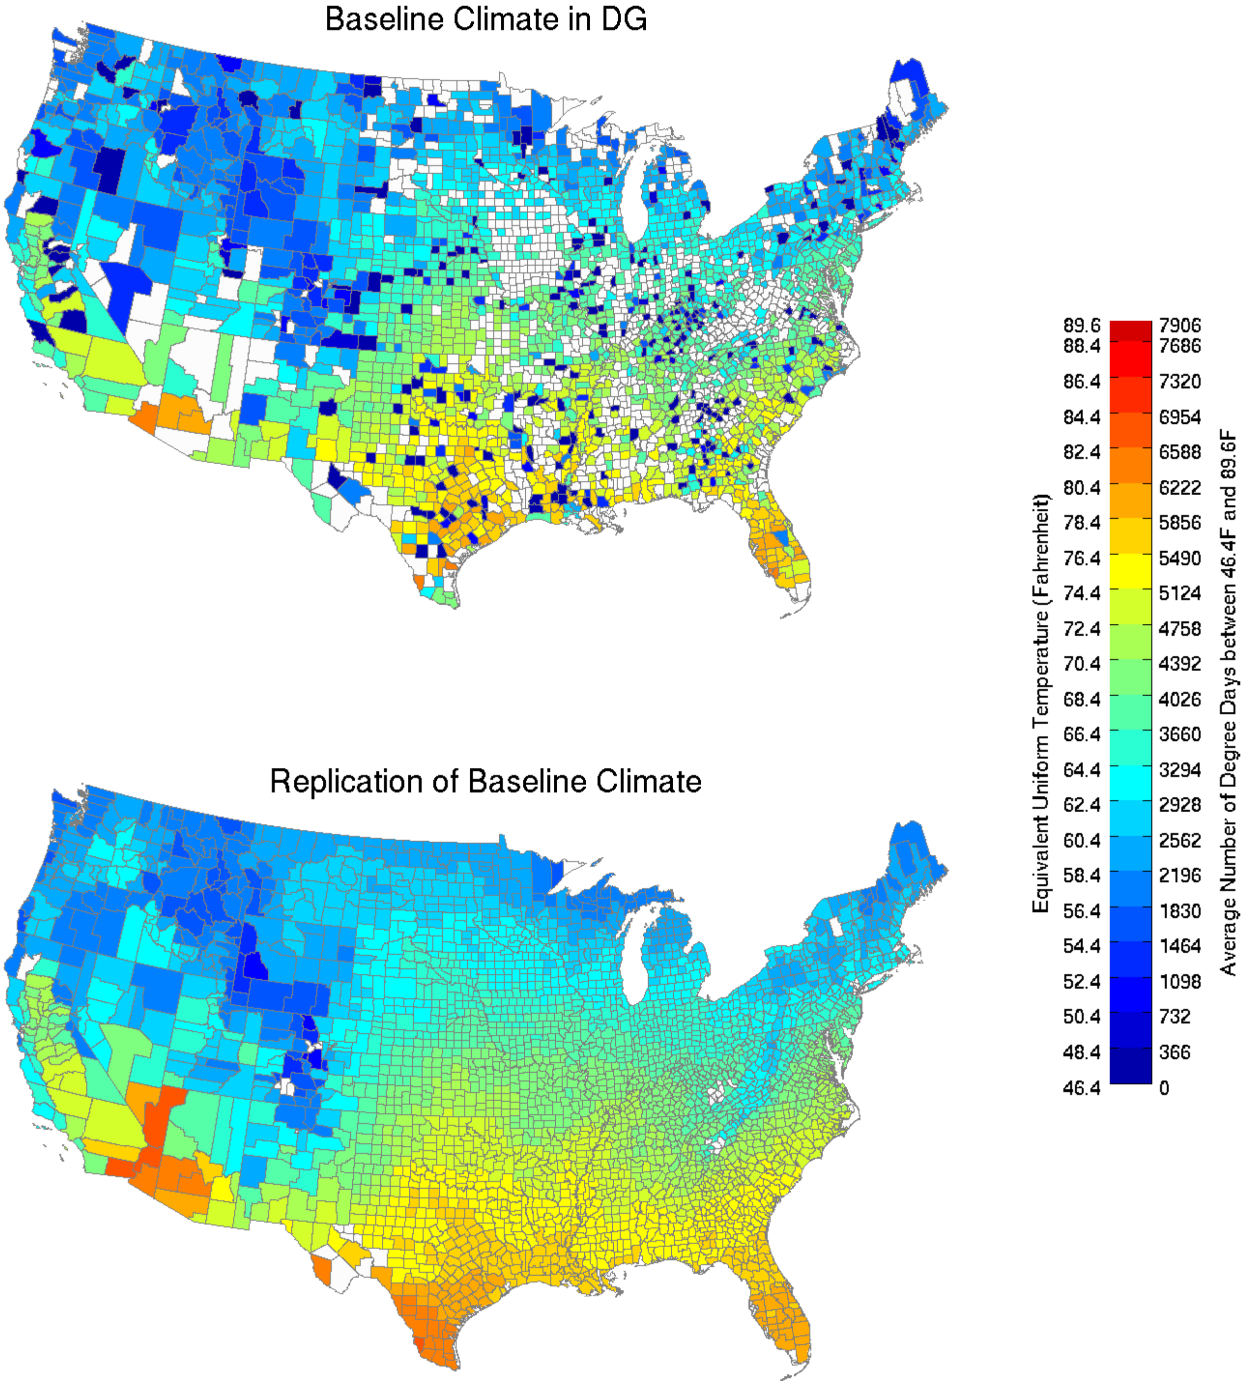

Notes: Graphs displays the baseline climate. The top panel are the data used in Deschenes and Greenstone, the bottom panel show our replications of the same variable degree days 8-32°C. The right index of the legend shows the number of degree days. Since degree days are difficult to interpret, we added another index at the left of the legend that shows the equivalent uniform temperature in degrees Fahrenheit, i.e., the equivalent constant temperature that would give the same number of degree days.

Figure 2: Climate Change Predictions in Deschenes and Greenstone

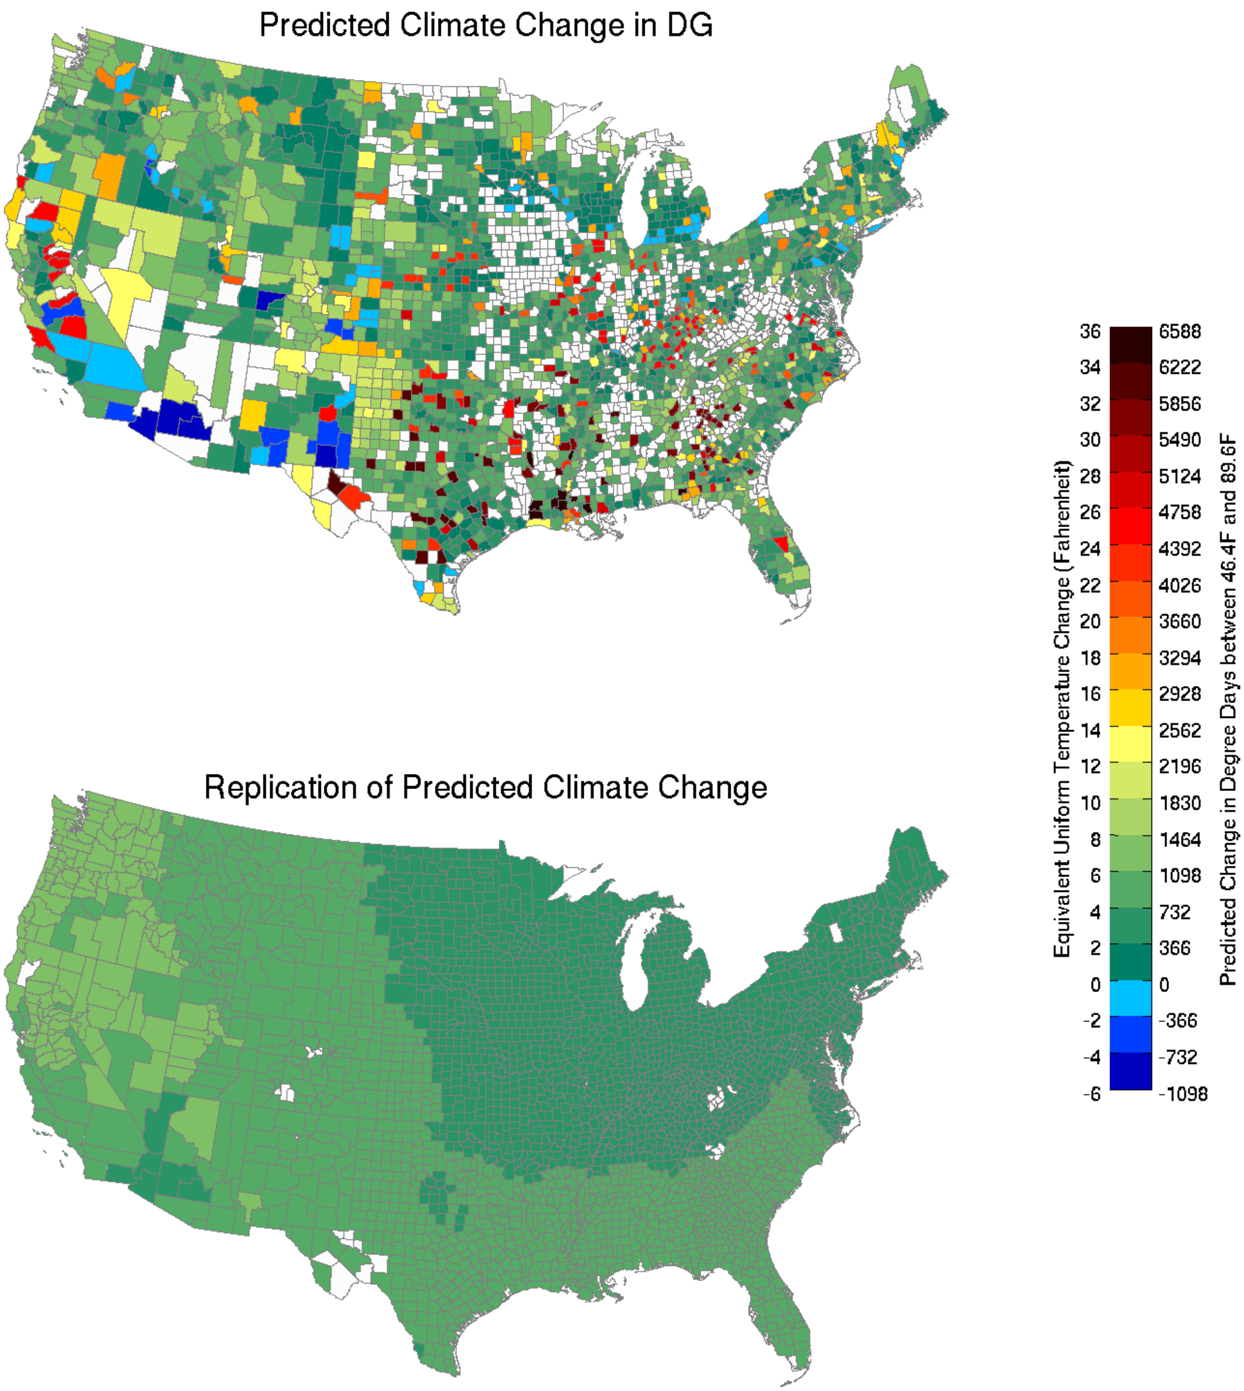

Notes: The top panel are the data used in Deschenes and Greenstone, the bottom panel shows our replications of the predicted changes in the same variable degree days 8-32°C. The right index of the legend shows the predicted change in the number of degree days. Since degree days are difficult to interpret, we added another index at the left of the legend that shows the equivalent uniform temperature change in degrees Fahrenheit.

Figure 3: Robustness of Hedonic Model using Degree Days

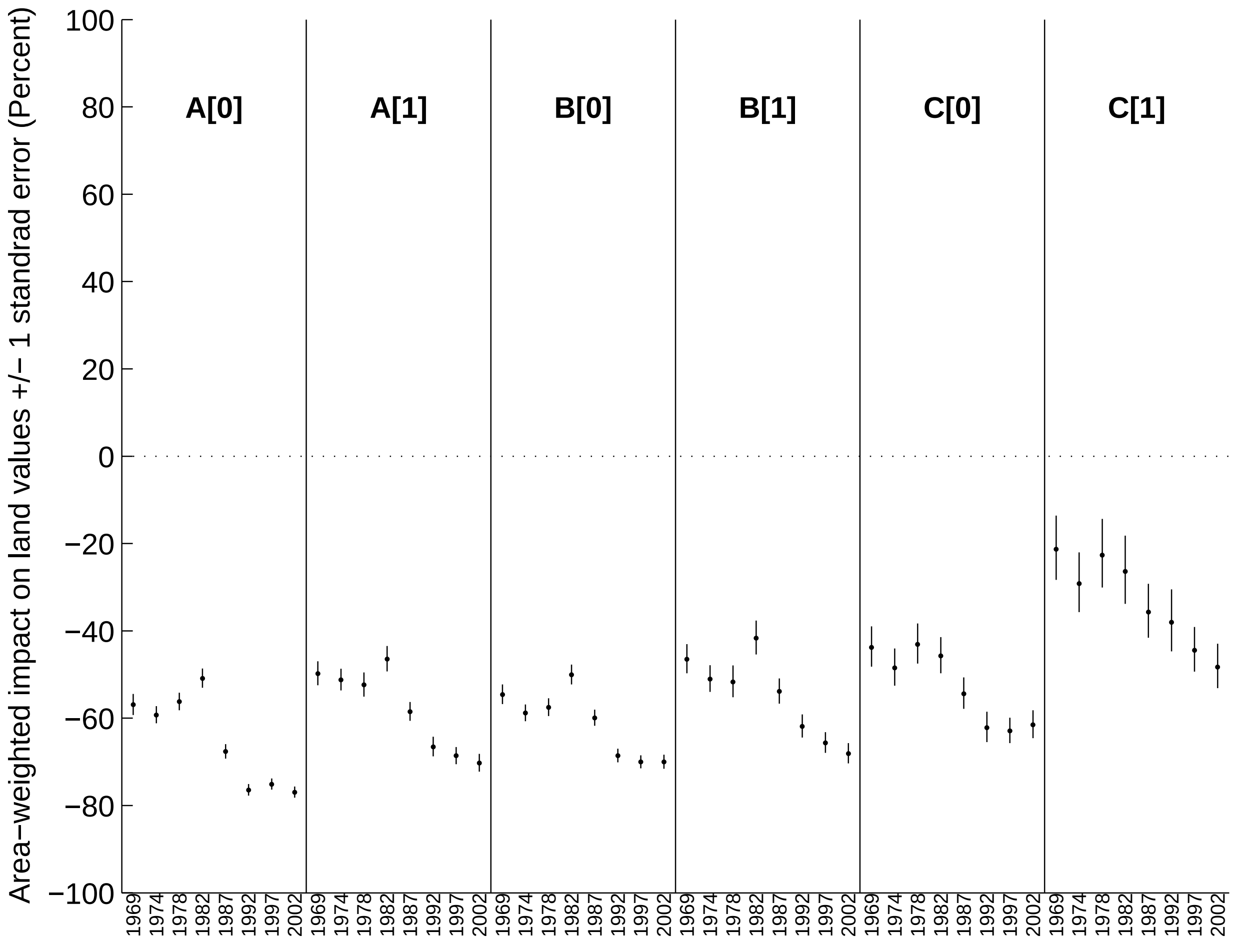

Notes: This is a replication of Figure 3 in Deschenes and Greenstone using the degree days from Schlenker et al. (2006) instead of average monthly temperatures. Panel A only includes climate variables with no other controls. Panel B includes controls for both soil and socio-economic variables to avoid potential omitted variable problems. Panel C additionally includes state fixed effects. Column marked with a [0] are unweighted regressions, while columns marked with a [1] are regressions weighted by the square root of acres of farmland. The x-axis lists the year in which the cross-section is estimated.

Table 1: Comparison of Various Data Sources in Yield Regressions

	Corn			Soybeans		
	(1a)	**(1b)**	**(1c)**	**(2a)**	**(2b)**	**(2c)**
Regression diagnostics						
Variance explained by weather	12.6%	21.3%	26.3%	15.3%	31.0%	38.1%
Model comparison tests (J-test)						
DG against other weather (t-value)		15.97	21.08		19.98	25.80
Other weather against DG (t-value)		1.85	1.66		1.59	1.16
Climate change impact (Percent)						
Hadley II-IS92a scenario	-0.98	-11.50	-14.03	-3.01	-16.49	-18.36
(s.e.)	(1.23)	(1.71)	(1.66)	(1.38)	(1.73)	(1.64)
[s.e. Conley]	[2.12]	[5.57]	[5.18]	[1.96]	[4.93]	[4.39]
Hadley III-B2 scenario		-44.46	-65.64		-53.90	-75.71
(s.e.)		(3.69)	(3.97)		(3.87)	(3.92)
[s.e. Conley]		[13.09]	[11.76]		[11.94]	[10.60]
Observations	6862	6862	6862	5141	5141	5141
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table summarizes and compares alternative regression models for the relationship between crop yields and weather. Corn yields are the dependent variable in the first three columns and soybean yields are the dependent variables in the next three columns. Columns (a) replicate the results in DG using their code and data (a quadratic in degree days 8-32°C and precipitation); columns (b) are the same models as (a) estimated with our reconstructed data; columns (c) account for within-day temperature variation and extremely warm temperatures (square root of degree days above 34°C as an additional variable). The variance explained by weather is the residual variance of the full specification over the residual variance of the model excluding weather. The non-nested J-tests compare models, taking DG's model (columns a) to be best as the null hypothesis and vice versa.

Table 2: Comparison of Various Data Sources in Profit Regressions

	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
Regression diagnostics						
Variance explained by weather	0.5%	1.4%	1.5%	0.5%	0.7%	0.9%
Model comparison tests (J-test)						
DG against other weather (t-value)		4.89	4.13		2.97	3.90
Other weather against DG (t-value)		0.44	0.63		0.72	0.65
Climate change impact (Percent)						
Hadley II-IS92a scenario	-6.08	-34.49	-29.53	3.92	0.52	3.86
(s.e.)	(3.03)	(6.90)	(6.63)	(2.87)	(12.85)	(12.86)
[s.e. Conley]	[4.07]	[12.25]	[12.14]	[2.56]	[12.53]	[12.80]
Hadley III-B2 scenario		-52.91	-41.70		-4.47	5.50
(s.e.)		(11.24)	(11.30)		(20.24)	(20.72)
[s.e. Conley]		[20.04]	[21.54]		[19.76]	[20.85]
Observations	9024	9024	9024	9024	9024	9024
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	No	No
State-By-Year FE	No	No	No	Yes	Yes	Yes

Notes: The table summarizes and compares alternative regression models for the relationship between agricultural profits (sales minus costs) and weather. Columns (a) replicates the results in DG using their code and data (a quadratic in degree days 8-32°C and precipitation); column (b) is our replication of DG's data (also a quadratic in degree days 8-32°C and precipitation); columns (c) use daily minimum and maximum temperatures instead of averages and includes the square root of degree days above 34°C as an additional variable besides the quadratic in degree days 8-32°C and precipitation. The first three columns and second three columns differ only in that the second three use state-by-year fixed effects in place of year fixed effects. The variance explained by weather is the residual variance of the full specification over the residual variance of the model excluding weather. The non-nested J-tests compare models, taking DG's model (columns a) as best as the null hypothesis and vice versa.

Table 3: Temperature Variation under Various Sets of Fixed Effects

	Variable dd89 in DG			**Replication of dd89**		
	$\mathbf{R}^2$	$\boldsymbol{\sigma_e}$	$\lvert\mathbf{e}\rvert > \mathbf{1F}$	$\mathbf{R}^2$	$\boldsymbol{\sigma_e}$	$\lvert\mathbf{e}\rvert > \mathbf{1F}$
	(1a)	**(1b)**	**(1c)**	**(2a)**	**(2b)**	**(2c)**
No Fixed Effects (F.E.)		6.82F	91.1%		6.07F	89.8%
County F.E.	0.845	2.69F	56.5%	0.940	1.49F	64.7%
County + Year F.E.	0.867	2.48F	54.8%	0.979	0.88F	24.2%
County + State-by-Year F.E.	0.879	2.38F	50.5%	0.997	0.35F	1.3%

Notes: Table regresses degree days on various sets of fixed effects and examines the residuals, i.e., the remaining variation in the temperature measure. The first three columns uses the variable dd89 from DG, while the last three columns use our recalculation of the same variable when data errors are corrected. Columns (a) report the R-square of the regression, columns (b) report the standard deviation of the residuals (remaining temperature variation) in degrees Fahrenheit during the growing season, and columns (c) report what fraction of the observations have a residual that is larger than 1 degree Fahrenheit over the growing season.

Table 4: Regressing Sales on Value of Production

	Combined (1)	**Corn (2)**	**Soybeans (3)**	**Wheat (4)**	**Cotton (5)**
Using all farms in a county					
Coefficient	0.780***	0.738***	0.770***	0.868***	0.923***
Std. Error	(0.010)	(0.012)	(0.007)	(0.010)	(0.019)
Observations	11595	10268	8241	10540	2573
Using farms with no livestock					
Coefficient	0.821***	0.772***	0.793***	0.871***	0.931***
Std. Error	(0.016)	(0.015)	(0.009)	(0.012)	(0.021)
Observations	11083	9303	7569	9490	2207

Notes: The table summarizes regressions of total sales on total value-of-production. Each coefficient and standard error is from a separate regression. All regressions include county and state-by-year fixed effects. The first column uses the combined sales of all crops where sales are reported in the Micro-files: corn, soybeans, wheat, cotton, oats, sorghum, barley, and tobacco. The next four columns report individual regressions for the four most important commodities. The first three rows aggregate all farms in a county. The second three rows only aggregate farms that do not have any livestock. Significance levels whether coefficients are different from 1 are indicated by: *** (1%); ** (5%); and * (10%).

Appendix

A1 Coding Issues

As outlined in the main paper, the variable dd89, which measures degree days for each year and county, and dd89_7000, which measures the average number of degree days in each county between 1970 and 2000 do not appear to consistent. Given average temperatures vary much more between regions of the country than across the census years used by DG, both correlations are surprisingly low. A 4-year average should pick up approximately the same regional variation as the 31-year average.

If the baseline weather variables are mismeasured, the error this creates is amplified by the way DG implement their climate change impacts. Yields y are modeled using a quadratic function of degree days d, i.e., $y = \beta_0 + \beta_1 d^2 + \beta_2 d$, so the effect of a change in degree days from an average of d_0 to $d_1 = d_0 + \Delta$ impacts yields by

$$\begin{aligned} y_1 - y_0 &= \beta_0 + \beta_1 d_1^2 + \beta_2 d_1 - \beta_0 - \beta_1 d_0^2 - \beta_2 d_0 \\ &= \beta_1 (d_0 + \Delta)^2 + \beta_2 (d_0 + \Delta) - \beta_1 d_0^2 - \beta_2 d_0 \\ &= \beta_1 (2\Delta d_0 + \Delta^2) + \beta_2 \Delta \end{aligned}$$

In their code, DG use the variable dd89_7000 to derive the change Δ but the variable dd89 to derive the average number of baseline degree days d_0. However, since those two variables don't match, it induces additional noise in the regression.

Second, we obtain slightly different damage estimates than DG. They derive the area-weighted sum of the changes in the weather variables. Summary statistics are provided using STATA's command `summ` and the mean value is then multiplied with the coefficient estimates of the corresponding weather varaiable. The log-files posted on the AER website reveal that sometimes the authors multiply the coefficients by numbers that differ from what was obtained in the `summ` command.

A2 Sensitivity Checks

We replicate Table 1 in the main paper by including state-by-year fixed effects instead of year fixed effects in Table A1. We replicate Table 4 of the main paper by using year fixed effects instead of state-by-year fixed effects in Table A4.

Table A1: Comparison of Various Data Sources in Yield Regressions Using State-By-Year Fixed Effects

	Corn			Soybeans		
	(1a)	**(1b)**	**(1c)**	**(2a)**	**(2b)**	**(2c)**
Regression diagnostics						
Variance explained by weather	7.2%	13.9%	18.5%	12.6%	20.3%	23.4%
Model comparison tests (J-test)						
DG against other weather (t-value)		10.71	15.19		11.45	14.26
Other weather against DG (t-value)		2.04	1.71		1.12	0.91
Climate change impact (Percent)						
Hadley II-IS92a scenario	0.15	-22.24	-20.55	1.06	-19.04	-14.93
(s.e.)	(1.08)	(3.31)	(3.05)	(0.95)	(2.74)	(2.67)
[s.e. Conley]	[1.36]	[7.24]	[6.01]	[0.88]	[5.22]	[4.21]
Hadley III-B2 scenario		-65.12	-78.28		-57.73	-63.51
(s.e.)		(7.23)	(6.83)		(6.05)	(5.79)
[s.e. Conley]		[16.54]	[13.93]		[12.73]	[9.85]
Observations	6862	6862	6862	5141	5141	5141
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
State-By-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table summarizes and compares alternative regression models for the relationship between crop yields and weather. Corn yields are the dependent variable in the first three columns and soybean yields are the dependent variables in the next three columns. Columns (a) replicate the results in DG using their code and data (a quadratic in degree days 8-32°C and precipitation); columns (b) are the same models as (a) estimated with our reconstructed data; columns (c) account for within-day temperature variation and extremely warm temperatures (square root of degree days above 34°C as an additional variable). The variance explained by weather is the residual variance of the full specification over the residual variance of the model excluding weather. The non-nested J-tests compare models, taking DG's model (columns a) to be best as the null hypothesis and vice versa.

Table A2: Comparison of Various Data Sources in Regressions Using Observations East of 100 Degree Meridian

	Corn				Soybeans				Profit			
	(1a)	**(1b)**	**(1c)**	**(1d)**	**(2a)**	**(2b)**	**(2c)**	**(2d)**	**(3a)**	**(3b)**	**(3c)**	**(3d)**
Regression diagnostics												
Variance explained by weather	10.6%	20.7%	25.0%	26.4%	15.1%	28.7%	35.8%	35.9%	0.1%	1.0%	1.4%	1.2%
Model comparison tests (J-test)												
DG against other weather (t-value)		15.65	18.71			16.23	24.00			7.59	8.81	
Other weather against DG (t-value)		0.93	0.91			1.48	1.32			1.23	1.20	
Climate change impact (Percent)												
Hadley II-IS92a scenario	-1.76	-14.93	-16.02	-17.14	-4.28	-16.85	-18.11	-16.39	4.93	-30.84	-38.49	-27.94
(s.e.)	(1.34)	(2.05)	(1.97)	(1.56)	(1.36)	(1.76)	(1.65)	(1.25)	(4.17)	(5.91)	(6.00)	(4.75)
[s.e. Conley]	[2.49]	[7.37]	[6.90]	[7.21]	[2.20]	[5.05]	[4.57]	[4.10]	[4.71]	[12.01]	[12.26]	[11.08]
Hadley III-B2 scenario		-55.36	-72.79	-77.72		-53.02	-75.34	-69.49		-56.15	-85.31	-64.46
(s.e.)		(4.67)	(4.92)	(4.07)		(4.01)	(3.93)	(3.11)		(10.33)	(11.21)	(9.00)
[s.e. Conley]		[17.60]	[15.82]	[16.91]		[12.50]	[10.89]	[10.16]		[21.90]	[24.31]	[21.93]
Observations	5717	5717	5717	7538	4912	4912	4912	6504	6960	6960	6960	9653
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table summarizes and compares alternative regression models for the relationship between crop yields or profit and weather. Corn yields are the dependent variable in the first three columns, soybean yields in the middle three columns and profit per acre in the last three columns. Columns (a) replicate the results in DG using their code and data (a quadratic in degree days 8-32°C and precipitation); columns (b) are the same models as (a) estimated with our reconstructed data; columns (c) account for within-day temperature variation and extremely warm temperatures (square root of degree days above 34°C as an additional variable); columns (d) include observations that are missing in DG's data. The variance explained by weather is the residual variance of the full specification over the residual variance of the model excluding weather. The non-nested J-tests compare models, taking DG's model (columns a) to be best as the null hypothesis and vice versa.

Table A3: Comparison of Various Data Sources in Regressions Using Observations East of 100 Degree Meridian

	Corn				Soybeans				Profit			
	(1a)	**(1b)**	**(1c)**	**(1d)**	**(2a)**	**(2b)**	**(2c)**	**(2d)**	**(3a)**	**(3b)**	**(3c)**	**(3d)**
Regression diagnostics												
Variance explained by weather	5.7%	13.5%	20.0%	21.4%	11.5%	18.6%	22.4%	21.4%	0.1%	0.3%	0.3%	0.3%
Model comparison tests (J-test)												
DG against other weather (t-value)		10.32	15.43			10.54	14.14			3.90	4.11	
Other weather against DG (t-value)		1.30	1.00			1.18	1.00			1.91	1.95	
Climate change impact (Percent)												
Hadley II-IS92a scenario	0.78	-20.36	-15.69	-15.07	0.66	-19.06	-14.98	-14.39	3.17	-0.03	-4.05	1.21
(s.e.)	(1.14)	(3.49)	(3.18)	(2.51)	(0.92)	(2.80)	(2.67)	(2.07)	(3.83)	(7.99)	(8.28)	(6.73)
[s.e. Conley]	[1.53]	[7.98]	[6.62]	[6.50]	[0.89]	[5.34]	[4.24]	[4.28]	[3.53]	[8.92]	[9.42]	[8.52]
Hadley III-B2 scenario		-65.57	-79.33	-78.95		-56.67	-65.38	-61.06		-7.61	-10.80	-6.86
(s.e.)		(7.96)	(7.31)	(6.00)		(6.19)	(5.82)	(4.99)		(14.64)	(14.88)	(12.47)
[s.e. Conley]		[18.75]	[15.15]	[15.52]		[13.00]	[10.10]	[10.49]		[16.41]	[17.39]	[16.46]
Observations	5717	5717	5717	7538	4912	4912	4912	6504	6960	6960	6960	9653
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-By-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table summarizes and compares alternative regression models for the relationship between crop yields or profit and weather. Corn yields are the dependent variable in the first three columns, soybean yields in the middle three columns and profit per acre in the last three columns. Columns (a) replicate the results in DG using their code and data (a quadratic in degree days 8-32°C and precipitation); columns (b) are the same models as (a) estimated with our reconstructed data; columns (c) account for within-day temperature variation and extremely warm temperatures (square root of degree days above 34°C as an additional variable); columns (d) include observations that are missing in DG's data. The variance explained by weather is the residual variance of the full specification over the residual variance of the model excluding weather. The non-nested J-tests compare models, taking DG's model (columns a) to be best as the null hypothesis and vice versa.

Table A4: Regressing Sales on Value of Production using Year Fixed Effects

	Combined (1)	**Corn (2)**	**Soybeans (3)**	**Wheat (4)**	**Cotton (5)**
Using all farms in a county					
Coefficient	0.718***	0.640***	0.702***	0.862***	0.819***
Std. Error	(0.012)	(0.012)	(0.007)	(0.009)	(0.025)
Observations	11595	10268	8241	10540	2573
Using farms with no livestock					
Coefficient	0.792***	0.693***	0.721***	0.864***	0.823***
Std. Error	(0.016)	(0.013)	(0.008)	(0.010)	(0.029)
Observations	11083	9303	7569	9490	2207

Notes: The table replicates Table 4 of the main paper except that state-by-year fixed effects are replaced with year fixed effects. We regress total sales on total value-of-production. Each coefficient and standard error is from a separate regression. All regressions include county and year fixed effects. The first column uses the combined sales of all crops where sales are reported in the Micro-files: corn, soybeans, wheat, cotton, oats, sorghum, barley, and tobacco. The next four columns report individual regressions for the four most important commodities. The first three rows aggregate all farms in a county. The second three rows only aggregate farms that do not have any livestock. Significance levels whether coefficients are different from 1 are indicated by: *** (1%); ** (5%); and * (10%).