EJW
ECON JOURNAL WATCH
Scholarly Comments on
Academic Economics

ECON JOURNAL WATCH 20(2)
September 2023: 234–253

Temperature Shocks and Economic Growth: Comment on Dell, Jones, and Olken

David Barker[1]

LINK TO ABSTRACT

In 2012, Melissa Dell, Benjamin Jones, and Benjamin Olken (DJO) published an article in *American Economic Journal: Macroeconomics* arguing that higher temperatures reduce the rate of economic growth. According to the Web of Science, the paper is in the top one percent for citations in academic economics and business publications. Google Scholar lists 1,895 citations. The same team of authors had published a briefer and related article in the *American Economic Review* (DJO 2009), which has 605 Google Scholar citations. The present commentary focuses on DJO (2012).

DJO (2012) was widely covered in the popular press and the authors are extremely well credentialed. Their result is important because the rate of economic growth is the most important determinant of future world wealth. The present paper does not take up the issue of whether temperature increases are coming. It only addresses the question of whether temperature increases reduce the rate of economic growth.

In a previous issue of this journal (Barker 2023) I commented on a Federal Reserve publication by Michael Kiley (2021), which is based on the methods and data of DJO, although it used more sophisticated econometrics. I found that Kiley's methods produced similar results from simulated data where no effect of temperature fluctuations on growth was present, and I showed that dropping countries with unusual events such as large oil discoveries and genocides reversed

1. Views expressed in this comment are my own and not those of the Iowa Board of Regents or any other organization I am affiliated with.

his results. In an earlier article I had debunked another study first published by the Federal Reserve that claimed to find effects of warm temperatures on economic growth using U.S. state-level data (Barker 2022; Colacito et al. 2019; 2018). The authors of the two criticized studies have been invited to reply but have not done so thus far; the invitation remains open.

The statistical methods of DJO are less complicated than in Kiley (2021) and their errors are not as subtle. Each of the following five statements is true as a single, standalone criticism or robustness check: (1) They use an untenable method of classifying countries by income; using more reasonable methods I find that their results disappear. (2) Their results are influenced by arbitrary methodological choices. (3) Their results are influenced by a small number of observations with unusual characteristics. (4) The inclusion of additional countries and more recent data weakens their results. (5) Alternative data does not support their hypothesis that high temperatures reduce economic growth.

Description of DJO (2012)

DJO (2012) regress growth on temperatures and claim to find a causal relationship between them in which higher temperatures lower economic growth. In their abstract DJO say that "higher temperatures substantially reduce economic growth." DJO uses annual data on average temperatures and rates of economic growth by country. The authors regress annual growth on annual average temperatures, with a sample of observations on 127 countries in years from 1961 to 2003 for a total of 4,924 observations. The panel is unbalanced, with some countries missing data for some years. There is no weighting of observations, so China, with a population of 1.2 billion, has the same influence on the results as St. Vincent and the Grenadines, with a population close to 100,000 and an area one-eighth the size of Rhode Island. Every country is assigned a single average temperature for each year.

DJO include a number of fixed-effect control variables in their model as independent variables. There are dummy variables for each country, and dummy variables for each year/region combination. For example, there is a dummy variable equal to 1 for all observations for which both of the following are true: the year is 1961 and the region is Middle East and North Africa; if one or both of those statements is not true, the dummy takes the value 0. There are also dummy variables for poor countries in each year; those dummy variables allow poor countries to have their own trend of per capita GDP growth that affects growth independent of temperature.

Average annual temperature is another independent variable. The main variable of interest, however, is average temperature multiplied by a dummy variable equal to 1 for countries with per capita GDP below the sample median and 0 otherwise. In this way DJO picks up what they claim is an overall effect of temperature on economic growth, and a separate effect for poor countries. The entire claimed effect of temperature on growth for poor countries is the sum of the coefficients on average temperature and average temperature interacted with the income dummy variable.

The independent variables described above comprise all of the variables included in the regression. Annual per capita GDP growth is the dependent variable. The estimation method is ordinary least squares, with standard errors adjusted for clustering by both country and year.[2]

The primary results of DJO are reported in their second table, titled "Main Panel Results" (DJO 2012, 75). Table 1 shows the results from the second column of their table. Other columns in their table show the results including precipitation and an agricultural country dummy variable, but the results for temperature are very similar in each column. DJO (ibid., 67–68) state that "Changes in precipitation have relatively mild effects on national growth," so I focus on the results in column (2) of their Table 2. The t-statistics and p-values are not reported in DJO's Table 2. Those shown in Table 1 below are from my replication of DJO. DJO report only whether a result is significant at the 1-, 5-, or 10-percent level.

TABLE 1. Results from 2 column (2) in DJO (2012, 75), with replicated t-statistics and p-values

Dependent variable: annual growth rate	
Temperature	0.261
Standard error	0.312
t-statistic	0.836
p-value	0.403
Temperature × poor	−1.655
Standard error	0.485
t-statistic	−3.410
p-value	0.001
Total effect on poor	−1.394
Standard error	0.408
t-statistic	−3.418
p-value	0.001
R^2	0.223
Observations	4924

2. The text in Table 2 of DJO (2012, 75) says only that standard errors are clustered by country, but they are actually clustered by both country and year.

For countries that are above median GDP per capita, higher temperatures are associated with higher economic growth, but the effect is statistically insignificant. For poor countries, the total relationship is the sum of the two coefficients "Temperature" and "Temperature × poor" in the table above; the sum is −1.394 and is statistically significant. For poor countries, an increase in temperature of one degree Celsius is associated with economic growth that is 1.394 percentage points lower. The R^2 for the regression is 0.22.

If we take DJO's claim of a causal relationship between temperature and growth seriously, it is interesting to note that the coefficient estimates of the model, although not statistically significant, imply that higher temperatures will increase worldwide wealth. For countries in the upper half of per capita GDP, each degree of temperature increase would raise annual growth by 0.261 percentage points. For countries in the lower half of per capita GDP, a degree of warming would lower growth by 1.394 percentage points. I obtained the predicted temperature increase to the year 2100 for all countries assuming no change in CO2 emissions from the replication files of Marshall Burke et al. (2015). I multiplied this temperature change by 0.00261 for wealthy countries and −0.01394 for poor countries and added the resulting growth adjustment to an assumed baseline growth rate of 0.02. I compounded the resulting growth rate over eighty years to obtain a growth factor up to the year 2100 for each country and multiplied that result by the country's population. The sum is more than $500 trillion higher with warming than without. This result is not noted or discussed in DJO.

Again taking seriously DJO's claim of causality, it is also interesting to note that an implication of DJO's result is that cooling temperatures would have a positive effect on economic growth in poor countries. If the positive effect of temperature on growth in non-poor countries is ignored because it is statistically insignificant, then spraying enough sulphate particulates into the air to lower temperatures by one degree Fahrenheit would increase annual growth in poor countries by 0.77 percentage points, increasing their annual GDP per capita by nearly 80 percent in the year 2100.[3]

Replication

Data and replication code are available from Melissa Dell's website (Dell 2013). I was able to exactly replicate the results in Table 2 of DJO. The Stata program the authors used named "cgmreg" is no longer easily available, but I was able to find and run an archived version. The program named "clus_nway"

3. The cost of such a program has been estimated to be only $2.5 billion per year (Smith and Wagner 2018).

has replaced cgmreg. It produces identical coefficient estimates, but the clustered standard errors are slightly different.[4]

The original data file produced by DJO named "climate_panel.dta" contains missing values for a variable containing the initial GDP per capita for Myanmar, and so, even though all data are available to include the country in the sample, it is missing from the analysis. Adding Myanmar slightly strengthens DJO's results.

In the original dataset used in DJO's Stata program, called "climate_panel," the countries Pakistan and Bangladesh are separate observations, but both are coded as "PBD" for the variable called "parent." As a result, standard errors are clustered by all countries except that these two are clustered together. In addition, the fixed-effect variables for these two countries are combined. Allowing these two countries to be separate clusters and to have different fixed effects slightly weakens DJO's results.

DJO use different regional classifications of countries than are used by the World Bank. They create a single region combining North American countries, Australia, New Zealand, and European countries. In the World Bank database, North America is a separate region, Australia and New Zealand are part of Southeast Asia, and Europe is combined with parts of Central Asia. DJO put Turkey into Europe, while the World Bank puts Turkey into the Middle East and North Africa region. Using World Bank regions slightly weakens DJO's results.

Classification of countries by income

A simple but crucial step in DJO's analysis is the creation of a dummy variable indicating whether a country has per capita GDP above or below the world median. In the footnote to Table 2, they state "Poor is defined as a dummy for a country having below median PPP GDP per capita in its first year in the data." In their sample, data for South Korea begins in 1961, when per capita GDP was $1,168 in 2022 dollars. At that time, South Korea was a poor country, recovering from the Korean War. By 1977, however, South Korean per capita GDP was above the median of countries in DJO's sample. South Korean per capita GDP was below the median for only 16 out of 43 years in the sample, yet it is classified as below the median for all 43 years. Jamaica, which was above the median in 1967, when data are first available, was below the median by the end of the sample period, yet it is classified as a non-poor country for all years.

I reclassified South Korea as poor for the years 1961–1976 and rich for the

4. Lines 37 and 38 of the file "maketable2.do" run without error when executed as a .do file, but need to be modified to run in the command line of Stata. The modification is simply to remove the comments in these two lines.

years 1977–2003. This single change nearly eliminates DJO's result, as shown in the second column of results in Table 2. The t-statistic on the temperature-income interaction variable changes from −3.4 to −1.3, and the p-value from 0.001 to 0.183. The p-value on the total association for poor countries is 0.086.

Why does South Korea have such a strong influence on DJO's results? As South Korea emerged from poverty in the 1960s, its rate of growth of per capita income was high. As it became wealthy, its rate of growth declined. This pattern is consistent with convergence, or catch-up growth, predicted by commonly used economic models (Solow 1956). The average temperature in South Korea was higher in the later years of the sample. By classifying South Korea as poor for the years 1977–2003, DJO cause the regression to interpret the simultaneous higher temperatures and lower growth as evidence of high temperatures damaging growth in poor countries. Interestingly, the warmest months in South Korea, May–October, were cooler by 1.8 degrees from 1977–2003 than from 1961–1976. It was the cold months, November–April, that were 4.4 degrees warmer.[5] If high temperatures hindered economic growth, it seems likely that it would be a result of summer temperatures. In other words, classifying South Korea as rich from 1977 to 2003 eliminates DJO's primary result, there is no justification for classifying it as poor during those years, and temperature likely had nothing to do with lower growth during those years. When all countries are allowed to have different income classifications in different years, DJO's results disappear. This result is shown in the third column of Table 2 here. Neither the separate association of temperature and growth for poor countries nor the overall association for poor countries is statistically significant.

Allowing country classification to change requires GDP data in a common currency so that GDP per capita can be compared between countries. These data are missing in the World Bank database used by DJO for some years in some countries. DJO calculates GDP growth using inflation-adjusted local currency, but per capita income in dollar terms is not always available for the same years. This is why the number of observations is smaller for the third column of results in Table 2. DJO's results hold up rerunning their regression using the same smaller sample, so the missing observations are not causing the change in the results. In most cases it is obvious whether a country is above or below the median when the data are missing. When I classified countries in these years in this way I was able to use the entire sample of 4924 observations and the results were similar. There are cases where a country changes from above or below the median, resulting in short periods of classification as poor or rich. In such periods of four years or fewer in a classification, I reclassified these years according to the classification of the

5. DJO do not use or discuss monthly temperature data. I explain in a later section how I obtained them.

surrounding years, and the results were also similar.

I also classified countries according to whether they are above or below the median for most years. The results are in the fourth column of numbers in Table 2 labeled “All mode.” I ignored missing observations to make this calculation, so that the full sample could be used. DJO’s results survive with slightly reduced statistical significance. This significance is eliminated, however, with a change to the set of fixed-effect variables used that is discussed in the next section of this paper.

TABLE 2. Results allowing income classifications to change over time

	Dependent variable: annual growth rate			
	DJO	Korea mixed	All mixed	All mode
Temperature	0.261	−0.232	−0.284	0.128
Standard error	0.312	0.293	0.284	0.317
t-statistic	0.836	−0.792	−1.000	0.405
p-value	0.403	0.428	0.317	0.685
Temperature × poor	−1.655	−0.290	−0.005	−1.327
Standard error	0.485	0.218	0.035	0.519
t-statistic	−3.410	−1.330	−0.142	−2.556
p-value	0.001	0.183	0.887	0.011
Total effect on poor	−1.394	−0.522	−0.289	−1.198
Standard error	0.408	0.304	0.282	0.436
t-statistic	−3.418	−1.718	−1.024	−2.749
p-value	0.001	0.086	0.306	0.006
R^2	0.223	0.220	0.227	0.221
Observations	4924	4924	4924	4924

Calculating median income using different starting points for different countries is clearly incorrect, but this error does not have a material effect on the results, as is shown in the fourth column of results in Table 2. Classifying countries as rich or poor for the entire 43-year period when many countries changed dramatically relative to other countries is also incorrect, and this error does have a material effect on DJO’s results. As was mentioned earlier, South Korea was poor in the early years of DJO’s sample, and rich in later years. Forcing South Korea to be one or the other for the entire sample is not reasonable. Even if income classifications of countries are constrained to be constant using a better method than DJO’s, their results depend on the inclusion of fixed effects with questionable justification, which will be discussed in a later section.

Instead of constraining the effect of temperature on growth to be discretely different between countries above and below median per capita GDP, it seems reasonable to check whether such an effect might vary continuously as income changes. In section A22 of DJO's appendix, they show results from using quintile dummy variables instead of the binary dummy variable used in their main specification, and they show results from interacting temperature with initial GDP from their sample. They apparently do not try the obvious specification of directly interacting temperature with current per capita GDP. In this specification, the effect of temperature on growth would be greatest for poor countries, and less for rich countries, with the effect declining in a linear fashion as per capita GDP is higher. The results from this interaction are shown in Table 3. There is no statistically significant effect of temperature on growth.

TABLE 3. Interaction of per capita GDP with temperature

Dependent variable: annual growth rate	
Temperature	−0.144
Standard error	0.276
t-statistic	−0.521
p-value	0.602
Temperature × per capita GDP	0.018
Standard error	0.029
t-statistic	0.615
p-value	0.539
R^2	0.222
Observations	4654

Another possibility is to allow the effect to vary continuously, but in a manner more like the discrete jump between per capita GDP classes in DJO's main specification using a logit function. Using a nonlinear estimation technique, I estimated the parameters a and k shown in equation 1. The variable g represents growth in per capita GDP, T represents temperature, P represents the level of per capita GDP, and P_0 represents median GDP per capita. If k is equal to zero, then there is no effect from the level of income, and the growth rate varies linearly with temperature. As k increases, the effect of T on g varies with P, at first linearly, and then with a small effect for countries below median GDP per capita, then a rapid jump in the effect, and a larger effect for countries above median GDP per capita.

$$g = \frac{\alpha T}{1 + e^{k\left(P - P_0\right)}} \qquad (1)$$

I used OLS to calculate residuals from regressing per capita GDP growth on all of the fixed-effect variables, then I used the nonlinear estimation technique to regress the residuals on temperature and the level of per capita GDP. I also checked to see if estimating DJO's equation on these residuals produced the same results, and doing so did not change their original results. The results of nonlinear estimation, shown in Table 4, show that neither parameter is statistically significant.

TABLE 4. Nonlinear estimate of equation 1

Parameter	Estimate
α estimate	−0.002
α standard error	0.007
α z statistic	−0.268
α p value	0.788
k estimate	0.000
k standard error	0.002
k z statistic	0.054
k p value	0.957
Observations	4654

Fixed effects

One set of fixed-effect variables included in DJO's regressions is dummy variables indicating year multiplied by dummy variables indicating whether a country is above or below median income. Including these dummy variables in the analysis holds constant the pattern of average growth of poor countries over time. DJO (2012, 79 n.21) discuss the removal of these variables, saying that "dropping the poor × year fixed effects produces similar estimated temperature effects in poor countries." These results are not reported in the paper, but in the online appendix Table A20, results are reported using urban-only data with and without the poor × year fixed effects. The t-statistic on the effect of temperature on growth in poor countries drops from 2.41 to 2.01 when the poor × year fixed effect is left out. This is not obvious in the table, since both results are only labeled as significant at the 5-percent level. In Table A30 the same specifications are reported for satellite temperature data, and the t-statistic without poor × year is only 1.03.

Why do the poor × year fixed-effect variables reduce the statistical significance of the result? I simulated a case in which high temperatures increase growth (as in DJO for rich countries), a few outlier poor countries offset this effect, and there is a warming trend in poor countries. Without the poor × year fixed-effect variables, the regression shows no statistically significant relationship, but when they are included, the association is negative and statistically significant. In DJO's data, unweighted average temperatures in poor countries tend to increase from 1976 to 2003, and at the same time unweighted average growth rates in poor countries also increase. This correlation pushes against DJO's hypothesis. By including poor × year fixed effects, DJO offset this effect, adding to the measured statistical significance. DJO discuss taking these variables out of the regression, and claim that it has no significant effect, but they never provide a justification for including them. Including these particular fixed-effect variables appears to be an arbitrary decision that happens to improve the statistical significance of their result.

TABLE 5. Results without poor × year fixed effects

	Dependent variable: annual growth rate			
	DJO	No poor × year FE	No poor × year FE, mixed	No poor × year FE, mode
Temperature	0.261	0.010	−0.280	−0.144
Standard error	0.312	0.329	0.287	0.326
t-statistic	0.836	0.031	−0.976	−0.441
p-value	0.403	0.976	0.329	0.660
Temperature × poor	−1.655	−0.884	−0.035	−0.465
Standard error	0.485	0.491	0.020	0.501
t-statistic	−3.410	−1.800	−1.705	−0.928
p-value	0.001	0.072	0.088	0.353
Total effect on poor	−1.394	−0.874	−0.315	−0.609
Standard error	0.408	0.398	0.285	0.421
t-statistic	−3.418	−2.199	−1.105	−1.445
p-value	0.001	0.028	0.269	0.148
R^2	0.223	0.209	0.209	0.209
Observations	4924	4924	4924	4924

Influential observations

In order to further assess the robustness of DJO's results I checked whether particular observations had unusual influence on the estimated coefficients using their income classification method and the same fixed effects variables that they use. The regression diagnostic statistic DFBETA, which stands for difference in beta values, can be used for this purpose. It gives the difference, as a fraction of a standard deviation of a regression coefficient, in the estimated coefficient if an observation is removed. David Belsey et al. (2004) suggest a cutoff of $\pm 2/\sqrt{n}$ for this statistic, where n is number of observations. Observations above this cutoff should be examined to see if they may be affected by factors outside of the model being tested. DJO use 4,924 observations, so the cutoff value is −0.0285. I calculated DFBETA values for each observation using the specification in Table 1.[6]

6. I used simple OLS regression not adjusted for clustered standard errors to compute DFBETA values.

The most influential observation was for Rwanda in 1994 with a DFBETA value of −0.462, sixteen times the suggested cutoff value. Out of a population of 6.7 million, approximately 500,000 Rwandans were killed over a 15-week period in 1994 (McDoom 2020). GDP dropped by 63 percent. The average temperature in the country was 0.52 degrees Fahrenheit warmer than the previous year, 0.13 degrees cooler than the following year, and 1.66 degrees warmer than the average, and the second highest for years in DJO's sample. This combination of an extreme economic downturn and higher than normal temperatures substantially influenced the model to conclude that temperature affects economic growth. In Rwanda in 1994 the warmest month was September, but the political events which caused the drop in GDP took place in April.

Other influential observations include Burundi in 1995, which was affected by events in Rwanda the previous year, and China in 1961, a year of famine caused primarily by agricultural mismanagement, and Kuwait in 1980, the first year that oil prices had ever fallen. That same year Kuwait reduced oil production by 25 percent[7] and GDP fell by 29 percent. In 1980 Kuwait had its fifth coolest year in DJO's sample of 35 years. Since Kuwait is classified as a rich country, low growth in a cool year there leads the model to increase the estimate of the negative association of temperature and growth in poor countries relative to rich countries, although it has little effect on the estimate of the total association between temperature and growth in poor countries.[8]

Dropping 13 influential observations out of 4,924 eliminates the statistical significance of the total association of temperature on poor countries at the 5-percent level, and dropping 16 observations eliminates it at the 10-percent level. That is, statistical significance disappears when we use that 0.997 portion of the set of observations. Dropping 23 observations eliminates both the total association of temperature and growth and the specific association for poor countries. These results are shown in Table 6. Sensitivity of DJO's results to the effects of a small number of observations does not necessarily invalidate their results. There are observations in the dataset that are influential in both directions, but the most influential observations that enhance DJO's results are more influential than those that diminish the results. The fact that there are such highly influential observations that have much more complicated stories than high temperature reducing growth adds uncertainty to DJO's results that are not captured in the calculated standard

---

7. *New York Times*, March 31, 1980, p. 63.

8. The effect of an observation on the total effect of temperature on poor countries is the sum of the coefficients on temperature and temperature multiplied by the dummy variable representing poor countries. To find the effect of each observation on the total effect of temperature it was necessary to recalculate the regression for each observation, since DFBETA is calculated for a single coefficient, not a linear combination of coefficients.

deviations of the coefficients that they estimate. At the very least, the existence of these highly influential observations indicates that checking DJO's results with additional data and alternative specifications is warranted.

TABLE 6. DJO results and without influential observations

	Dependent variable: annual growth rate			
	DJO	Without 13 influential	Without 16 influential	Without 23 influential
Temperature	0.261	0.282	0.286	−0.021
Standard error	0.312	0.310	0.310	0.284
t-statistic	0.836	0.910	0.923	−0.075
p-value	0.403	0.363	0.356	0.940
Temperature × poor	−1.655	−0.969	−0.881	−0.593
Standard error	0.485	0.452	0.450	0.437
t-statistic	−3.410	−2.146	−1.958	−1.355
p-value	0.001	0.032	0.050	0.175
Total effect on poor	−1.394	−0.687	−0.595	−0.614
Standard error	0.408	0.369	0.367	0.373
t-statistic	−3.418	−1.864	−1.622	−1.644
p-value	0.001	0.062	0.105	0.100
R^2	0.223	0.237	0.237	0.235
Observations	4924	4911	4908	4901

Extended data

The source that DJO used for average annual temperatures for countries has, since the publication of DJO in 2012, released data through the year 2017. DJO only used data through 2003. Their source, *Terrestrial Air Temperature and Precipitation: 1900–2006 Gridded Monthly Time Series, Version 1.01*, provided monthly temperature data, interpolated from weather stations, for 85,794 sections of the globe, each one measuring 0.5 × 0.5 degrees. Another data source, the Global Rural-Urban Mapping Project, provides population data for 741,312,000 sections of the world, each measuring 30 arc-seconds by 30 arc-seconds. DJO used GIS software to match these datasets and calculate population weighted temperatures for each country in each year. I used the raw data and matched it using my own Stata code. I also wrote a Java program to call a function called coordinate_to_country using the node.js environment. This allowed me to match each segment to a country. The temperature data have been updated since the publication of DJO, but the correlation coefficient of my calculated temperatures and those of DJO is 0.9978.

By calculating temperatures in this way, I discovered that DJO missed several small countries that do not incorporate the centroid of any 0.5-by-0.5 degree segment. Other small countries were included by DJO. For example, St. Vincent and the Grenadines is included in DJO's data, but Barbados is not, because the area of Barbados just misses the centroid of the surrounding global segments. I used the segment with a centroid nearest the capital city of these countries to provide temperature data.

I also used updated data for per capita economic growth from the World Bank's *World Development Indicators*. Data in this database is regularly changed by countries providing data, sometimes including the deletions of years of observations, sometimes for political purposes (Kadri 2016, 160). Data are also missing for other reasons. For example, GDP data are missing for Canada for the years 1970–1996 in the current version of the World Bank database, but they were available in the 2003 version of the database.

I was also able to expand the dataset by using dollar-denominated GDP from the World Bank database. DJO used the local currency unit-denominated GDP data. Data from some years in some countries are missing using the local currency unit data but are available using the dollar-denominated data. The growth rates are the same between the two series, because the IMF converts GDP to dollars in a base year and then applies local currency growth rates backwards to construct the time series of GDP for each country (**link**). For the purposes of categorizing countries by income, DJO used GDP data from the Penn World Table. These data are constructed to purchasing power parity instead of exchange rates.

Using the updated, recalculated data that I compiled, matching observations to those of DJO so that only years and countries used by DJO are included, I find results that are similar to DJO's results, indicating that the data adjustments I made are not driving the differences in results using extended data or alternative specifications.

All but the last column in the analysis reported in Table 7 use DJO's flawed method of using the same classification by income across years for countries. In the second column of results, using the extended data, the statistical significance of the total association of temperature and growth for poor countries is barely significant at the 5-percent level. Using World Bank regions and dropping the poor × year fixed effects, the association is not significant at the 10-percent level. Using changing income classifications, the association disappears.

Going back to the original data source also allowed me to use monthly temperature data. DJO use only annual average temperature. If high temperatures cause a reduction in per capita GDP, it seems likely that the warmest temperatures of the year would have the greatest impact. Because some different countries have the highest temperatures in different months, particularly those in different hemi-

TABLE 7. Comparing results from 2 in DJO (2012) with results using additional data

	Dependent variable: annual growth rate					
	DJO	Extended data	Without poor × year	WB regions	WB regions w/o poor × year	Changing income classes
Temperature	0.261	0.410	0.372	0.757	0.568	−0.215
Standard error	0.312	0.325	0.319	0.431	0.403	0.281
t-statistic	0.843	1.261	1.167	1.757	1.407	−0.766
p-value	0.403	0.207	0.243	0.079	0.159	0.444
Temperature × poor	−1.655	−1.199	−1.016	−1.571	−1.139	0.029
Standard error	0.485	0.497	0.419	0.579	0.493	0.052
t-statistic	−3.410	−2.410	−2.422	−2.714	−2.312	0.560
p-value	0.001	0.016	0.015	0.007	0.021	0.575
Total effect on poor	−1.394	−0.789	−0.644	−0.815	−0.571	−0.186
Standard error	0.408	0.401	0.334	0.410	0.374	0.276
t-statistic	−3.418	−1.967	−1.929	−1.989	−1.528	−0.674
p-value	0.001	0.049	0.054	0.047	0.126	0.500
R^2	0.223	0.201	0.195	0.194	0.187	0.206
Observations	4924	9033	9033	9033	9033	8934

TABLE 8. Monthly data

	Dependent variable: annual growth rate					
	Warmest	Coolest	Third-warmest	Third-coolest	Sixth-warmest	Sixth-coolest
Temperature	0.322	−0.127	0.412	−0.213	0.432	−0.195
Standard error	0.254	0.100	0.381	0.152	0.508	0.224
t-statistic	1.266	−1.271	1.083	−1.406	0.850	−0.871
p-value	0.205	0.204	0.279	0.160	0.396	0.384
Temperature × poor	−0.733	−0.426	−1.124	−0.697	−1.661	−1.170
Standard error	0.344	0.243	0.489	0.324	0.640	0.405
t-statistic	−2.131	−1.752	−2.300	−2.150	−2.595	−2.886
p-value	0.033	0.080	0.021	0.032	0.009	0.004
Total effect on poor	−0.411	−0.553	−0.712	−0.911	−1.229	−1.365
Standard error	0.245	0.264	0.327	0.351	0.423	0.464
t-statistic	−1.677	−2.092	−2.174	−2.594	−2.909	−2.942
p-value	0.094	0.036	0.030	0.009	0.004	0.003
R^2	0.221	0.221	0.222	0.221	0.223	0.222
Observations	4924	4924	4924	4924	4924	4924

spheres, I ranked the months for each country and put the highest monthly temperatures in one variable, the second highest in another variable, etc. Using a variety of specifications, I was unable to find evidence that warmer months had any more association with economic growth than cooler months. Table 8 shows the results for the warmest month, the coolest month, the warmest three months, the coolest three months, the warmest six months and the coolest three months.

Alternative data

To see whether DJO's results can be replicated using different data, I used Robert Tamura's dataset containing estimates of output per worker by country going back to the 19th century (Tamura et al. 2019). Tamura's data is not annual; it is mostly decennial. I calculated the annualized average growth rate between observations and matched these growth rates with the average temperature over the same period. I did two different analyses, one using all of Tamura's data, and another that was restricted to 20th- and 21st-century data. For the 20th- and 21st-century analysis, I used the same temperature data as in the previous section. For the analysis that includes 19th-century data I used temperature data published by the National Oceanic and Atmospheric Administration (NOAA) called the "Global Historical Climatology Network daily" (**link**). This dataset contains monthly average temperatures from more than 100,000 weather stations from 180 countries over up to 175 years. I took averages of the temperatures by year and country.

To classify countries as rich or poor I ran a fixed-effect regression using year and country dummy variables and took the median coefficient on countries. Countries with coefficients above the median were classified as rich and countries below the median were classified as poor. I also restricted the sample to countries included in DJO's dataset and used their income classifications and obtained similar negative results. I also checked to see whether a continuous interaction of temperature and income would show an association between temperature and growth that varied by income, and it did not.

Table 9 shows the results. The first column shows the DJO results, the second column shows the results using 19th-, 20th- and 21st-century data, and the third column shows the results using 20th- and 21st-century data.

The alternative data show no statistically significant relationship between temperature and growth. This is true whether all fixed-effect variables from DJO are used, or if a subset of them are used, or if no fixed-effect variables are used at all.

TABLE 9. Alternative data

	DJO	19th–21st centuries	20th–21st centuries
Temperature	0.261	0.085	−0.286
Standard error	0.312	0.070	0.746
t-statistic	0.843	1.211	−0.384
p-value	0.403	0.226	0.701
Temperature x poor	−1.655	0.002	1.171
Standard error	0.485	0.032	3.045
t-statistic	−3.410	0.078	0.385
p-value	0.001	0.938	0.701
Total effect on poor	−1.394	0.087	0.885
Standard error	0.408	0.078	2.665
t-statistic	−3.418	1.121	0.332
p-value	0.001	0.262	0.740
R^2	0.223	0.332	0.418
Observations	4924	933	1339

Temperature and political economy

DJO claim that a mechanism through which temperature affects economic growth is that higher temperatures cause political instability. They use the same method of classifying countries by income as described in the previous sections of this paper. The dependent variables are a measure of political instability, with a dummy variable marking "years when the political system is in flux and no clear political regime has emerged," and another that marks a year in which an "irregular" transition of power occurs, such as a coup. Table 10 shows the original DJO results and the results when different methods of classifying countries by income are used.

For the political instability dependent variable, simply reclassifying South Korea as poor when it was poor and non-poor in other years eliminates the statistical significance of the total association with temperature. The interaction of temperature and the dummy variable indicating whether a country is poor is statistically significant, but the total association, which is the sum of that coefficient and the coefficient of temperature alone is not; the sum just spoken of has a large standard error, and that is why the statistical significance is lost for the total association. The results also disappear when all countries are assigned a single income category based on whether the country is above or below the median in most years of the sample.

TABLE 10. Political-economy effects: political instability

	Dependent variable: Any change in POLITY score			
	DJO	Korea mixed	All mixed income class	Mode income class
Temperature	−0.013	−0.006	−0.000	−0.016
Standard error	0.009	0.008	0.010	0.010
t-statistic	−1.447	−0.684	−0.042	−1.522
p-value	0.148	0.494	0.966	0.128
Temperature × poor	0.040	0.021	0.005	0.037
Standard error	0.016	0.007	0.003	0.018
t-statistic	2.456	3.142	1.572	2.022
p-value	0.014	0.002	0.116	0.043
Total effect on poor	0.027	0.016	0.005	0.022
Standard error	0.015	0.011	0.010	0.017
t-statistic	1.782	1.443	0.486	1.243
p-value	0.075	0.149	0.627	0.214
R^2	0.156	0.156	0.157	0.154
Observations	5388	5388	5388	4734

Table 11 shows the results for irregular changes of government. In this case, if countries are given the same income classification over the entire sample period, the association between temperature and growth remains statistically significant. This is shown in the columns labeled "DJO" and "Mode income class." The association is also statistically significant if South Korea's classification is allowed to change during the sample period, as shown in the column labeled "Korea mixed." If all countries are allowed to change, as shown in the column labeled "All mixed income class," the differential effect of temperature on poor versus non-poor countries disappears, but the total association is still statistically significant. Eliminating only three influential observations, however, is enough to make the total association statistically insignificant at the 10-percent level. The observations are Ecuador in 1976, Paraguay in 1954, and Rwanda in 1994. A coup in Ecuador took place on January 11, 1976. The year 1975 was the second coolest in DJO's sample period of 1951–2003, but 1976 was above the median temperature in the sample. By coding the coup as taking place in 1976, it appeared that the coup took place following warm temperatures, when in fact it followed cool temperatures. In Paraguay in 1954, monthly temperatures were all within 2.5 degrees of the mean temperature for that month from 1951–2003 except for November, which was 4.8 degrees warmer than average. The coup in Paraguay took place in May. As noted earlier, in Rwanda in 1994 the warmest month was September, but the irregular change of power took place in April.

TABLE 11. Political economy effects: irregular changes of government

	Dependent variable: Irregular leader transition				
	DJO	Korea mixed	All mixed income class	Mode income class	Mixed, drop three
Temperature	−0.005	0.009	0.012	−0.006	0.009
Standard error	0.004	0.006	0.006	0.005	0.006
t-statistic	−1.315	1.554	1.990	−1.089	1.364
p-value	0.188	0.120	0.047	0.276	0.173
Temperature × poor	0.050	0.012	0.002	0.064	0.001
Standard error	0.013	0.005	0.002	0.016	0.002
t-statistic	3.898	2.113	0.886	3.920	0.773
p-value	0.000	0.035	0.376	0.000	0.439
Total effect on poor	0.044	0.021	0.014	0.059	0.010
Standard error	0.013	0.007	0.006	0.016	0.007
t-statistic	3.482	2.792	2.166	3.548	1.522
p-value	0.000	0.005	0.030	0.000	0.128
R^2	0.113	0.111	0.108	0.119	0.108
Observations	6677	6677	6677	5427	6674

Conclusion

Each of the following five statements is true as a single, standalone criticism or robustness check of DJO (2012): (1) They use an untenable method of classifying countries by income; using more reasonable methods I find that their results disappear. (2) Their results are influenced by arbitrary methodological choices. (3) Their results are influenced by a small number of observations with unusual characteristics. (4) The inclusion of additional countries and more recent data weakens their results. (5) Alternative data does not support their hypothesis that high temperatures reduce economic growth. Thus, on the matter at hand, namely, whether higher temperatures reduce economic growth, DJO (2012) is not helpful and quite possibly has misled people. It is important that we not be misled on the matter at hand: If climate change does not reduce the rate of economic growth, then any likely effect of warming on the level of economic activity will be outweighed by long-term growth.

DJO's finding that higher temperatures would reduce growth is based on several arbitrary choices of method. One is the classification of countries as rich or poor for the entire 1961–2003 time period. Several countries changed from rich to poor or poor to rich during this period. Allowing them to change during the sample period eliminates DJO's results. Another is the inclusion of fixed-effect variables that control for changes in temperature that are common to many poor countries. If high temperatures reduce growth, they should reduce growth even when they occur in many poor countries at the same time, but these fixed-effect variables control for this effect, causing the remaining influence of temperature to be exaggerated if the effects move in opposite directions, which they apparently do.

Adding additional data by time and country weakens DJO's results, and an alternative data set provides no support to DJO's hypothesis. A mechanism that DJO proposes for temperature to influence growth, political instability caused by heat, also fails robustness tests.

DJO began an important area of research. If higher temperatures significantly reduce the rate of economic growth, and if humans have the ability to lower temperatures, then doing so could greatly increase world wealth. DJO's hypothesis was worth testing, but their tests and those of subsequent authors were flawed, and do not support the hypothesis that higher temperatures reduce economic growth.

Data and code

Data and code used in this research are available from the journal website (**link**).

References

Barker, David. 2022. Temperature and U.S. Economic Growth: Comment on Colacito, Hoffmann, and Phan. *Econ Journal Watch* 19(2): 176–189. **Link**

Barker, David. 2023. Temperature and Economic Growth: Comment on Kiley. *Econ Journal Watch* 20(1): 69–84. **Link**

Belsey, David A., Edwin Kuh, and Roy E. Welsch. 2004. Regression Diagnostics: Identifying Influential Data and Sources of Collinearity. Hoboken: John Wiley & Sons.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. Global Non-Linear Effect of Temperature on Economic Production. *Nature* 527: 235–239. **Link**

Colacito, Riccardo, Bridget Hoffman, and Toan Phan. 2018. Temperature and Growth: A Panel Analysis of the United States. Working Paper 18-09. Federal Reserve Bank of Richmond (Richmond, Va.). **Link**

Colacito, Riccardo, Bridget Hoffman, and Toan Phan. 2019. Temperature and Growth: A Panel Analysis of the United States. *Journal of Money, Credit and Banking* 51(2–3): 313–368.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (DJO). 2009. Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates. *American Economic Review* 99 (2): 198–204. **Link**

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (DJO). 2012. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics* 4(3): 66–95. **Link**

Kadri, Ali. 2016. The Unmaking of Arab Socialism. Anthem Press.

Kiley, Michael. 2021. Growth at Risk From Climate Change. Finance and Economics Discussion Series 2021-054. Washington: Board of Governors of the Federal Reserve System. **Link**

McDoom, Omar Shahabudin. 2020. Contested Counting: Toward a Rigorous Estimate of the Death Toll in the Rwandan Genocide. *Journal of Genocide Research* 22(1): 83–93. **Link**

Smith, Wake, and Gernot Wagner. Stratospheric Aerosol Injection Tactics and Costs in the First 15 years of Deployment. *Environmental Research Letters* 13(12): 1–11. **Link**

Solow, Robert M. 1956. A Contribution to the Theory of Economic Growth. *Quarterly Journal of Economics* 70 (1): 65–94.

Tamura, Robert, Jerry Dwyer, John Devereux, and Scott Baier. 2019. Economic Growth in the Long Run. *Journal of Development Economics* 137: 1–35.

About the Author

David Barker taught economics and finance at the University of Chicago and the University of Iowa. His Ph.D. is from the University of Chicago and he worked as an Economist at the Federal Reserve Bank of New York. He currently runs a real estate and finance company in Iowa and is a member of the Iowa Board of Regents. His email address is drb@barkerapartments.com.

Go to archive of Comments section
Go to September 2023 issue