EJW
ECON JOURNAL WATCH
Scholarly Comments on
Academic Economics

ECON JOURNAL WATCH 21(1)
March 2024: 35–68

Global Non-Linear Effect of Temperature on Economic Production: Comment on Burke, Hsiang, and Miguel

David Barker[1]

LINK TO ABSTRACT

The journal *Nature* published an influential article in 2015 by Marshall Burke, Solomon M. Hsiang, and Edward Miguel (hereafter BHM) purporting to show that higher temperatures will lower economic growth in warm countries. The headline result is that unrestrained global warming will reduce world GDP per capita by 23 percent in the year 2100, approximately nine times larger than the estimate of William Nordhaus (2018).

The Web of Science reports that the paper is in the top six one hundredths of one percent of economics and business publications by citations,[2] and Google Scholar shows 2,269 citations. BHM (2015) also received significant attention in the popular press.[3] Hsiang further developed this work and cowrote a chapter of the National Climate Assessment (Hsiang et al. 2023) claiming that higher temperatures would reduce the rate of economic growth.

BHM's analysis is shallow and misleading. The authors use data with characteristics that are known to create spurious regression results without making proper adjustments or even acknowledging these characteristics. They estimate

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.
2. Based on interpolation of published data.
3. Examples include the *Washington Post* (link), PBS (link), Bloomberg (link), the *Economist* (link), and *Time* (link).

parameters of a quadratic curve relating temperature to growth, and then cherry-pick countries to include in a chart that appears to confirm the shape of this curve. The curve is then used to project growth rates into the distant future using temperature scenarios that a more recent comment in *Nature* described as either "extremely unlikely" or "unlikely" (Hausfather 2020). BHM say that results from the subsamples 1960–1989 and 1990–2010 are similar, indicating that the effects of higher temperatures are not being mitigated, but do not report that the results from subsamples 1960–1990 and 1991–2010 are different.[4] The headline result, that warming will reduce global GDP per capita by 23 percent, is more than double the mean estimate of BHM's bootstrap estimation, which they do not report. BHM claim that their result is "globally representative" (2015, 237), but it does not hold without Greenland and the regions of the Sahara and Central Africa, and it does not hold in large regions of the world.

Simulations support the hypothesis that spatial autocorrelation may be the cause of BHM's results, and robustness checks also suggest that their results may be spurious. BHM has been the subject of methodological criticism (Newell et al. 2021; Tol 2019; Rosen 2019), but my paper is the first to precisely document its deceptive practices.

I begin with a review of previous criticism of BHM, then I describe BHM, and then I discuss problems with their analysis.

Previous criticism of BHM

BHM has received thousands of positive citations and glowing coverage in the popular press. The paper has, however, faced some criticism. In a short letter in PNAS, Richard Rosen (2019) argues that using year dummies in a regression of economic growth is invalid because it ignores the many economic factors that affect growth. Rosen (2019) also argues that using a single equation to estimate optimal temperatures around the world is invalid because the economic conditions of countries are very different. Finally, he argues that equally weighting all countries distorts the results, because very small countries are given the same importance as very large countries. Rosen (2019) describes these issues in his one-page letter, but does not perform independent econometric analysis.

Richard Tol (2019) has written a footnote stating the following:

> The econometrics of Burke et al. (2015) do not stand up to scrutiny. They regress the difference of the log of per capita income, a stationary variable, on

4. BHM's data on GDP per capita begin in 1960, so the first observation of GDP growth is in 1961.

> temperature, a non-stationary variable, and year dummies. As a non-stationary variable cannot explain a stationary one, Burke's year dummies must have de facto detrended temperature. This is indeed the case, as confirmed by their replication package. To the best of my knowledge, the statistical properties of regressing a stationary variable on a cointegrating vector are not known. However, that cointegration vector is measured with error. Errors-in-variables induce bias, of unknown sign in non-linear models (Griliches and Ringstad, 1970). Burke's unusual procedure works fine in-sample, but goes off the rails out-of-sample (Newell et al.) as the year dummies cannot be predicted. (Tol 2019, 556 n.1)

Tol (2019) also does not perform independent econometric analysis of BHM. Richard Newell et al. (2021), as Tol (2019) describes, finds that BHM's model does not predict well out of sample. They test the model's ability to predict out of sample by dividing the data into different portions, estimating the model on one portion of data, and using the model to predict the other portions.

There have also been discussions of BHM on economics blogs. Tol blogged in 2015 that "The graphs are fantastic. The analysis is less impressive" (**link**). Tol also stated that "Burke and co fail to test for spatial autocorrelation, a feature of both economic growth and temperature. Their standard errors are thus downward biased." Marshall Burke replied on his own blog that Tol's criticism about "picking up spurious time trends" was "annoying" (**link**). Burke went on to say:

> If you still think we messed this up, then download our data and show us. The onus is on you at this point, and just making claims about the potential for spurious trends does a disservice to the debate.

In this paper I take up Burke's challenge. I go beyond the work of Rosen (2019), Tol (2019), and Newell et al. (2021) to look carefully at BHM's model and a variety of econometric issues that cause BHM (2015), as Tol (2019) puts it, to go off the rails.

Description of BHM (2015)

BHM (2015) use annual data representing 166 countries from 1961 to 2010 on temperature and economic growth. All countries are equally weighted, and every country is assigned a single average temperature for each year. Because some data are missing, there is a total of 6,584 country/year observations instead of the 8,300 that could be used if data from all years in all countries were available.

The basic idea of BHM (2015) is to estimate GDP per capita growth as a quadratic function of temperature, and then use this function and CMIP5[5] estimates of future temperature increases to estimate year-2100 GDP for each country. BHM find that 77 percent of all countries would be poorer with temperature increases than without increases, and 5 percent of countries would be poorer in 2100 than they are today because of temperature increases. Using more pessimistic IPCC[6] growth scenarios,[7] but without adjusting temperature changes to account for lower growth, BHM unsurprisingly find even larger effects.

The quadratic curve BHM estimate has an inverted U shape, which means that cool countries will grow more rapidly if they warm, countries with moderate temperatures will see little change, and warm countries will grow more slowly. BHM write:

> We find country-level economic production is smooth, non-linear, and concave in temperature (Fig. 2a), with a maximum at 13° C, well below the threshold values recovered in micro-level analyses and consistent with predictions from equation (1). Cold-country productivity increases as annual temperature increases, until the optimum. Productivity declines gradually with further warming, and this decline accelerates at higher temperatures (Extended Data Fig. 1a–g). This result is globally representative and not driven by outliers. (BHM 2015, 236–237)

To find GDP per capita growth as a quadratic function of temperature, BHM regress annual growth of GDP per capita on temperature, temperature squared, precipitation, precipitation squared, and four sets of control variables. The first set of control variables consists of dummy variables for each country. The second consists of dummy variables for each year. The third consists of the dummy variables for each country multiplied by an index value for year, with observations for 1961 taking a value of 1, 1962 a value of 2, and so on. The fourth set of control variables consists of the dummy variables for each country multiplied by the square of the index value for year. Standard errors are clustered by country.

The precipitation variables are not statistically significant at the 5-percent level, and leaving them out of the regression makes very little difference to the

5. The Coupled Model Intercomparison Project is a project of the World Climate Research Programme, which is made up of organizations that are either part of or funded by the United Nations.

6. The Intergovernmental Panel on Climate Change (IPCC) describes itself as "the United Nations body for assessing the science related to climate change."

7. The scenarios, known as Shared Socioeconomic Pathways (SSPs), are developed by the International Institute for Applied Systems Analysis (IIASA), a research institute that is funded by scientific organizations in member countries. The IIASA is a member of the IPCC. The economic growth projections associated with SSPs used in BHM (2015) are developed by the OECD, an intergovernmental organization.

estimated coefficients that are used in subsequent analysis.

The country dummy variables control for the mean of each country's growth over time, and the year dummy variables control for each year's average growth across countries. The rationale for these control variables is to ensure that the estimate of the effect of temperature on growth is not affected by the fact that some countries have persistent differences in growth rates, nor the fact that growth rates in all countries have common trends.

The final two sets of control variables are included to allow the growth rate of each country to have a quadratic trend independent of temperature. For example, if the growth in a country tends to increase or decrease with a concave or convex pattern, or if it follows a U or an inverted U shape, then any measured effect of temperature on growth for a particular country will be relative to this pattern.

Another way of describing the effects of these control variables is that the coefficients on temperature and temperature squared will show the effect of temperature independent of means of growth by country and year, and independent of any country-specific quadratic trend of growth over time. China, for example, had GDP per capita growth of −33 percent in 1962, the second year of the sample. Growth was then weak and variable until the late 1970s, when growth became consistently positive, and high. Because of this trend, the coefficients on the control variables indicating a quadratic trend in growth for China are highly statistically significant. This quadratic trend is assumed to be independent of temperature, and any effect of temperature on growth that is found is relative to this baseline pattern of growth.

This estimated quadratic curve is the basis of all forecasts in the paper. Expected temperature increases under RCP8.5 calculated by the CMIP5 for global grid points from the 1986–2005 average to the 2080–2100 expected average are weighted by population and aggregated by country. Temperature increases for each year are linear interpolations of this overall expected temperature increase. These temperature increases for each country are plugged into the estimated quadratic curve to form a predicted path of GDP per capita for each year until 2100. Some countries show higher income in 2100 than they would have with no temperature increase, while others show lower income. Overall, under the IPCC's RCP8.5 scenario, BHM find that 77 percent of countries will be poorer in 2100 because of temperature increases than if there is no climate change.

Using Burkina Faso as an example, BHM use 28.07 degrees Celsius as the starting temperature for their forecasts that begin in 2010 based on the 1986–2005 average. The starting growth rate is 3.3 percent, based on the SSP3 and SSP5 scenarios with growth projected for each country by the OECD. The temperature in Burkina Faso is expected to increase by 4.5 degrees by the year 2100 based on the RCP8.5 emissions scenario, and so they are expected to increase by 4.5/90, or

0.05 each year between 2010 and 2100. For the first year, the temperatures 28.07 and 28.02 are plugged into equation 1.

$$g = 0.0127184t - 0.0004871t^2 \qquad (1)$$

Equation 1 predicts growth of −2.7 percent without warming and −2.8 percent with warming of 0.05 degrees, for a growth loss of 0.1 percent. This growth is subtracted from the beginning growth rate of 3.3 percent and multiplied by the beginning value of GDP per capita. This process is repeated each year. By the year 2100, the equation predicts growth of −10.3 percent at a temperature of 32.6. BHM cap temperatures at 30 so as not to predict out of sample, so for the year 2100 growth is estimated at a temperature of 30, for a growth reduction of 3 percent. SSP3 and SSP5 predict declining growth rates for Burkina Faso without climate change, down to 1.7 percent per year by the year 2100, so the BHM model predicts growth of −1.3 percent.

Burkina Faso is well above the estimated optimal temperature of 13, so temperature increases reduce growth. For a cold country such as Iceland, the quadratic curve predicts that temperature increases will increase growth. For a country near 13 degrees Celsius, like France, a small change in temperature would have little effect on economic growth.

Using the Stata and R code provided on BHM's website, I was able to exactly replicate the results of the paper. Some data is common to that of Melissa Dell et al. (2012), and some errors, such as the inadvertent omission of Burma, are also common to both papers.

BHM present the quadratic curve differently than simply plotting equation 1. Figure 1 shows BHM's Figure 2a, which is the estimate of the total effect of temperature on growth of GDP per capita over the period 1960–2010.[8] The blue area is a 90-percent confidence interval.

Growth of GDP per capita peaks at approximately 13 degrees Celsius. Projected GDP per capita is close to 15 percent lower for cold and hot countries than for countries near the temperate optimal. The figure may remind one of Aristotle's view of virtue as moderation between extremes.

To produce the curve in Figure 1, BHM calculated, for each observation and using values and coefficients for all independent variables in the regression, predicted growth when temperatures were allowed to range from −5 to 35 degrees Celsius. For each temperature in that range, they took the average predicted growth across observations. More specifically, they set the temperature at −5, −4, etc., and calculated the predicted growth for Argentina in 1962, Canada in 1970, etc., and

8. BHM's data on GDP per capita begin in 1960, so the first observation of GDP growth is in 1961.

averaged them all. The result is overall predicted growth for each temperature with each country weighted equally. The pattern of growth with temperature calculated in this way is fashioned using the estimated coefficients of growth on temperature and temperature squared, and so it must, by construction, follow a smooth quadratic path. The estimated regression coefficients result in this curve having an inverted U shape, with growth peaking at approximately 13 degrees Celsius. Any data, no matter how noisy, will generate a smooth quadratic curve if one variable is regressed on another and its square and the predicted values of the dependent variable are plotted against possible values of the independent variable.

Figure 1. Quadratic relationship between growth and temperature from BHM (2015, 236 Figure 2a)

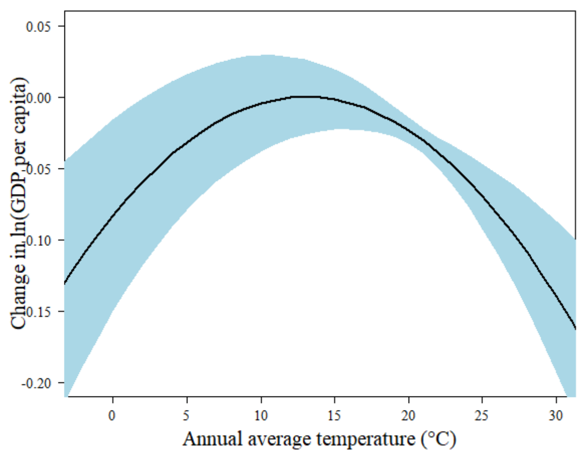

The headline result of a 23 percent reduction in GDP comes from taking each country's projected GDP per capita with and without climate change, then taking the weighted average by population, and then taking the percentage difference between the weighted sum with and without climate change.

Problems with BHM's analysis

An overly ambitious graph

As Tol (2019) commented, the graphs in BHM are impressive. One of them, BHM's Extended Data Figure 1, panels a–f, is reproduced here as Figure 2. The graph appears to confirm the quadratic relationship between growth and temperature described in the previous section in Equation 1 and Figure 1. BHM estimate the quadratic relationship using the full panel of 6,584 country/year observations, but Figure 2 shows the average relationship between growth and temperature for 166 countries to illustrate the overall pattern. It also highlights five countries that clarify the pattern. The overall pattern, however, is not statistically significant, and the five countries are cherry-picked to make the relationship appear to be significant. While the figure is not a crucial part of BHM's analysis, it is indicative of the misleading approach of the paper, and suggests alternative methods of measuring the relationship between growth and temperature.

Panel (a) at the top of Figure 2 shows the quadratic curve discussed earlier and temperature ranges of five countries on different continents. The panels b–f show individual country data on what BHM call the marginal effect of temperature fluctuations on growth. To calculate this marginal effect for each country, separate regressions for each country are run that are similar to the regression on all countries that generated the quadratic curve shown in Figure 1, but of course the control variables that are based on country dummy variables cannot be included on regressions run separately for each country. The regressions include precipitation, year, and year squared. For reasons that are not explained, precipitation squared is not included. The plots are produced as follows:

1. Regress growth on precipitation, year, and year squared, but not temperature.
2. Regress temperature on precipitation, year, and year squared.
3. Draw scatter plots of the residuals from #1 and the residuals from #2.

Figure 2. Reproduction of Extended Data Figure 1 panels a–f from BHM (2015) showing relationship between growth and temperature for countries

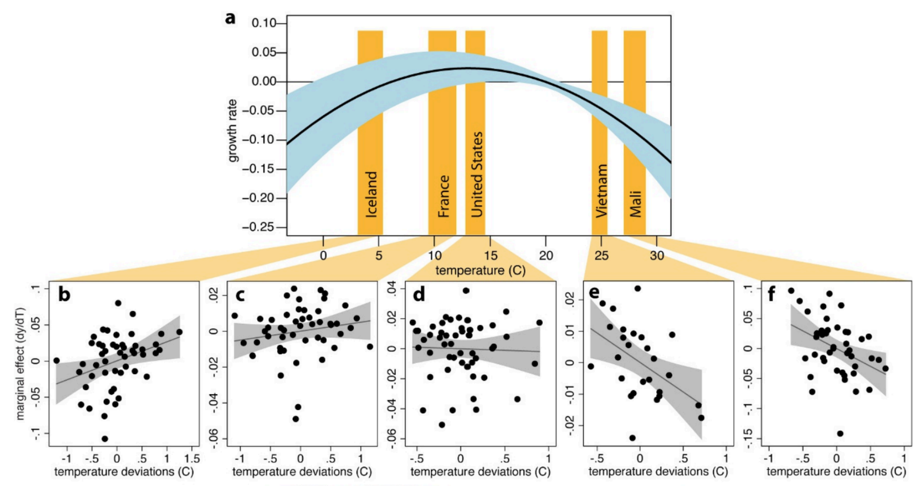

The resulting scatterplot points are detrended growth and temperature, controlled for precipitation. For Iceland, the coldest country shown, the slope of points is upward, corresponding to the point on which Iceland lands on the quadratic curve, meaning that higher temperatures were associated with higher growth. For France and the United States, near the maximum of the quadratic curve, the slopes are much flatter. For the warmer countries, Vietnam and Mali, the slopes are downward.

Panel (h) of BHM's Extended Data Figure 1 shows the slopes for all 166 countries, and adds an OLS line showing a downward slope. I recreated the scatterplot in panel (h) as Figure 3, and highlighted the five countries selected by BHM, as BHM also do in panel (h). A negative slope of the OLS line shows that the slope is positive and high for cold countries, near zero for intermediate countries, and negative for warm countries. Figure 3 shows these slopes plotted against temperature for all countries. The examples shown in Figure 2, Iceland, France, the United States, Vietnam and Mali, are shown in red, while other countries are shown in blue.

Figure 3. Slopes of growth-temperature regressions by country

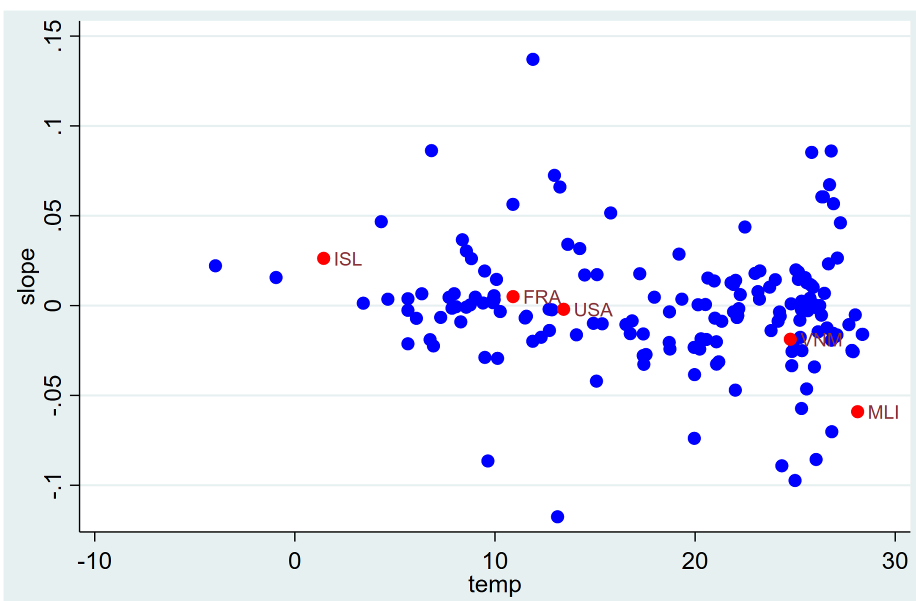

Panels a–f are misleading because they cherry-pick five countries that make it appear that slope is strongly negatively related to temperature, when the actual relationship is quite noisy. The correlation of slope and temperature for all countries is −0.13, while the correlation for the five selected countries is −0.94. The country-level regression coefficient on temperature for the five selected countries is 74.5 percent higher than the coefficient obtained by running the regression on all countries. The regression slope for all countries is not statistically significant at the 10-percent level. A regression test of a quadratic relationship shows no significance of either temperature or temperature squared.

A closer examination of the data shown in Figure 3 reveals that the very weak relationship between slope and temperature is not consistent over the range of temperature. The median temperature of the 166 countries is 20.6 degrees, while the range of Figure 3 is from −10 to 30 degrees, so most countries are on the far-right side of the graph. For countries above the median, an OLS line would have a positive slope, although it is statistically insignificant. In other words, for the warmest half of all countries, the relationship that BHM highlight in their Extended Data Figure 1 does not exist. This is particularly important, since it is only in the warmest countries that BHM claim that warming temperatures will reduce growth.

Figure 4. Local polynomial regression of slope and temperature

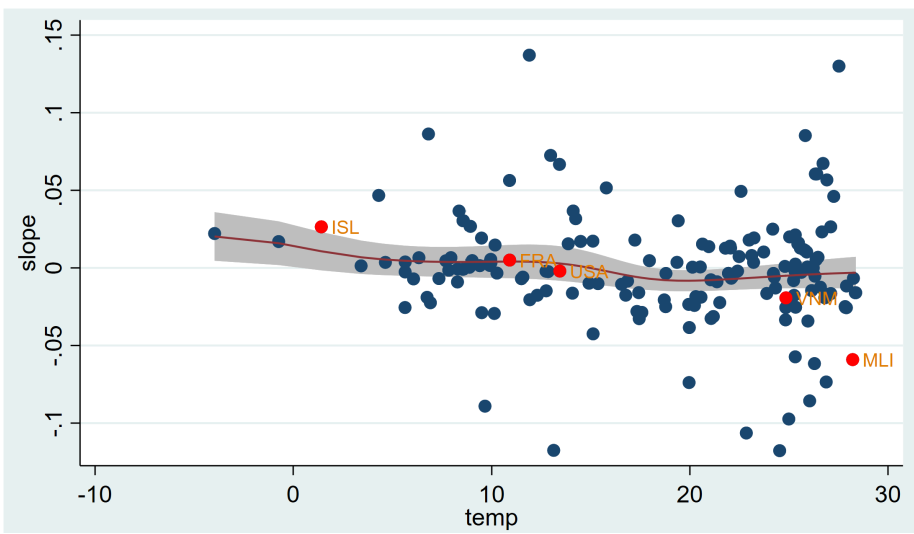

Figure 4 shows a local polynomial regression of the same data from Figure 3. The gray band represents a 95-percent confidence interval. With the exception of the two coldest countries, Greenland and Mongolia, the slope is not statistically different from zero at any temperature. In other words, for nearly all countries, these data do not show a statistically significant relationship between growth and temperature. For the warmest half of all countries, those with temperatures above 20.6 degrees, the slope increases with temperature. Three of the five countries selected by BHM to illustrate the relationship are outside of the 95-percent confidence interval, and the other two are at the visual midpoint of the graph, creating the impression of a downward relationship across all temperatures.

For the purposes of their Extended Data Figure 1, BHM examine the relationship between detrended growth and detrended temperature. In their main regression analysis, however, BHM use unadjusted growth and temperature with fixed effect controls. I will show later that BHM's results disappear if detrended growth and temperature are used. BHM do not report this result.

It is worth repeating that this figure is not a crucial part of BHM. I discuss it here only as an illustration of the character of the presentation of data in BHM. In the following sections I examine BHM's data and results in greater detail.

Data characteristics

The full panel of data that BHM use to estimate the quadratic curve shown in Figure 1 present a number of difficulties for estimation. The measures of GDP per capita growth contain extreme outliers, with annual growth ranging from a 70 percent drop in GDP per capita in Liberia in 1990 to an 88 percent increase in Equatorial Guinea in 1997. All countries are equally weighted, so unusual growth in a small country can have a large effect on the results. The panel data is unbalanced because data are missing in many years in some countries. Because of the changing mix of countries over time, the average yearly temperature across countries falls from 1960 to 1996 and then increases. Average growth generally falls until 1992 and then generally rises.

Residuals from the basic regression that produces the quadratic curve show a skewness coefficient of −0.67 and kurtosis of 28.4. That skewness level is high, and that kurtosis level is extreme. With a very high level of confidence, the residuals fail a variety of normal normality tests. Failures of these tests do not invalidate BHM's results, but they should suggest additional caution in evaluating them. BHM do not mention anything about the distribution of their residuals.

The Variance Inflation Factors (VIF) in the basic regression are 1,310 for temperature and 1,484 for temperature squared. A rule of thumb for VIF is that a value larger than ten is cause for concern over multicollinearity (Forthofer 2007). The VIFs here, then, are more than 130 times higher than the threshold for cause for concern.

A Wooldridge test for autocorrelation (Wooldridge 2002) of the residuals from BHM's main regression, a test that takes account of the panel structure of the data, rejects the null hypothesis of no autocorrelation with a very high level of confidence. The F statistic is 28.6, while the critical value at a 99-percent level of confidence is 6.1. The same test for temperature produces a test statistic of 276.5. In the presence of a unit root, regressing an autocorrelated variable on another autocorrelated variable can produce spurious findings in the sense that the results appear to be statistically significant even when there is no underlying relationship between the two variables (Granger 1974).

Spurious regression results can also be caused by spatial autocorrelation (Muller 2023). Spatial autocorrelation also exists in BHM's data. For each country in the sample, I determined the nearest neighboring countries and regressed annual GDP per capita growth on growth in the neighboring countries, as shown in Table 1. The first five spatial lags were all statistically significant at a 95-percent or higher level of confidence. Regressing temperature on its spatial lags produced even stronger results.

TABLE 1. Regressions of growth and temperature on spatial lags

	Growth			Temperature		
Spatial lag	Coeff.	T-stat	P-value	Coeff.	T-stat	P-value
1	0.152	5.70	0.000	0.295	12.44	0.000
2	0.044	2.21	0.027	0.372	11.14	0.000
3	0.049	2.03	0.043	−0.072	−2.98	0.003
4	0.108	4.54	0.000	0.207	9.56	0.000
5	0.085	3.05	0.002	0.231	6.57	0.000
Constant	0.012	8.11	0.000	−1.725	−5.41	0.000

The data are also heteroskedastic, meaning that the variance of economic growth is very different for different countries. France, for example, has a standard deviation of growth of 1.9 percent, while Liberia has a standard deviation of 19.3 percent. Other than heteroskedasticity and temporal autocorrelation, none of these data characteristics are discussed in BHM, other than to say that "this result is globally representative and not driven by outliers" (2015, 236–237). The only corrections applied to the data are the fixed-effect control variables discussed earlier and the adjusting of standard errors for clustering by country to account for heteroskedasticity. These corrections make little difference to the results. A simple regression of growth on temperature and temperature squared with no controls and no corrections of standard errors shows the same quadratic relationship that is statistically significant, with an optimal temperature of 10.3 instead of 13.0 degrees.

BHM mention the spurious regression problem in their supplementary materials (**link**, p. 39), but only with regard to temporal serial correlation. There is no mention of potential spatial or other cross-sectional autocorrelation that might result in unit roots that could produce spurious results.

BHM do perform robustness checks for temporal autocorrelation by including lags of the dependent variable, GDP per capita growth. A graph in their Extended Data Figure 2 shows that including lags dramatically widens the confidence interval around the effect of temperature on growth, but the finding is not discussed in the text of the paper. In their Extended Data Table 1, columns 9 and 10 also show that the inclusion of lags of the dependent variable substantially reduces the estimated magnitude of the effect of temperature on growth and its statistical significance, but, when including three lags, the result is still practically and statistically significant. This robustness check, however, says nothing about the possibility of spatial or other cross-sectional autocorrelation. Spatial autocorrelation is more likely to be a problem than temporal autocorrelation, since there are 166 countries in the sample and only 50 years. The sample is unbalanced, so some countries have as few as eight years of observations.

The presence of a unit root in economic growth or in temperature could cause spurious regression results. The measure of growth used by BHM is the first difference of GDP per capita, so in a single-country regression growth would not be expected to contain a unit root. In panel data with multiple countries, however, a unit root in economic growth is possible. Using the Levin-Lin-Chu (LLC) test for a unit root in panel data (Levin et al. 2002) that takes account of cross-sectional correlation, a unit root in economic growth cannot be rejected, and a unit root in temperature also cannot be rejected. The failure to reject occurred using a variety of lag structures and testing methods. The null hypothesis of the LLC test is that all panels contain unit roots. Even if this null hypothesis is rejected, it is possible that enough panels contain unit roots to cause spurious regression results. The Hadri test (Hadri 2000) has as a null hypothesis that all panels are stationary. This hypothesis is rejected, again using a variety of lag structures and testing methods for both growth and temperature. Both tests require balanced panels, so I only used countries with 50 years of data, which eliminated 80 out of 166 countries.

Unit roots in growth could be the result of the effects of location on growth that are independent of temperature. When latitude and longitude are added to BHM's primary regression with growth of GDP per capita as the dependent variable, they are highly statistically significant, even though temperature and rainfall are controlled for. These results are shown in Table 2. When temperature is regressed on latitude and longitude, rainfall, and all of the fixed effect control variables in BHM's primary regression, latitude and longitude are extraordinarily statistically significant. Rainfall and the control variables are also highly significant. Table 3 shows that the effects are nonlinear, with latitude squared significant in the growth regression, and both latitude and longitude squared and the product of latitude and longitude significant in the temperature regression. BHM's analysis clearly has an endogeneity problem, and there is a possibility of spurious results coming from regressing dependent and independent variables that may contain unit roots.

TABLE 2. Growth and temperature regressed on latitude and longitude

	BHM			Growth			Temperature		
	Coeff.	T-stat	P-value	Coeff.	T-stat	P-value	Coeff.	T-stat	P-value
Temp	0.013	3.358	0.001	0.013	3.301	0.001			
$Temp^2$	0.000	−4.114	0.000	0.000	−4.018	0.000			
Lat				0.013	2.596	0.010	1.249	19.838	0.000
Long				0.000	2.700	0.008	0.210	10.164	0.000
Rain	0.000	1.440	0.152	0.005	1.527	0.129	−0.001	4.335	0.000
$Rain^2$	0.000	−1.861	0.065	0.000	1.423	0.157	0.000	9.917	0.000
R^2	0.286			0.288			0.997		
Obs	6584			6519			6519		

TABLE 3. Growth and temperature regressed on latitude and longitude with nonlinearity

	BHM			Growth			Temperature		
	Coeff.	T-stat	P-value	Coeff.	T-stat	P-value	Coeff.	T-stat	P-value
Temp	0.013	3.358	0.001	0.013	3.301	0.001			
$Temp^2$	0.000	−4.114	0.000	0.000	−4.018	0.000			
Lat				0.013	2.596	0.010	1.249	19.838	0.000
Lat^2				0.000	2.700	0.008	0.021	10.164	0.000
Long				0.005	1.527	0.129	0.210	4.335	0.000
$Long^2$				0.000	1.423	0.157	0.003	9.917	0.000
Lat×Long				0.000	0.880	0.380	−0.012	−18.192	0.000
Rain	0.000	1.440	0.152	0.000	1.447	0.150	−0.001	−4.896	0.000
$Rain^2$	0.000	−1.861	0.065	0.000	−1.873	0.063	0.000	3.878	0.000
R^2	0.286			0.288			0.997		
Obs	6584			6519			6519		

The quadratic relationship between growth and temperature

Basic robustness checks

All of BHM's results depend on the quadratic function of growth that they estimate with respect to temperature using these data. The growth~temperature relationship is specified in Equation 1 and illustrated in Figure 1. BHM estimate Equation 1 by regressing growth on temperature and control variables using panel data consisting of annual observations from 166 countries. Average temperatures vary considerably between countries, and temperatures within in each country vary by year, but not as much as country averages differ from each other. The data characteristics described in the previous section recommend some basic robustness checks on the estimated quadratic relationship between temperature and growth.

Weighting

Table 2 shows BHM's results, along with results from weighting residuals to be minimized by population to see whether small countries have a disproportionate effect on BHM's estimates. In order to calculate country average annual temperatures from gridded temperature data, BHM weight the gridded temperatures by gridded population, but countries themselves are not weighted by population in their analysis. Weighting at the sub-country level but not at the country level seems arbitrary, and seeing whether population weighting of residuals affects BHM's results would be a reasonable robustness check. While weighting by population creates its own problems by allowing large countries to dominate the results, substantial differences between population-weighted and non-weighted regression are indicators of potential problems. Results that are driven by small countries may be irrelevant for most of the world.

In addition, to see if heteroskedasticity affects the results, it seems reasonable to weight residuals by the inverse of the variance of growth by country. The second and third columns of numbers in Table 4 show that these weightings substantially reduce the statistical significance of BHM's results. Weighting by population and variance increases the p-value of temperature and temperature squared from less than one tenth of one percent to over 10 percent.

TABLE 4. Robustness checks for BHM results

	BHM	Pop weighted	Pop, var wgt	Growth lags	W/o 2 obs.
Temp	.0127	.0136	.00602	.00345	.000231
Se	.00379	.0087	.00389	.00412	.00443
T-stat	3.36	1.56	1.55	.838	.0521
P-value	.000975	.12	.124	.403	.959
Temp^2	−.000487	−.000499	−.000223	−.000124	.0000424
Se	.000118	.00024	.000141	.000152	.000182
T-stat	−4.11	−2.08	−1.58	−.816	.233
P-value	.0000612	.0389	.116	.416	.816
Opt. temp	13.1	13.6	13.5	13.9	−2.72
Temp diff	.235	.228	.103	.0552	−.0603
R^2	.286	.454	.425	.471	.488
Obs	6584	6584	6584	6252	6250

Lagged dependent variables

The fourth column of numbers shows the regression results with the inclusion of first and second lagged values of growth as independent variables. BHM allow growth to have quadratic trends, but in their main specification they do not allow for an autocorrelated growth process. Their Extended Data Figure 2 shows that confidence intervals for the effect of temperature are wider with the addition of lagged growth. Moreover, in their Extended Data Table 1, columns 9 and 10 show a 50-percent reduction in the coefficient on temperature and a reduction in statistical significance from the 1-percent to the 10-percent level, and a 40-percent reduction in the coefficient on temperature squared. But in the text of the article BHM only say that their results are "robust to estimation procedures that…account for multiple lags of growth" (2015, 237). In their Supplementary Materials (**link**) they say that "our main result is robust across models that use alternative set[s] of controls" (p. 43).

The results in Table 4 show that in the weighted regression, the addition of two lags of growth eliminate the statistical significance of temperature and reduce the effect of changing temperature from the optimal level to 35 degrees by 76 percent. Dropping just two out of 6,252 observations—namely, India 1979 and Indonesia 1998—changes the sign of the coefficient on temperature squared, implying a positive effect of rising temperature, although with no statistical significance. In 1979 GDP per capita in India dropped by 7.7 percent, the worst year for India in BHM's sample. The year 1979 was the 7th warmest year in BHM's sample of 50 years. The decline in India's GDP can be attributed to a balance of payments crisis triggered by oil price increases. Indonesia's GDP per capita dropped by 15.5 percent in 1998, by far the worst in the sample. The year 1998 also happened to be the warmest year in the sample of 50 years. The 1998 drop in GDP can be attributed to the Asian financial crisis.

Weighting by country size and volatility, adding temporal and spatial lags of growth, and dropping two outlier observations out of 6,252 are reasonable model modifications that completely eliminate BHM's results.

Country effects vs. temporal effects

A natural question to ask is how much of the measured effect of temperature on growth in BHM comes from differences in country averages versus annual fluctuations within countries. Dell et al. (2012) and Michael Kiley (2021) explain that observers have for centuries noted an association between the long-term average temperature of countries and economic development, which may be due to "spurious associations of temperature with national characteristics such as institutional quality" (Dell et al. 2012, 66). Country fixed-effect variables are included in their models of the effect of temperature on growth in order to ensure that their results are due to short-term temperature fluctuations, not "spurious associations" with long term averages. Kiley (2021, 4) says about the inclusion of fixed effect variables that "this specification eliminates the 'permanent' component of weather, and hence may control for concerns regarding the link between the average temperature and the level of income across countries." BHM explain their use of country fixed effects as follows:

> We deconvolve economic growth to account for: (1) all constant differences between countries, for example, culture or history; (2) all common contemporaneous shocks, for example, global price changes or technological innovations; (3) country-specific quadratic trends in growth rates, which may arise, for example, from changing political institutions or economic policies; and (4) the possibly non-linear effects of annual average temperature and rainfall. (BHM 2015, 236)

If random temperature fluctuations affect economic growth independently of an association of average country temperature with average growth, then demeaned growth should show a relationship with demeaned temperature. In simulated data, if random temperature fluctuations affect growth while at the same time there is a correlation between average growth and average temperature, then fixed-effect controls remove the effect of average temperature on average growth, and the effect of temperature on growth is statistically significant. In addition, a regression of demeaned growth on demeaned temperature will show a statistically significant effect. If in simulated data there are quadratic trends in growth and/or temperature, then BHM's fixed-effect controls will remove the effects of these trends and show any true effect of random temperature fluctuations affecting economic growth. I find, however, that after demeaning and detrending growth and temperature, there is no effect of temperature on growth.

In Table 5, the first column of numbers shows BHM's results. The second column shows the results using as the dependent variable the demeaned growth, that is, the temperature for a country in a particular year minus the mean temperature of the country over the entire sample period. The third column shows the results from removing the mean and country-specific trend from growth. In other words, the dependent variable is the residual of a regression of growth on country fixed-effect variables, year fixed-effect variables, and interacted country fixed-effect variables (one set of them multiplied by time and another set multiplied by time squared). Time is an index, with 1961 equal to one, 1962 equal to two, etc.

The ninth row of numbers shows the optimal temperature calculated from the regression coefficients, and the tenth row shows the difference in the growth rate from the optimal temperature to 35 degrees Celsius, which is approximately the maximum projected 2100 temperature for any country.

TABLE 5. Regression results for quadratic relationship between growth and temperature

	BHM	Demeaned	Detrended	Clustered	Weighted	W/o 1 obs
Temp	.0127	.00368	.000202	.000202	.000121	.0000614
Se	.00379	.00159	.00039	.0000625	.000044	.0000657
T-stat	3.36	2.32	.519	3.24	2.74	.935
P-value	.000975	.0216	.604	.00145	.00681	.351
$Temp^2$	−.000487	−.000109	−6.26e−06	−6.26e−06	−3.46e−06	−1.74e−06
Se	.000118	.0000538	.0000116	1.88e−06	1.50e−06	2.05e−06
T-stat	−4.11	−2.03	−.54	−3.33	−2.3	−.849
P-value	.0000612	.0437	.589	.00107	.0228	.397
Opt. temp	13.1	16.8	16.2	16.2	17.4	17.6
Temp diff	.235	.0361	.00222	.00222	.00107	.000527
R^2	.286	.147	.000179	.000179	.0000583	.0000329
Obs	6584	6584	6584	6584	6584	6583

The calculated optimal temperature is significantly higher once country means are removed, and the effect of temperature on growth is reduced by 84.6 percent. In BHM's results, controlling for rainfall, fixed effects and trends in growth, annual growth is 23.5 percentage points lower in very hot countries compared to countries at the optimal temperature. Demeaning growth lowers this estimate to 3.6 percentage points. Removing growth trends from the dependent variable lowers the estimate to 0.22 percentage points. Removing growth trends also reduces R^2 from 28.6 percent to 0.018 percent.

Weighting residuals by the inverse of variance reduces the estimate to 0.1 percentage points. In the weighted regression, this small effect is statistically significant, but removing one out of the 6,584 observations, Greenland 1990, eliminates this significance. Greenland is the coldest country in the sample, and 1990 was the worst year for growth in that country. In 1990 the Black Angel mine in Greenland closed after financial losses that began in 1985 and exhaustion of extractable reserves, not because of the temperature in 1990. Annual production had been rising since the 1970s, and suddenly stopped when the mine was closed (Thomassen 2003). Gross sales from the mine were $112.6 million in 1989, *12.1 percent of Greenland's GDP at the time*. Greenland's GDP per capita fell by 13.0 percent in 1990. The ocean fishing harvest was also significantly lower in 1990 than in 1989 (Booth and Knip 2014). The ocean fishing harvest seems unlikely to have been affected by air temperatures on land.

The measured effect of temperature on growth appears to be primarily the result of differences in country averages, not annual temperature fluctuations within countries. When country-specific quadratic trends are also removed, the effect of temperature on growth is reduced to statistical insignificance.

Interestingly, when standard errors are adjusted for clustering, these small effects are still statistically significant, but with standard OLS regression they are not. This difference can be explained by heteroskedasticity of growth by country. Adjusting standard errors for clustering normally increases estimated standard errors, because variation within panels is normally smaller than variation across panels. Once the mean of growth is removed, there is more variation within some countries than between countries, leading the adjustment for clustering to lower standard errors and inflate the statistical significance of temperature. Weighting residuals by the inverse of the variance of growth of each country helps to control for heteroskedasticity, which reduces the cluster-adjusted statistical significance of temperature.

We can also demean and detrend temperature. Table 6 shows the results. Statistical significance is eliminated, although the coefficient estimates suggest extremely large effects. Weighting by the inverse of variance flips the sign of the effect. Similarly to the results in Table 5, when growth and temperature are detrended, R^2 is significantly reduced.

A regression of growth on temperature for 166 country averages produces a weak quadratic relationship between temperature and growth. Using instead, as BHM do, up to 50 annual observations for each country, country averages are replicated over and over, inflating statistical significance without adding additional information.

TABLE 6. Regression results for quadratic relationship between growth and temperature

	BHM	Demeaned	Detrended	Weighted
Temp	0.0127	−0.00141	−0.00086	0.000748
Se	0.00379	0.0022	0.00161	0.000928
T-stat	3.36	−0.642	−0.532	0.806
P-value	0.000975	0.522	0.595	0.42
Temp^2	−0.00049	−0.00522	−0.00155	0.00189
Se	0.000118	0.00275	0.00196	0.00105
T-stat	−4.11	−1.9	−0.791	1.81
P-value	6.12e−05	0.0589	0.429	0.0699
Opt. temp	13.1	−0.135	−0.276	−0.197
Temp diff	0.235	6.45	1.93	−2.35
R^2	0.286	0.147	0.000262	0.000559
Obs	6584	6584	6584	6584

To ensure that the ‘generated regressor problem’ is not causing the reduction in the effect of temperature on growth, I generated bootstrapped standard errors using bootstrap samples for both stages of the regression. The results are similar.

Country averages

The previous section showed that BHM's results come largely from country-by-country averages, not annual temperature fluctuations. But is there truly a relationship between average country growth and average temperature? There is some sign of such a relationship, but the evidence is statistically weak, as shown in Table 7.

TABLE 7. Regression of average growth on country average temperature and temperature squared

	OLS	Var wgt	Pop wgt	Both	Pop w/o 3
Temp	0.000879	0.000249	0.00662	0.000669	0.000896
Se	0.000884	0.000696	0.00129	0.00107	0.000862
T-stat	0.995	0.357	5.12	0.627	1.04
P-value	0.321	0.721	8.38e−07	0.532	0.3
$Temp^2$	−0.0000427	−0.0000179	−0.000202	−0.0000129	−0.0000121
Se	0.0000264	0.0000211	0.0000363	0.0000301	0.000025
T-stat	−1.61	−0.849	−5.56	−0.427	−0.485
P-value	0.109	0.397	1.07e−07	0.67	0.628
Opt. temp	10.3	6.93	16.4	26	37
Temp diff	0.026	0.0141	0.0698	0.00104	0.0000484
R^2	0.0575	0.0402	0.179	0.0144	0.0702
Obs	166	166	166	166	163

The first column of numbers shows the results from unweighted regression of country average temperature and temperature squared on average growth. Temperature squared is just short of being statistically significant at the 10-percent level. The coefficient estimates are consistent with a quadratic relationship with an optimal temperature of 10.3 degrees, reasonably close to BHM's estimate. However, if residuals by country are weighted by the inverse of variance of growth, the significance is eliminated. Weighing by population produces very significant coefficient estimates, and looming large is China, which has a large population, mid-level temperatures, and high growth. Removing China, Japan, and South Korea eliminates the significance of temperature squared, as shown in the last column of numbers. The fourth column of numbers shows the results of weighting by both population and the inverse of variance of growth, where the significance of temperature is also eliminated.

Table 8 shows the results of the same regressions, but without temperature squared, to see if there is a simple linear relationship between country average temperatures and growth. The row labeled "Temp Diff" shows the difference in predicted growth when the average country temperature is −5 versus 35 degrees.

TABLE 8. Regression of average growth on country average temperature

	OLS	Var wgt	Pop wgt	Both
Temp	−0.000513	−0.000332	−0.000465	0.000218
Se	0.00019	0.000134	0.000235	0.000146
T-stat	−2.7	−2.47	−1.98	1.49
P-value	0.00775	0.0144	0.0493	0.139
Temp diff	0.0205	0.0133	0.0186	−0.00871
R^2	0.0424	0.036	0.0234	0.0133
Obs	166	166	166	166

In the unweighted regression, temperature has a negative relationship with growth, but the effect is reduced when residuals by country are weighted by population or the inverse of variance, and the sign reverses when residuals are weighted by both. This is primarily due to high-population countries with moderate, stable growth and low temperatures, such as Germany, France, the United States, and Japan, and other countries with high populations, high and reasonably stable growth, and high temperatures, such as Pakistan, India, Indonesia, and Vietnam.

Regional effects

Figure 5. Map of large world regions

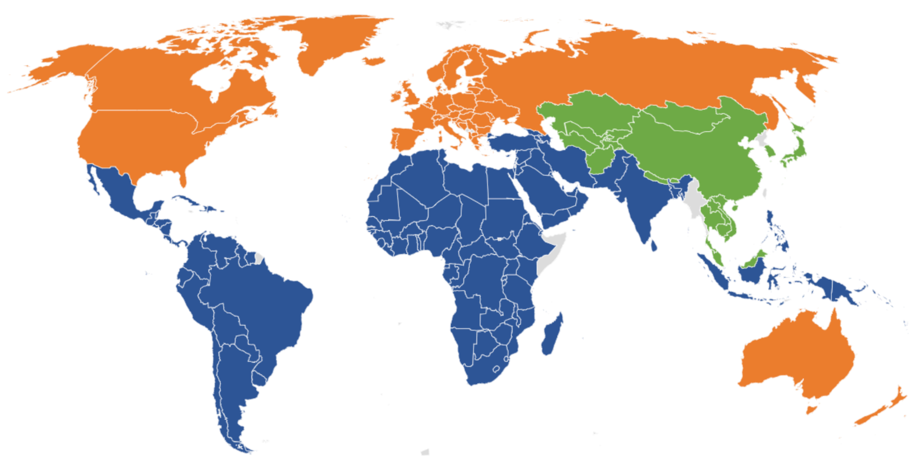

The influence of China, Japan, and South Korea on the results in Table 7 suggests that regional factors may play a part in the measured relationship between temperature and growth. In this section, I first divide the world into three large regions, which I will call Orange, Green, and Blue, following the colors on the map in Figure 5. Orange includes North America except for Mexico, Europe including Russia, and Australia and New Zealand. Green includes continental countries in eastern, southeastern, and central Asia, not including southwestern Asian countries on the Indian Ocean. All other countries are included in Blue. I tested different regions to see if BHM's results are consistent throughout the world.

Orange is relatively cool and grew at a moderate pace from 1960–2010. Green has intermediate temperatures and grew quickly, and Blue is relatively warm and grew slowly.

Data for Somalia and some very small countries are missing, and Burma is missing because of a coding error in BHM that came from Dell et al. (2012). Figure 6 shows the average growth and temperatures of the three regions over the period 1960–2010 with a quadratic curve fitting the three points, and Figure 5 shows a map of these regions. Figure 6 is constructed with actual growth rates and temperatures without adjustment for the fixed effects variables in BHM.

Figure 6. Fitted quadratic relationship between growth and temperature for three world regions

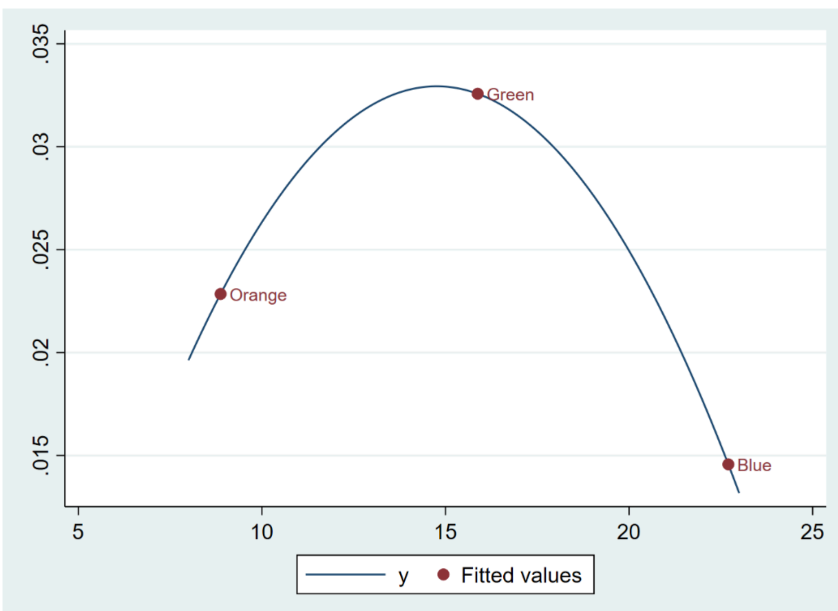

Figure 6 shows that a quadratic relationship between growth and temperature can appear to be the case, even when other factors are far more important than temperature. The relative unimportance of temperature is clear from more distant history. China and surrounding countries grew more slowly than Europe, North America, and Australia during the 19th and early 20th centuries (Pomeranz 2001), and GDP of Africa and South America grew rapidly during the early 20th century. In other words, relative growth of the regions has changed over time, as relative temperatures have changed little.

TABLE 9. Regression results by region

	BHM	Orange	Green	Blue	Green w/o 98	Blue w/o 2
Temp	0.0127	0.00721	0.0372	0.0266	0.0336	0.00237
Se	0.00379	0.00304	0.0172	0.0212	0.0172	0.0132
T-stat	3.36	2.37	2.17	1.25	1.95	0.179
P-value	0.000975	0.0224	0.0457	0.213	0.0685	0.858
$Temp^2$	−0.00049	−0.0003	−0.00093	−0.00077	−0.00044	−0.00024
Se	0.000118	0.000236	0.000387	0.000474	0.000419	0.000312
T-stat	−4.11	−1.28	−2.4	−1.63	−1.05	−0.783
P-value	6.12e−05	0.209	0.0288	0.107	0.312	0.435
Opt. temp	13.1	12	20	17.2	38.4	4.85
Temp diff	0.235	0.159	0.21	0.245	0.00498	0.222
R^2	0.286	0.511	0.477	0.235	0.499	0.232
Obs	6584	1572	546	4466	530	4464

Table 9 shows the regression results from BHM for the entire world and for the three regions. The effect of temperature on growth that BHM reports for the entire world is highly statistically significant. Within regions, however, the relationship between growth and temperature is weak and the coefficient estimates are unstable. Removing 1998, the year following the Asian financial crisis, from the regression for Green eliminates the statistical significance of temperature. Removing two out of 4,466 observations from Blue eliminates the statistical significance of temperature. The observations removed are the two neighboring countries of Georgia and Armenia in 1992, a year of war and upheaval for both.

Another example of the lack of consistency of the effect of temperature on growth by region can be seen by removing Greenland and central Africa from the sample. Table 10 shows the results of BHM, results with Greenland and central Africa removed, and then also removing the year 1992, a year of economic turmoil in many countries following the collapse of the USSR. Figure 7 shows a map of the countries removed.

TABLE 10. Regression results by region

	BHM	−GL, AF	−1992
Temp	0.0127	0.00736	0.00549
Se	0.00379	0.00437	0.00336
T-stat	3.36	1.68	1.63
P-value	0.000975	0.095	0.105
Temp^2	−0.000487	−0.000187	−0.000111
Se	0.000118	0.00013	0.000116
T-stat	−4.11	−1.44	−0.962
P-value	0.0000612	0.152	0.338
Opt. temp	13.1	19.6	24.7
Temp diff	0.235	0.0441	0.0118
R^2	0.286	0.303	0.29
Obs	6584	5503	5373

Figure 7. Greenland and Saharan and Central Africa

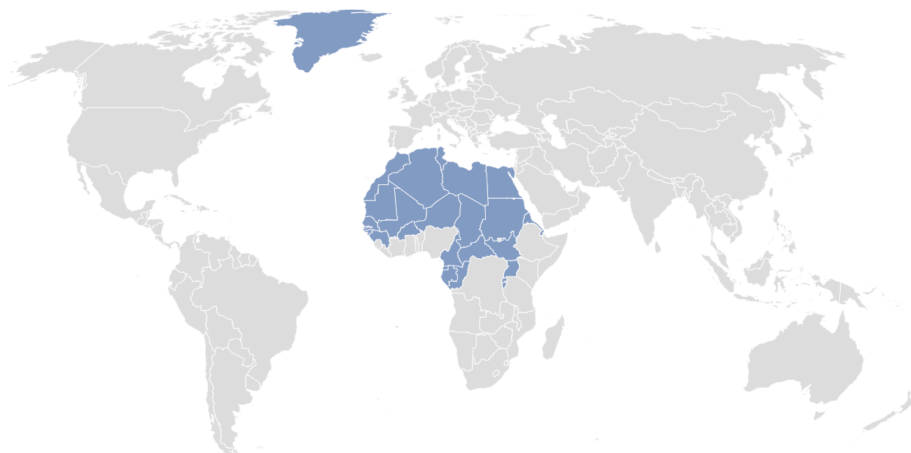

Simulation

The control variables used by BHM are supposed to extract the signal of temperature affecting growth that is independent of any noise from growth trends or country or year averages. But is it possible for a regression that includes BHM's extensive control variables to show an effect of temperature on growth when no such effect exists?

I simulated random temperatures over 50 years for 166 countries, with growth and temperatures autocorrelated by location and time, with location being one dimensional. In other words, the simulated countries are labeled by number from 1 to 166, and countries with labels numerically close to each other have growth rates and temperatures that are correlated with each other more than do countries with numbers that are further apart. The temperature of the first country is normally distributed. The second country's temperature is the first country's temperature plus another normally distributed random number. The third country adds another random variable. This is repeated for each year. In this way, countries that are closer in distance have more correlated temperatures. Simulated growth is constructed in the same way. There is no relationship between growth and temperature, but both are spatially autocorrelated.

When growth is regressed on temperature and temperature squared using 1,000 different random number draws, even with all of the controls in BHM, in 71.3 percent of the 1,000 simulations temperature and temperature squared are both statistically significant at the 5 percent level. They are both significant at the one percent level 62.7 percent of the time, and they are both significant at the 0.1 percent level 55.3 percent of the time. The coefficient on temperature is significantly positive and temperature squared significantly negative at the 5 percent level 17.6 percent of the time. Adding temporal autocorrelation increases the false statistical significance of the results. Adding a small effect of temperature on growth that is reversed in the following year produces positive coefficients on temperature and negative coefficients on temperature squared with p-values less than 5 percent in 76.1 percent of the simulations.

Using the two-step process of first regressing growth on all of the control variables without temperature and obtaining residuals, and second, regressing temperature and temperature squared on those residuals, the coefficient estimates and their statistical significance are greatly diminished, just as was the case with BHM's results.

Figure 8 shows the distribution of t-statistics in 1,000 simulations using the BHM fixed-effect model (blue line) and regressing temperature and temperature squared on the residuals of a regression of simulated growth on all of the fixed-effect controls but without temperature (red line). Using residuals reduces the statistical significance of temperature squared. (A figure for temperature is similar.) The results are similar to those obtained using BHM's actual data, where the effect of temperature on growth is strong in a model containing both temperature and fixed effects, but much weaker when residuals from a regression using fixed effects without temperature are regressed on temperature.

The results of these simulations suggest that the large number of fixed-effect control variables in BHM might not be sufficient to correct for effects, such as spatial autocorrelation, that can produce spurious results.

Figure 8. T statistics for squared temperature with fixed effects and for residuals, 1,000 simulations

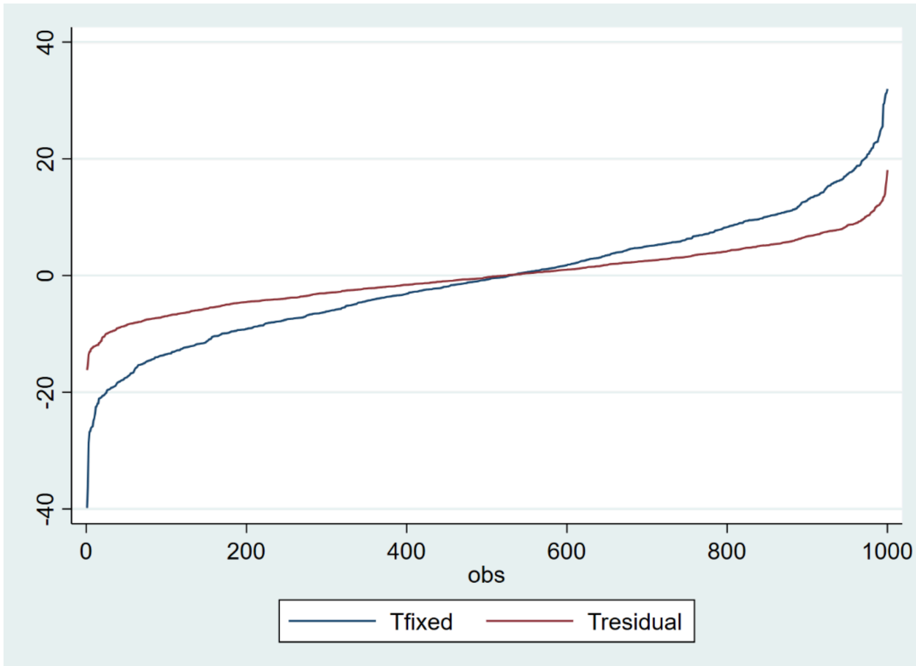

Mitigation of effect over time

Persistence of effects

In section C.2 of the online Supplementary Materials for BHM is a discussion of a test for level vs. growth effects in which lagged temperatures are included as independent variables. The tests are not discussed in the text of BHM, although their Extended Data Figure 2 contains an illustration of the results without any discussion of what the results show or their importance. BHM admit (in the Supplementary Materials) that "in models that account for lagged effects…, projections become more uncertain" (p. 26).

The idea of the test is that if a hot year reduces growth in the current year, but growth is higher than normal the next year, then there will be no persistent effect of temperature. This would be reflected in a negative coefficient on temperature in the current year and a positive coefficient on lagged temperature. If the coefficient on lagged temperature is zero while the coefficient on temperature is negative, then the loss of output in the hot year will not be recouped. Additionally, if temperatures rise over time, growth rates would fall over time. But if the signs are opposite and of similar magnitude, then there would be no long-term effect of temperature on GDP.

In their Supplementary Materials (p. 15), BHM include Supplementary Table S2, which shows the sum of the marginal effect of current and lagged temperatures for different numbers of included lags at different temperatures. The table shows estimated total effects and their standard deviations, but no p-values, t-statistics, or discussion of their statistical significance. Extended Data Figure 2 in BHM shows blue shading to illustrate the widening confidence intervals as lags are added, but there is no discussion of this in the paper. Calculating t-statistics and p-values from their Table S2 shows that adding one lag of temperature dramatically reduces the statistical significance of the effect. At 5 degrees Celsius, the p-value goes from 1 percent to 25.3 percent, and at 25 degrees it goes from 0.12 percent to 10.8 percent.

For Table S2, BHM include lagged temperature and lagged temperature squared and evaluate the marginal effects at different temperatures. Table 11 shows the results of a simpler approach. BHM model growth as a quadratic function of temperature, with higher temperatures increasing growth up to the optimal temperature and decreasing growth at temperatures above the optimum. Below the optimum, therefore, a linear model should show a positive relationship between growth and temperature, and it should show a negative linear relationship for temperatures above the optimum. A dummy variable called *cold* equal to zero for temperatures above the optimal temperature of 13.1 and another called *hot* that is the reverse of cold can be multiplied by the temperature variable and lagged temperature. Table 11 shows that the positive effect of temperature on growth for cold countries is a little more than offset by the effect in the next year, and the negative effect is almost offset in hot countries. A test of equality for each of the pairs of coefficients, however, shows no statistically significant difference.

TABLE 11. Lagged temperatures, effect varying by temperature

	Coefficient	Std. error	T-stat	p-value
cold×temp	0.004877	0.0026105	1.87	0.064
cold×lag temp	−0.006098	0.0025042	−2.44	0.016
hot×temp	−0.0081514	0.0025813	−3.16	0.002
hot×lag temp	0.0077645	0.0023425	3.31	0.001
Test cold×temp=cold×lag temp	−0.00039	0.002814	−0.14	0.891
Test hot×temp=hot×lag temp	−0.00122	0.002872	−0.43	0.671
R^2	0.3068			
Obs	6418			

If there is an effect of temperature on growth, it is clear from these results that growth bounces back from any reduction resulting from warm years, eliminating any significant net effect. BHM apparently found this result and relegated it to an online appendix and a graph that is not explained or mentioned in the paper.

Reduction in effect over time

BHM claim that there is no difference in the effect of temperature on growth over time. In other words, the effect is just as strong in the early years of their sample as in the later years. BHM write:

> We do not find that technological advances or the accumulation of wealth and experience since 1960 has fundamentally altered the relationship between productivity and temperature. Results using data from 1960–1989 and 1990–2010 are nearly identical. (BHM 2015, 237)

"Nearly identical" perhaps means that the coefficient estimate for temperature falls by only 8 percent and the coefficient for temperature squared falls by only 23 percent and both remain statistically significant, but at reduced levels of confidence. This change, however, results in an increase in the optimal temperature from 10.9 degrees to 15.2 degrees and a 48 percent decrease in the change in the growth rate from the optimal temperature to 35 degrees. Moving the cutoff date by one year, from 1989 to 1990, eliminates the statistical significance of both estimates at the 5-percent level. Because data are available for more countries later in the sample, a cutoff date of 1991 instead of 1990 does a better job of balancing the number of observations between the two subsamples of early years and late years. Table 12 shows the results of breaking the sample in 1989 and also in 1990.

TABLE 12. Regression results for early and late years[9]

	BHM	1961–1989	1990–2010	1961–1990	1991–2010
Temp	0.0127	0.0129	0.0139	0.0121	0.00819
Se	0.00379	0.00354	0.00659	0.00371	0.0069
T-stat	3.36	3.66	2.11	3.26	1.19
P-value	0.000975	0.000363	0.036	0.0014	0.237
$Temp^2$	−0.00049	−0.000595	−0.000458	−0.000579	−0.000343
Se	0.000118	0.000167	0.000183	0.000175	0.00019
T-stat	−4.11	−3.57	−2.5	−3.3	−1.81
P-value	6.12e−05	0.000488	0.0134	0.00121	0.0726
Opt. temp	13.1	10.9	15.2	10.4	11.9
Temp diff	0.235	0.347	0.18	0.349	0.183
R^2	0.286	0.305	0.489	0.317	0.502
Obs	6584	3212	3372	3353	3231

9. BHM's data on GDP per capita begin in 1960, so the first observation of GDP growth is in 1961.

The increase in the optimal temperature between the two subsamples suggests further tests of whether the estimated optimal temperature changes over time. If growth is a quadratic function of temperature, and the coefficients on temperature and temperature squared change over time, then the following equation with interactions between time and temperature could be estimated. Growth is signified by *g*, temperature by *h*, and time by *t*.

$$g = ah + bh^2 + cht + dh^2t \qquad (2)$$

Equation 2 can be rewritten as follows:

$$g = h(a + ct) + h^2(b + dt) \qquad (3)$$

The quadratic form of this function means that the optimal temperature will be minus the coefficient on h divided by twice the coefficient on h^2.

$$\text{Optimal temperature} = \frac{-(a + ct)}{2(b + dt)} \qquad (4)$$

Table 13 shows the result of this estimation. The interaction terms are statistically insignificant, but the coefficient estimates imply an increase in the optimal temperature over time, as shown in Figure 9.

TABLE 13. Interactions of time and temperature

	Coefficient	T-stat	p-value
Temperature	0.0187601	3.06	0.003
Temperature^2	−0.0007781	−2.88	0.004
Time × Temperature	−0.0002144	−1.04	0.300
Time × Temperature^2	0.0000102	1.29	0.199

Figure 10 shows the estimated optimal temperature from rolling eight-year windows of observations beginning with 1961–1968 and ending with 2003–2010. A best-fit line showing an upward slope is also pictured. An eight-year window was chosen because shorter windows resulted in regressions with few observations and noisy results because many countries are missing data from the early years of the sample. Different length windows can produce different results, but most window sizes result in an upward sloping line.

These results from dividing the sample in two suggest that any effect that existed in the early years of BHM's sample may be weaker in the later years. The results from time-and-temperature interaction and rolling window estimates are weak, but suggest that it is possible that either more countries benefit from higher temperatures over time, or that countries have become better able to cope with higher temperatures. At the very least, BHM's dismissal of the possibility that the data indicate the possibility of effect mitigation over time seems overly aggressive.

Figure 9. Optimal temperature over time

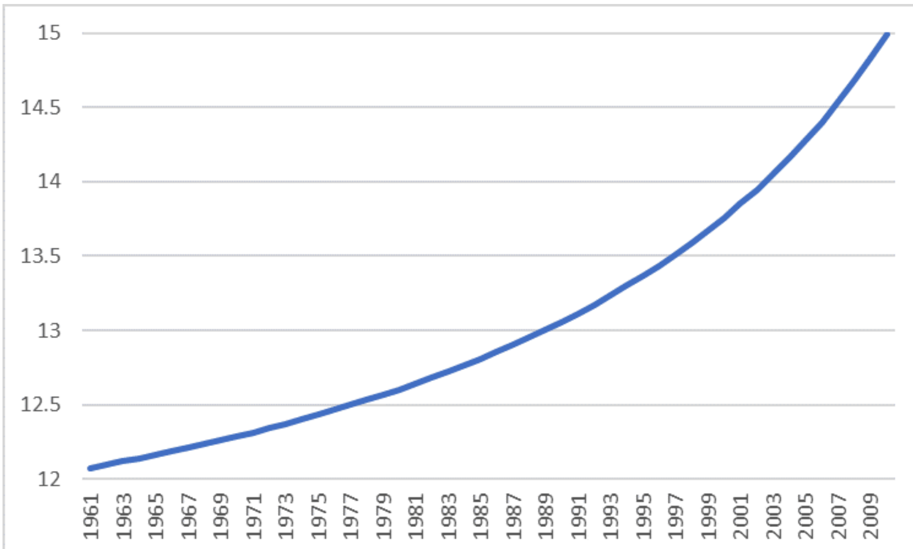

Figure 10. Optimal temperature in 8-year rolling windows of observations

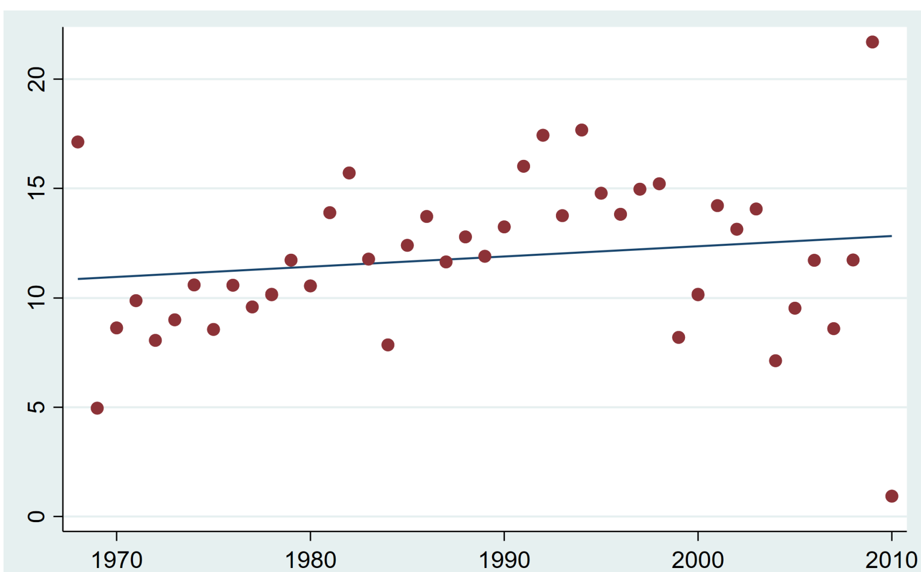

The headline result

The headline result of BHM (2015, 235) is that "unmitigated warming is expected to reshape the global economy by reducing average global incomes roughly 23% by 2100." The world average of this warming is 3.7 degrees. To calculate the 23 percent figure, BHM use the RCP8.5 SSP5 scenario to estimate average temperature for each country each year until 2100, and then plug that temperature into the quadratic function they estimated to predict growth for each country each year. GDP per capita in 2100 is calculated using these annual growth rates. Growth rates are then weighted by projected population in 2100 to calculate the overall percentage difference in GDP per capita for the entire world.

To determine the precision of the estimate, BHM perform a bootstrap procedure, where a different selection of countries is chosen 1,000 times to re-estimate the quadratic relationship between growth and temperature. The entire procedure is done each time, resulting in 1,000 different estimates of the percentage difference in world GDP per capita between a scenario with no temperature change and the temperature changes predicted by the RCP8.5 SSP5 scenario.

Twenty-three percent is the "point estimate," which means the result obtained by using the coefficient estimates of the quadratic relationship between growth and temperature from the non-bootstrapped sample. In Extended Data Table 3 in BHM, they report the median of the bootstrap estimates (21 percent) and several percentile values, including the 50th percentile value, which is the median (opposed to the mean). BHM do not report the mean estimate of the 1,000 bootstrap estimates, which is 0.11, less than half of the point estimate that they report as their headline result. While bootstrap estimates are typically used to evaluate standard errors of coefficients rather than mean estimates, such a large difference between the point estimate and the mean of the bootstrap estimates is an indication that something is amiss, perhaps the influence of countries with extreme swings in GDP per capita that are due to factors other than temperature, and also perhaps the effect of spatial and temporal autocorrelation.

BHM also do not discuss that the distribution of their estimated effect of higher temperatures is highly skewed and has extreme kurtosis, meaning that the distribution has a fat tail to the right. From Extended Data Table 3, one can calculate that there is a 29 percent chance that unmitigated warming will result in higher world GDP per capita in 2100, and a 5 percent chance that it will be more than 66 percent higher. The highest estimate in their 1,000 runs is that GDP per capita would be four times higher with warming than without. The model also predicts a one in ten chance that warming will reduce world output by more than 50 percent.

The headline result is also based on the RCP8.5 and SSP5 scenarios of future emissions. The RCP8.5 scenario is estimated to be higher than the 98th percentile of expected concentrations of greenhouse gases by the year 2100 (van Vuuren 2011).

Conclusion

BHM (2015) is a complicated paper that makes strong claims. The authors use thousands of lines of code to run regressions containing over 500 variables to test a nonlinear model of temperature and growth for 166 countries and forecast economic growth out to the year 2100. Careful analysis of their work shows that they bury inconvenient results, use misleading charts to confuse readers, and fail to report obvious robustness checks. Simulations suggest that the statistical significance of their results is inflated.

BHM's results are unreliable and their data do not support predictions of large economic losses due to rising temperatures. BHM (2015), along with Dell et al. (2012), Colacito et al. (2019), and Kiley (2021; forthcoming 2024) attempted to show that warming will reduce the rate of growth of GDP per capita. I have shown that all of these papers are seriously flawed (Barker 2022; 2023a; 2023b), and there is no credible evidence suggesting that warming will reduce the rate of growth of GDP per capita.

Continued economic growth at levels similar to what the world has experienced in recent years would increase the level of future economic activity by far more than Nordhaus' (2018) estimate of the effect of warming on future world GDP. If warming does not affect the rate of economic growth, then the world is likely to be much richer in the future, with or without warming temperatures.

Code

Code used in this research is 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. 2023a. Temperature and Economic Growth: Comment on Kiley. *Econ Journal Watch* 20(1): 69–84. **Link**

Barker, David. 2023b. Temperature Shocks and Economic Growth: Comment on Dell, Jones, and Olken. *Econ Journal Watch* 20(2): 234–253. **Link**

Booth, Shawn, and Danielle Knip. 2014. The Catch of Living Marine Resources Around Greenland from 1950 to 2010. In *Fisheries Catch Reconstructions: Islands, Part IV*, eds. Kyrstn Zylich, Dirk Zeller, Melanie Ang, and Daniel Pauly, 55–71. *Fisheries Centre Research Reports* (Fisheries Centre, University of British Columbia, Vancouver) 22(2). **Link**

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (BHM). 2015. Global Non-Linear Effect of Temperature on Economic Production. *Nature* 527(7577): 235–239.

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. 2012. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics* 4(3): 66–95.

Forthofer, Ronald, Eun Sul Lee, and Mike Hernandez. 2007. *Biostatistics: A Guide to Design, Analysis and Discovery*, 2nd ed. Burlington, Mass.: Elsevier Academic Press.

Granger, Clive. 1974. Spurious Regressions in Econometrics. *Journal of Econometrics* 2(2): 111–120.

Hadri, K. 2000. Testing for Stationarity in Heterogeneous Panel Data. *Econometrics Journal* 3: 148–161.

Hausfather, Zeke, and Glen Peters. 2020. Emissions—the 'Business as Usual' Story Is Misleading. *Nature* 577(7792): 618–620.

Hsiang, Solomon, Simon Greenhill, Jeremy Martinich, Monica Grasso, Rudy M. Schuster, Lint Barrage, Delavane B. Diaz, Harrison Hong, Carolyn Kousky, Toan Phan, Marcus C. Sarofim, Wolfram Schlenker, Benjamin Simon, and Stacy E. Sneeringer. 2023. Economics. In *Fifth National Climate Assessment*, eds. Allison R. Crimmins, Christopher W. Avery, David R. Easterling, Kenneth E. Kunkel, Brooke C. Stewart, and Thomas K. Maycock, ch. 19. Washington, D.C.: U.S. Global Change Research Program. **Link**

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

Kiley, Michael T. 2024 (forthcoming). Growth at Risk From Climate Change. *Economic Inquiry*. **Link**

Levin, Andrew, Chien-Fu Lin, and Chia-Shang James Chu. 2002. Unit Root Tests in Panel Data: Asymptotic and Finite-Sample Properties. *Journal of Econometrics* 108: 1–24.

Müller, Ulrich K., and Mark W. Watson. 2023. Spatial Unit Roots. Working paper, March. **Link**

Newell, Richard G., Brian C. Prest, and Steven E. Sexton. 2021. The GDP-Temperature Relationship: Implications for Climate Change Damages. *Journal of Environmental Economics and Management* 108: 102445.

Nordhaus, William. 2018. Projections and Uncertainties About Climate Change in an Era of Minimal Climate Policies. *American Economic Journal: Economic Policy* 10(3): 333–360.

Pomeranz, Kenneth. 2001. *The Great Divergence: China, Europe, and the Making of the Modern World Economy*. Princeton, N.J.: Princeton University Press.

Rosen, Richard. 2019. Temperature Impact on GDP Growth Is Overestimated. *PNAS* 116(33): 16170. **Link**

Thomassen, Bjørn. 2003. The Black Angel Lead-Zinc Mine at Maarmorilik in West Greenland. *Geology and Ore* (Nationale Geologiske Undersøgelser for Danmark og Grønland, Copenhagen), no. 2. **Link**

Tol, Richard. 2019. A Social Cost of Carbon for (Almost) Every Country. *Energy Economics* 83: 555–566. **Link**

van Vuuren, Detlef P., Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C. Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque, Toshihiko Masui, Malte Meinshausen, Nebojsa Nakicenovic, Steven J. Smith, and Steven K. Rose. 2011. The Representative Concentration Pathways: An Overview. *Climatic Change* 109: 5. **Link**

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass.: MIT Press.

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 March 2024 issue