UC Berkeley
UC Berkeley Previously Published Works

Title

The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment

Permalink

Journal

ISSN

Authors

Fisher, Anthony C
Hanemann, W Michael
Roberts, Michael J
et al.

Publication Date

DOI

The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment[†]

By Anthony C. Fisher, W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker*

Agriculture is the sector of the economy most directly linked to climate and, thus, likely to be affected by climate change. To date, however, there exists considerable disagreement about not only the magnitude of potential impacts but also the sign.

A recent paper by Olivier Deschênes and Michael Greenstone (2007b), henceforth DG, criticizes the hedonic model, a cross-sectional approach that regresses farmland values on climate variables for the United States first introduced by Mendelsohn, Nordhaus, and Shaw (1994), and proposes to use random year-to-year weather fluctuations in a panel of US agricultural profits and yields.[1] DG find *no* statistically significant relationship between US agricultural profits and weather variables in the same years. DG also find *no* statistically significant relationship between corn and soybean yields (output per acre) and weather. They argue 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 even prove beneficial. They conclude that "the preferred estimates indicate that climate change will lead to a $1.3 billion (2002 US dollars), or 4.0 percent, increase in annual agricultural sector profits. (...) The basic finding of an economically and statistically small effect is robust to a wide variety of specification checks (...). Additionally, the analysis indicates that the predicted increases in temperature and precipitation will have virtually no effect on yields among the most important crops (i.e., corn for grain and soybeans) (...)."

In this comment we revisit their paper in an attempt to reconcile their findings with others in the literature, which suggest a less optimistic outcome.[2] We present

[1] The first part of DG's paper argues that the hedonic approach does not produce robust results. We replicate the same checks using a well-specified hedonic model in the online Appendix and show them to be robust.

[2] For example, in Schlenker and Roberts (2009) we find a strong relationship between corn, soybean, and cotton yields and weather. The relationship is robust and very similar if derived from time-series variations in weather or cross-sectional variations in climate. Holding fixed the locations where crops are grown, we predict yields losses of 30–46 percent by the end of the century under the slowest warming scenario and 63–82 percent under the fastest warming scenario. These predictions also accord with our research that uses the hedonic approach, e.g., Schlenker, Hanemann, and Fisher (2006).

evidence showing that the differences stem mainly from three sources: (i) data and coding errors in DG's weather data, agricultural data, and the construction of climate-change scenarios; (ii) the particular climate change scenario which is used for impact predictions; and (iii) standard errors that are biased due to spatial correlation. Correcting DG's data and coding errors makes predictions for climate-change impacts unambiguously negative in all but one specification. The exception is a profit regression with state-by-year fixed effects where the standard errors are very large because state-by-year fixed effects absorb almost all variation in weather.

DG's measure of profits is reported sales in a given year minus reported production expenditures in that year, each variable being derived from the Census of Agriculture. This measure of profits creates a potential problem for DG's method of analysis. It does not include implicit costs like farm household labor or inventory adjustments during the year being reported. It does not control for crops produced in the reporting year but not sold until a later year, or for crops sold in the reporting year that had been harvested in a previous year and stored.[3] The problem is that storage and other inventory decisions, like the holding of animal breeding stocks, are captured by the error term and are also correlated with weather, the key explanatory variable. The induced correlation between the error term and key explanatory variables violates the identification assumption and causes the estimated effect of weather to be biased toward zero.

I. Data Irregularities

To investigate the differences between DG's findings and others in the literature 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 differences in findings.

DG have two weather variables in their dataset: 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.[4] 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 four-year and 31-year averages across counties. We reconstruct the same weather variables from raw data sources and find a correlation of 0.996. 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 of April through

---

[3] Since prices are low during years when production is high and vice versa, there is an incentive to smooth sales over time. Progressive income taxes and imperfect insurance also create an incentive to smooth sales and profits.

[4] Growing degree days integrate the product of temperature and time (measured in days) above a baseline temperature and below an upper threshold. For example, DG's baseline temperature is 8°C, so one day at a temperature of 13°C would equal five 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.
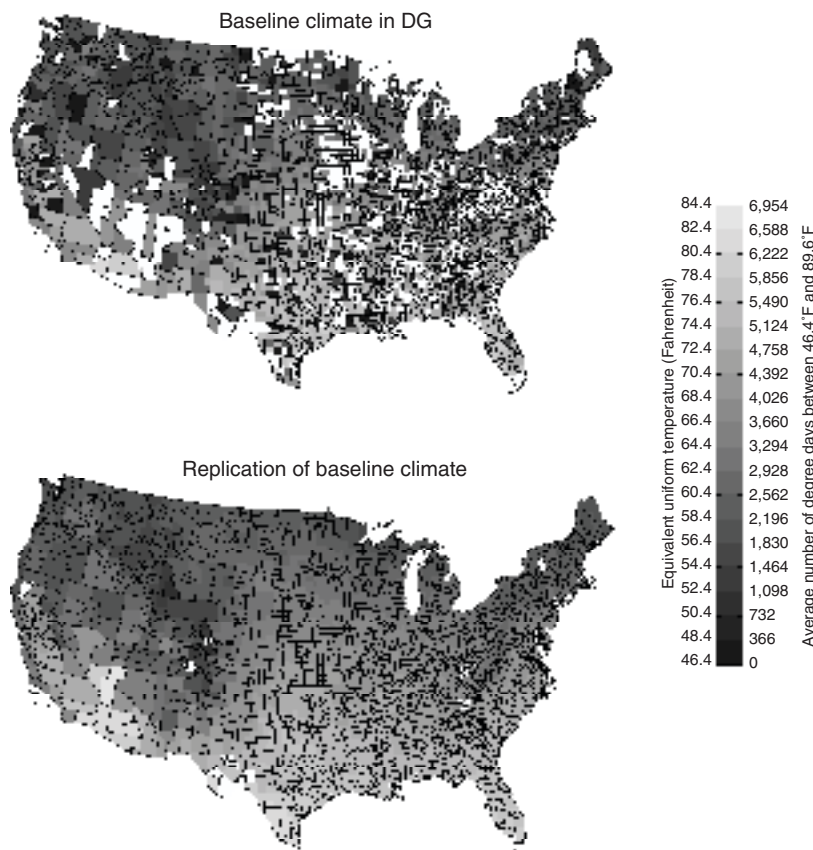
FIGURE 1. BASELINE CLIMATE IN DESCHÊNES AND GREENSTONE

Notes: Graphs display the baseline climate. The top panel is the data used in DG; the bottom panel shows 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.

September. Temperatures this low would seem implausible in any state, yet many of these counties are in warm southern states such as 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.[5] 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 and latitudes further south are warmer. Natural exceptions are mountain chains like the Rockies in the West or the Appalachians in the East, where temperatures are cooler due to gains in

---

[5] We derive degree days 8–32°C by first calculating average daily temperatures in each of the 2.5 × 2.5 mile grids of the data in Schlenker and Roberts (2009) and average over all cells with positive agricultural area.
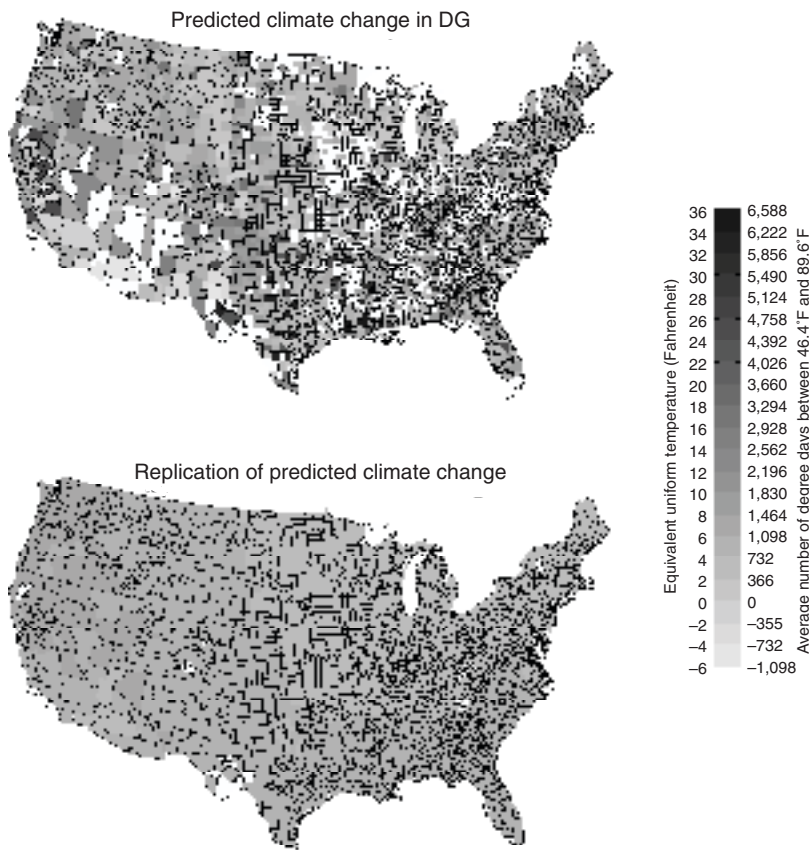
Figure 2. Climate Change Predictions in Deschênes and Greenstone

Notes: The top panel is the data used in DG; 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.

altitude. The discontinuous pattern in DG's data induces excess weather variation, which causes significant attenuation bias on their parameter estimates, especially in regression models that use state-by-year fixed effects. Within-state-year temperature deviations in our replicated dataset are approximately one-seventh the size of DG's.

Third, DG's predicted changes in climate vary widely and abruptly over the contiguous United States. They range from a decrease of 880 growing degree days (equivalent to a uniform 4.8°F decrease during the growing season) to a 6,572 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 United States 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 both inaccurate baseline values of zero for some observations and the way they combine observed weather and climate-change forecasts. 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, Los Angeles and San Francisco, the Salinas Valley and the San Joaquin Valley, Mount Whitney and Death Valley, are all assumed to have the same climate, after climate change occurs, since they are all in the same state. 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 temperature change of $-2.3°F$); Kings county has an increase of 403 degree days ($+2.2°F$) and Tulare an increase of 4,685 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 implausibly different for neighboring 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 US location. The standard approach in the climate science literature is *not* to compare GCM projections with historic climate: it is to add GCM projections of regional climate *change* between, say, 1970–1999 and 2070–2099 to the subregional baselines, thereby preserving subregional variation and avoiding regression toward the mean.

II. Replication and Comparison

While there are other differences between DG's model of yields and our own model, much of the difference in the predicted impacts of climate change stems from the data issues described above. Comparisons of the original and replicated yield and profit regressions are summarized in Table 1.

In our replications we fix the sample so they exactly match those used by DG. This excludes some agriculturally important counties that are missing in DG's data. For example, 66 of Iowa's 99 counties are missing from their dataset, yet Iowa is the largest producer of corn and soybeans, the nation's two largest crops. In an online Appendix we present results where these counties are included, but the results change little. The online Appendix also replicates the analysis for the subsample of dryland counties east of the one-hundredth meridian, the historic boundary between (primarily) irrigated and (primarily) rainfed agriculture in the United States.[6]

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 either year or state-by-year fixed effects, and

---

[6] In irrigated areas west of the one hundredth meridian, water comes mostly from aquifers or from snow or rain falling in distant locations, so local precipitation is a poor measure of water availability, and predicted climatic changes in local precipitation do not measure predicted changes in access to irrigation water.

Table 1—Comparison of Various Data Sources in Yield and Profit Regressions

	Corn		Soybeans		Profit (sales − expenditures)			
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Regression diagnostics								
Variance explained by weather	11.6%	19.6%	14.4%	30.6%	0.4%	1.5%	0.4%	0.6%
Climate change impact (percent)								
Hadley II-IS92a scenario	−0.80	−10.61	−2.73	−15.63	−6.63	−36.50	3.75	1.21
(SE)	(1.24)	(1.45)	(1.38)	(1.60)	(3.03)	(5.41)	(2.82)	(12.88)
[SE clustered by state]	[2.08]	[4.18]	[2.08]	[4.93]	[4.98]	[10.34]	[3.98]	[15.18]
Hadley III-B2 scenario		−42.01		−51.59		−55.99		−3.28
(SE)		(3.23)		(3.65)		(8.93)		(20.61)
[SE clustered by state]		[11.14]		[11.80]		[16.58]		[25.12]
Observations	6,623	6,623	5,140	5,140	9,024	9,024	9,024	9,024
Soil controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	No	No
State-by-year FE	No	No	No	No	No	No	Yes	Yes

Notes: This table summarizes and compares alternative regression models. 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. The variance explained by weather is 1 minus the ratio of the residual variance of the full specification over the residual variance of the model excluding weather. Standard errors in parentheses cluster by county following DG, while standard errors in brackets cluster by state.

regressions are area weighted.[7] To avoid confounding our comparison with changes in specification, we use the same variables, the same weighting, and the same observations. There is one exception to this rule: 241 of DG's yield observations (3.5 percent of the observations for corn and 0.02 percent for soybeans) are zero even though the Agricultural Statistical Service shows that production was positive. We drop these observations in the yield regression.[8]

Columns (a) of Table 1 replicate results in DG using their original dataset. Our replication of DG differs slightly from DG's original due to (i) a coding error in DG that we corrected (this is detailed in the Appendix) and (ii) the fact that we drop observations with zero yields. Columns (b) replicate DG's regression model using our reconstruction of the weather data using their specification.[9]

The first row of Table 1 reports the variance explained by the weather variables.[10] For all models using year fixed effects (columns 1a and 1b for corn, 2a and 2b for soybeans, and 3a and 3b for profits) our replicated weather variables explain roughly twice as much of the variance in the dependent variable. In the profit model using state-by-year fixed effects (columns 4a and 4b), our replication explains approximately 50 percent more. Recall that our replicate weather variable has a lower

[7]Yield regression use cropland-area weights, while profit regressions use farmland-area weights. Since the results are comparable whether we use year or state-by-year fixed effects in the yield regression, we report them only for year fixed effects here. The interested reader is referred to the online Appendix for results using state-by-year fixed effects.

[8]However, we obtain similar point estimates if they are included.

[9]In a sensitivity check, DG include a variable to measure the potentially harmful effect of extreme heat on profits in Table 6 of their paper. In the online Appendix we therefore also include one additional variable, the square root of degree days above 34°C in columns (c) of the regression to account for extreme temperatures. The results are comparable to columns (b) if predicted temperature increases are limited but diverge further for larger, nonmarginal changes under the Hadley III model projections.

[10]We calculate the variance explained by weather as one minus the ratio of the residual variance from the full specification over the residual variance from a model with all controls but no weather variables.

variance than DG's original data yet explains a larger share of the variance of the dependent variable. This suggests that DG's weather data had significant measurement error, which likely results in attenuation bias towards zero.

The next six rows replicate the predicted climate change impacts. The first three rows give predicted impacts under the Hadley II scenario used by DG. The predicted impacts are insignificant if we use DG's data in columns (a) but are statistically significant in columns 1b, 2b, and 3b if we use our replicated weather variables and year fixed effects. We report two sets of standard errors. The first set uses the specification of DG and clusters the error by county. This allows for heteroskedasticity and autocorrelation of counties across year but assumes that observations are identically distributed in space. In the second set of standard errors [square brackets] we therefore cluster by state after specifying the panel structure of our data, allowing for spatial correlation of counties within a state in a year.[11] This increases standard errors considerably, yet our predicted impacts are still significant at the 5 or even 1 percent level. Corn yields, soybean yields, and profits are predicted to decline by 11 percent, 16 percent, and 37 percent, respectively. While the coefficients are statistically significant, the impacts are smaller in magnitude than earlier estimates in the literature that DG use as benchmarks.

We note, however, that DG used different climate change predictions than earlier studies to which they compare their results. We therefore replicate the predicted impacts using the Hadley III model in the next three rows, as an appropriate comparison should leave the climate forecasts fixed.[12] Predicted impacts remain significant and become larger in magnitude under the Hadley III model, comparable to earlier estimates in the literature. A discussion of the validity or the accuracy of either model is beyond the scope of this article, but we stress that these differences in impacts are due not to differences in modeling assumptions but to differences in predicted climate change.

In columns 4a and 4b of Table 1 we present the same set of results as in columns 3a and 3b with one critical difference: we use state-by-year fixed effects in place of year 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 datasets 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 United States. A countrywide yearly fixed effect would not capture the fact that crops specific to California might face unique price shocks. However, any crop-specific price

---

[11] An alternative would be to directly account for spatial correlation, e.g., using Conley (1999). We note that the results are comparable to the ones obtained using errors that are clustered by state. We report the latter as they simply require one minor modification in DG's code. DG had one table that adjusted standard errors for spatial correlation of the error terms, but the code posted on the *AER* website did not specify how their standard errors were derived.

[12] 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 larger temperature increases in North America, especially in summer, and a less optimistic forecast of changes in precipitation. In other research, DG have also used the Hadley III model (Deschênes and Greenstone 2007a).

TABLE 2—TEMPERATURE VARIATION UNDER VARIOUS SETS OF FIXED EFFECTS

	Variable dd89 in DG			Replication of dd89		
	R^2	σ_e	$|e| > 1\mathbf{F}$	R^2	σ_e	$|e| > 1\mathbf{F}$
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
No fixed effects (FE)		6.85F	91.2%		6.10F	89.9%
County FE	0.845	2.70F	56.8%	0.940	1.50F	65.0%
County + year FE	0.867	2.50F	55.0%	0.979	0.88F	24.4%
County + state-by-year FE	0.879	2.39F	50.8%	0.997	0.35F	1.3%

Notes: This table summarizes regressions of degree days, a temperature measure, on various sets of fixed effects and how much of the variation they absorb. The first three columns use the variable dd89 from DG, and the last three columns use our recalculation of the same variable when data errors are corrected. Columns (a) report the R^2 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°F over the growing season.

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.[13] Accounting for region-specific price responses should therefore make predicted impacts *more negative* as it cancels out the counterbalancing price response. It is counterintuitive that predicted changes in profits are negative and significant in a regression using year fixed effects yet turn positive 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? DG provide no clear answer. One hypothesis is that there is no statistical significance because there is too little statistical power. Thus, while we fail to reject no impact, we also fail to reject large negative impacts.

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 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 dataset and our own are reported in Table 2. The table summarizes regressions of degree days against different sets of fixed effects: (i) an intercept; (ii) county fixed effects; (iii) county plus year fixed effects; and (iv) county plus state-by-year fixed effects. The table reports the R^2, the standard deviation of the residual weather variation not absorbed by the fixed effects (in Fahrenheit equivalent),[14] and the fraction of residuals with an absolute value greater than 1° Fahrenheit. While the overall standard deviations of temperatures are similar in DG's measure and our replication (6.85°F versus 6.10°F), the

---

[13] This is especially true for speciality crops like lettuce where world trade is limited in volume.

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

two differ markedly once we include fixed effects. The residual standard deviation of DG's temperature measure is 2.70°F with county fixed effects and 2.39°F with county plus state-by-year fixed effects. Our measure, on the other hand, has a residual standard deviation of 1.50°F with county fixed effects and just 0.35°F with county plus state-by-year fixed effects. These differences suggest a noise-to-signal ratio of DG's temperature measure of about 7 to 1 in their preferred fixed-effects model.[15]

III. 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 are used in a model of corn and soybean yields. The results from the profit regressions are mixed and depend on whether one uses year or state-by-year fixed effects. We find significant negative impacts if year fixed effects are used and insignificant impacts (with large standard errors) for the case of state-by-year fixed effects.

As outlined at the end of the preceding section, state-by-year fixed effects absorb almost all variation and the identification rests on very slim margins, so even small amounts of measurement error will be greatly amplified. The reason why impacts in the yield regression are insensitive to the inclusion of state-by-year fixed effects and not those in the profit regression is related to DG's profit measure, the difference between agricultural sales reported for a given year and production expense reported for that same year. While production expense is essentially the cost associated with the crops grown in that year, the sales revenue is not necessarily the revenue from crops *grown* in that year—it is revenue from crops *sold* in that year. With major field crops in the United States, such as corn, soybeans, and wheat, farmers accumulate stocks in high-yielding years; 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.[16]

The most plausible argument against these dynamic considerations is DG's use of year or state-by-year fixed effects.[17] These fixed effects account for the incentive to

---

[15] Recall that our replication of their degree days measure explains approximately twice the variance in year-to-year crop yields (Table 1). It is therefore unlikely that our replication smooths temperatures too much in space, as it is superior at explaining yield variation in space.

[16] Other factors besides storage could cause the short-run response to weather to understate the long-run response to climate. For example, 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, Murphy, and Scheinkman 1994).

[17] 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 little information about how prices will adjust (Cline 1996).

Table 3—Regressing Sales on Value of Production

	(1)	(2)	(3)	(4)
Coefficient	0.822***	0.870***	1.015	0.978
(SE)	(0.029)	(0.028)	(0.039)	(0.034)
p-val. for coeff. $= 1$	<0.0001	<0.0001	0.70	0.52
Observations	10,891	10,891	10,891	10,891
County FE	Yes	Yes	No	No
Year FE	Yes	No	Yes	No
State-by-year FE	No	Yes	No	Yes

Notes: This table summarizes regressions of county-level sales on value-of-production for corn, soybeans, wheat, cotton, oats, sorghum, and barley. Sales and production data for all farms that do *not* have any livestock are added for each county. Following DG, both the dependent and independent variables are on a per-acre basis, and the regression is weighted by acres.

 ***Significant at the 1 percent level.
  **Significant at the 5 percent level.
   *Significant at the 10 percent level.

accumulate or deplete inventories, which are connected to prices (Williams and Wright 1991; Deaton and Laroque 1992). However, this would work only if prices of all commodities within a county move together and there is no substate price variation.[18]

To examine this issue empirically, we conducted the following exercise, the results of which are reported in Table 3. We regress *sales* against the *value of production*.[19] Columns 1 and 3 include year fixed effects, while columns 2 and 4 include state-by-year fixed effects. The first two columns also include county fixed effects, while the last two do not. Hence, the first two columns use deviations from county means for identification: the regressions capture how much sales differ from average in relation to production value relative to its average. If storage variations are fully accounted for in the model, the coefficient should be one. That is, sales should increase one-for-one by the value of each extra unit that is produced. Columns 3 and 4, on the other hand, drop county fixed effects, and the identification relies on the cross-section: these regressions show whether, on average, counties sell as much as they produce.

The first two columns reveal that storage is an important factor in sales. The coefficients are significantly different from one. State-by-year fixed effects do account for some of the tendency to store yield shocks, as the coefficient in column 2 is closer to one than the coefficient in column 1. One alternative explanation for why the coefficients could be less than one is measurement error and attenuation bias.

---

[18] Brennan, Williams, and Wright (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 local price variations for which futures markets do not exist. While such local price variation may be small relative to overall price fluctuations, true within-state weather variations in a given year are surely small too.

[19] To conduct this exercise we obtained access to individual farm-level data from the Agricultural Census Microfiles. These are the raw data used to construct the aggregate sales data used in DG. The survey asks farmers to report their sales of seven storable crops: corn, soybeans, wheat, cotton, oats, sorghum, and barley. It also asks for the yield of these 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 drop all farms that have livestock as sales could be less than the value of production when farmers feed these crops to animals. We then aggregate all nonlivestock 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 do DG, and cluster by state to get more conservative standard errors.

However, note that if we drop county fixed effects in columns 3 and 4 and, hence, no longer rely on year-to-year deviations that give incentives for storage, the coefficient is no longer different from one. If measurement error was a pervasive problem, these coefficients should also be biased toward zero.[20]

Finally, one might wonder whether a coefficient of 0.8 is different from 1 in an economically meaningful way. Recall that we are looking at sales in Table 3, and profits are the difference between sales and expenses. For comparison, consider that production expenses in the 2002 census equaled 86 percent of agricultural sales, which is about the coefficient in the first two columns of Table 3.

IV. 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, given 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.

A recent study by Deschênes and Greenstone (2007b) argues that predicted impacts are not economically significant: estimated relationships between weather variables, yields, and profits are taken to imply that the impact of long-run climate change on US agriculture will be either insignificant or modestly beneficial. On the other hand, earlier research by us and others has found large potential impacts, estimated from both cross-sectional and time-series variations in weather and climate.

Likely explanations for the divergence in findings include: (i) 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; (ii) DG's unusual and, in our judgment, incorrect treatment of climate-change predictions that assume a uniform future climate within each state; (iii) DG's implicit assumption that errors are not spatially correlated; and (iv) conceptual difficulties in their profit-based approach due to the confounding effects of storage and possibly also capital and inventory adjustments or 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 US 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 evidence from recent experiments that more realistically simulate fertilization suggests that their impact on crop yields will be much more 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, Wang, and Curtis 2002).

---

[20] DG argue in their reply that including lagged weather variables will solve the storage problem. We believe this is incorrect for two reasons. First, it does not break the contemporaneous correlation between weather and the error term. Second, any model with temporal fixed effects captures at least part of the storage decision in the fixed effect instead of the variable of interest: historic weather.

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 (or data errors) 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.** 1997. "Convenience Yield without the Convenience: A Spatial-Temporal Interpretation of Storage under Backwardation." *Economic Journal* 107 (443): 1009–22.

Cline, **William R.** 1996. "The Impact of Global Warming on Agriculture: Comment." *American Economic Review* 86 (5): 1309–11.

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

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

Deschênes, **Olivier**, **and Michael Greenstone.** 2007a. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." National Bureau of Economic Research Working Paper 13178.

Deschênes, **Olivier**, **and Michael Greenstone.** 2007b. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review* 97 (1): 354–85.

Deschênes, **Olivier**, **and Michael Greenstone.** 2007c. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Dataset." *American Economic Review.* http://www.aeaweb.org/articles.php?doi=10.1257/aer.97.1.354.

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. 2012. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment: Dataset." *American Economic Review.* http://dx.doi.org/10.1257/aer.102.7.3749.

Jablonski, **Leanne M.**, **Xianzhong Wang**, **and Peter S. Curtis.** 2002. "Plant Reproduction under Elevated CO_2 Conditions: A Meta-Analysis of Reports on 79 Crop and Wild Species." *New Phytologist* 156 (1): 9–26.

Long, Stephen P., Elizabeth A. Ainsworth, Andrew D. B. Leakey, Josef Nosberger, and Donald R. Ort. 2006. "Food for Thought: Lower-Than-Expected Crop Yield Stimulation with Rising CO_2 Concentrations." *Science* 312 (5782): 1918–21.

Mendelsohn, **Robert**, **William D. Nordhaus**, **and Daigee Shaw.** 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *American Economic Review*, 84 (4): 753–71.

Rosen, **Sherwin**, **Kevin M. Murphy**, **and Jose A. Scheinkman.** 1994. "Cattle Cycles." *Journal of Political Economy* 102 (3): 468–92.

Schlenker, **Wolfram**, **W. Michael Hanemann**, **and Anthony C. Fisher.** 2006. "The Impact of Global Warming on US Agriculture: An Econometric Analysis of Optimal Growing Conditions." *Review of Economics and Statistics* 88 (1): 113–25.

Schlenker, **Wolfram**, **and Michael J. Roberts.** 2009. "Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change." *Proceedings of the National Academy of Sciences* 106 (37): 15594–8.

Williams, **Jeffrey C.**, **and Brian D. Wright.** 1991. *Storage and Commodity Markets*. New York: Cambridge University Press.

This article has been cited by:

1. Melissa Dell, Benjamin F. Jones, Benjamin A. Olken. 2014. What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature* **52**:3, 740-798. [Abstract] [View PDF article] [PDF with links]