

Temperature and Income:

Reconciling New Cross-Sectional and Panel Estimates

Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken¹

It has long been observed that hot countries tend to be poor. A correlation between heat and poverty was noted as early as Montesquieu (1750) and Huntington (1915), and it has been repeatedly demonstrated in contemporary data (e.g. Nordhaus 2006). Looking at a cross-section of the world in the year 2000, national income per-capita falls 8.5% per degree Celsius rise in temperature (see Table 2 below). In fact, temperature alone can explain 23% of the variation in cross-country income today.

Despite the strength of this correlation, substantial debate continues over whether climatic factors can explain contemporary economic activity, or whether other correlated variables, such as a country's institutions or trade policy, drive prosperity in contemporary times, leaving no important role for geography (see, e.g., Sachs 2003; Acemoglu, Johnson and Robinson 2002; Rodrik, Subramanian and Trebbi 2004). Given the small number of countries in the world, and the many ways in which they vary, conclusively answering these questions using cross-sectional, cross-country regressions is challenging.

This paper offers two new insights into the climate-income relationship. First, we provide novel cross-sectional evidence by considering the temperature-income relationship not just using cross-country data but also using *sub-national* data at the municipal level for 12 countries in the Americas. Remarkably, we find that a negative relationship between income and temperature exists when looking within countries, and even when looking within states within countries. The within-country cross-sectional relationship is substantially weaker than the cross-

¹ Dell: Massachusetts Institute of Technology (mdell@mit.edu); Jones: Northwestern University and NBER (bjones@kellogg.northwestern.edu); Olken: Massachusetts Institute of Technology and NBER (bolken@mit.edu).

country correlation, but it remains statistically significant and of an economically important magnitude, with a 1 degree C rise in temperature associated with a 1.2-1.9% decline in municipal per-capita income. The fact that the cross-sectional relationship holds within countries, as well as between countries, suggests that omitted country characteristics are not wholly driving the cross-sectional relationship between temperature and income.²

Second, we provide a theoretical framework for reconciling these strong cross-sectional effects of temperature with the even stronger short-run effects of temperature shown in panel models. In related work (Dell, Jones, and Olken 2008; henceforth DJO), we build a climate and income panel at the country-year level and examine what happens to the national growth path when countries have unusually hot or cold years. The primary finding in DJO (2008) is that, in poor countries over the 1950-2003 period, a 1 degree Celsius rise in temperature in a given year reduced economic growth in that year by 1.1 percentage points. Moreover, the estimated temperature effects over 10 or 15-year time horizons are similar to the annual panel estimate, suggesting that these effects represent changes to growth rates, rather than level effects on income. These temperature effects on growth are sufficiently large that, in the absence of offsetting forces, they would quickly produce a much steeper relationship than we actually see between temperature and income: if an extra 1 degree C reduces growth by 1.1 percentage points, then it would take only 8 years of sustained temperature differences to explain the overall cross-sectional relationship between temperature and income observed in the world today.

To reconcile the cross-sectional and panel results, we consider a simple theory that emphasizes two forces, adaptation and convergence, and shows how the causative estimate of the temperature-growth effect in DJO (2008) can be reconciled with the long-run temperature-

² Albouy (2008) similarly finds a negative correlation between temperature and firm productivity within the United States.

income findings. The estimates suggest that, in the cross-country context, adaptation offsets about half of the negative effects of higher temperatures.

I. Cross-Sectional Evidence at the Sub-National Level

A. Data

To examine the temperature-income relationship at the sub-national level, we use municipal-level labor income data for 12 countries in the Western Hemisphere, as constructed from household surveys by Acemoglu and Dell (forthcoming).³ To make the data comparable across municipalities and countries, Acemoglu and Dell account for regional price dispersion and adjust each country's wage data so that it averages to GDP per worker in constant international dollars, taken from the 2003 Penn World Tables.

We use all countries in the Acemoglu and Dell dataset where the labor income data can be geo-referenced to a municipality, and merge this data with climate and geography data. The list of countries in the dataset, along with summary statistics, are shown in Table 1. Climate data are at 30 arc second resolution (approximately 1 km) and averaged over the 1950-2000 period, as calculated by Hijmans et al. (2005). Country-level climate variables aggregate the municipal-level variables, weighting by 2000 municipal population. Details can be found in the appendix.

B. Results

Using this data, we estimate the cross-sectional relationship between climate variables – mean temperature and mean precipitation levels – and log income, i.e.

$$(1) \quad \text{LOGY}_{rm} = \alpha_r + \beta_1 \text{TEMP}_{rm} + \beta_2 \text{PRECIP}_{rm} + X_{rm}' \gamma + \varepsilon_{rm}$$

where LOGY is the mean log labor income, r represents a region, m represents a municipality, and X represents other geographic variables. We estimate equation (1) using OLS. Standard errors are calculated clustering observations by state (shown in parentheses) and, alternatively,

³ Acemoglu and Dell focus on labor income since the errors in reporting are less severe than for total income.

using corrections for spatial correlation (Conley 1999) (shown in brackets)⁴.

The results from estimating equation (1) are presented in Table 2. As a benchmark, we begin in column (1) of Table 2 with a cross-country regression for the whole world. Specifically, we use all 134 countries in the DJO (2008) sample, and calculate *LOGY* as log gdp per capita from the Penn World Tables. This regression shows that each additional 1 degree C is associated with a statistically significant reduction of 8.5 percentage points of per-capita GDP. In Column (2), we limit the sample to the 12 countries in our labor income dataset shown in Table 1. The point estimate for the effect of temperature remains virtually unchanged at 8.9 percentage points of per-capita GDP per degree C, although with only 12 data points the standard errors increase substantially and the result is no longer statistically significant.

In column (3), we switch to our labor income data set. Column (3) examines the same set of countries as column (2) but at the municipality level. We regress mean municipal labor income on municipal temperature and precipitation, and add additional geographic controls for elevation, slope, and the distance from the municipality to the sea. The temperature coefficient is -0.085 log points, which is virtually identical to the coefficient using country-level data, and is now statistically significant with standard errors either clustered by state or corrected for spatial correlation. Remarkably, the five explanatory variables in this regression – temperature, precipitation, elevation, slope, and distance to the sea – explain 61 percent of the variation in municipal income across these 12 countries.

Columns (4) and (5) examine the relationship between temperature and income within countries. In column (4), we add country fixed effects. The point estimate falls substantially to -

⁴ The Conley covariance matrix is a weighted average of spatial autocovariances, where the weights are the product of Bartlett kernels in two dimensions (North/South and East/West). They start at one and decline linearly to zero when a pre-specified cut point is reached. We choose the cutoff in both dimensions to be one degree (approximately 100 kilometers); choosing other cut points produces qualitatively similar results.

0.012 but remains statistically significant; that is, a 1 degree C increase in temperature is associated with a 1.2% decline in labor income. Remarkably, when we add state fixed effects in column (5), so that we are using only variation in temperatures within individual states, the point estimate on temperature remains very similar to the estimate with country fixed effects (at -0.019) and is significant when using spatial standard errors.⁵ These results confirm that the cross-sectional relationship between temperature and income holds within countries, as well as across countries, though the relationship is substantially smaller in magnitude within countries than across countries.

II. Theory: Adaptation and Convergence

In this section we consider means of reconciling the long-run cross-sectional relationships documented in Section I with the short-run growth effects of temperature estimated in DJO (2008). As discussed above, DJO (2008) use panel data to show that a poor country's growth in a given year is 1.1 percentage points lower when its temperature is 1 degree Celsius higher that year. Moreover, as discussed in that paper, the persistent effect of temperature shocks suggests that temperature affects the growth rate, not simply the level of income, at least over 10 to 15 year time horizons.

To reconcile these large growth effects of temperature with the more modest (though still substantial) long-run cross-sectional relationship between temperature and income, we consider two mechanisms: convergence and adaptation. First, convergence forces may pull lagging countries and regions toward the frontier. Convergence effects offset temperature effects, so that convergence limits the cross-sectional income differences that can be sustained. If rates of convergence are larger within countries than across them, then the long-run effect of climate will

⁵ A state is defined as the first administrative level political unit below the central government. In Brazil, Mexico, and the United States, the term for these political units is state, whereas in other countries they are alternatively called departments or provinces. See the appendix for maps of the state boundaries.

be more muted within countries than across them. While data on within-country convergence for much of the world is limited, faster within than across country convergence is consistent with the smaller income variance within countries and is natural given greater opportunities for migration, public good provision, transfers, and idea exchange within countries.⁶

Second, over longer periods, regions may adapt to their climate. The panel growth estimates reflect responses to climate shocks. To the extent that individuals adjust their behavior to permanent temperature changes, e.g. by switching to more appropriate crops, industries, and technologies and by migrating away from difficult environments altogether, the short-run estimates may be larger than the longer-run response.

To fix ideas, imagine that growth in per-capita income proceeds as

$$(2) \quad \frac{\log y_i(t)}{dt} = g + \gamma(T_i(t) - \bar{T}_i) + (\gamma + \rho)\bar{T}_i + \varphi(\log y_*(t) - \log y_i(t)) \text{ for } t \geq 0$$

where $\log y_i(t)$ is the log per-capita income in geographic area i at time t , $T_i(t)$ is the temperature in area i at time t , \bar{T}_i is the average temperature level in area i , and $\log y_*(t)$ is the relevant frontier level of income to which the area converges.⁷ The parameter γ captures the causative short-run effect of temperature shocks on growth, as would be identified in a panel specification such as DJO (2008). The parameter ρ captures the degree of adaptation over the long-run to average temperature levels, potentially offsetting the short-run temperature effects. The parameter $\varphi \in (0,1)$ captures the rate of convergence. We further assume that all countries start, in antiquity at time zero, with the same level of per-capita income, $\log y_i(0) = c$ for all i .

Note that since equation (2) applies to all countries, including country $*$, $E[\log y_*(t)] = c +$

⁶ Note that within-country studies do not show faster rates of convergence, though estimates vary substantially depending on methodology (e.g. see Barro and Sala-i-Martin (1995) versus Higgins, Levy, and Young (2006)).

⁷ Note that while (2) is a very simple description of growth, it departs from the usual neoclassical assumption, where all countries have the same growth rate in total factor productivity, and convergence drives countries not towards a distribution of income, but to a common income level. We model growth in (2) to accommodate the empirical finding of DJO (2008), where temperature affects the growth rate (e.g. the ability to invent or absorb new ideas).

$$(g + (\gamma + \rho)\bar{T}_*)t.$$

Integrating the differential equation (2) with the initial condition and taking expectations, we have

$$(3) \quad \mathbf{E}[\log y_i(t)] = \mathbf{E}[\log y_*(t)] + \frac{\gamma + \rho}{\varphi} (\bar{T}_i - T_*) (1 - e^{-\varphi t})$$

(This derivation is shown formally in the appendix.) Therefore, in the long run as $t \rightarrow \infty$, the cross-sectional relationship between income and temperature is

$$(4) \quad \frac{d\mathbf{E}[\log y_i]}{d\bar{T}_i} = \frac{\gamma + \rho}{\varphi}$$

Equation (4) is an equation with four unknowns, and we have estimates for three of them. The left-hand side of (4) is the cross-sectional regression parameter in the regression of income on temperature; i.e., $\beta = -.085$ in a cross-country context and $\beta = -.012$ in a within-country context (see Table 2). As discussed above, the short-run growth coefficient is approximately $\gamma = -.011$ (DJO 2008). The convergence parameter, much analyzed in the growth literature, is typically estimated in the cross-country context in the range $\varphi \in [.02, .10]$ (Barro and Sala-i-Martin 1995; Caselli et al. 1996).

A. The Convergence Mechanism

We first consider turning off the adaptation channel (setting $\rho = 0$ in (4)) to examine the implications of convergence alone. In this setting, reconciling the short-run and long-run temperature effects is achieved when $\varphi = \gamma/\beta$. In a cross-country context, this requires $\varphi = .129$ [i.e. $-.011/-.085$], which appears somewhat high given estimates in the literature. At a within-country level, we have no panel estimate of the short-run growth effect γ . If one applies the cross-country estimate of γ , then we require $\varphi = .917$ [i.e. $-.011/-.012$]. While it is reasonable that convergence rates might be substantially higher in a within-country context,

this estimate appears extremely high.⁸ These calculations suggest that adaptation is likely to be important in reconciling the data.

B. The Adaptation Mechanism

Over the long run, areas may adapt to difficult geographic conditions. Technologies, skills, and physical capital can all be tailored to a given climatic regime. Moreover, population can react, either through fertility, death rates, or migration, thus altering the local per-capita intensity of the factors of production.

We now relax the strong assumption of no adaptation ($\rho = 0$), and instead estimate ρ using our findings for β and γ , and a chosen convergence rate, φ . Rearranging (4) shows that $\rho = \beta\varphi - \gamma$. In the cross-country context, taking a middle-of-the-road convergence rate of $\varphi = .06$ yields an estimate of $\rho = .0059$. This suggests that 54% of the short-run effect is offset in the long-run, so that the long-run growth rate effect of being 1 degree C warmer is $\gamma + \rho = -.0051$, or half of one percentage point per annum.

In the within-country context, there is more uncertainty, both because the short-run within-country growth effect has not been estimated in panel data and because the convergence rate may be greater. If we apply the country-level panel estimate of $\gamma = -.011$ and take the upper-bound cross-country convergence estimate of $\varphi = .10$ internally, we find $\rho = .0098$, so that 89% of the short-run growth effect is offset within countries. Thus if the short-run growth estimate was the same within countries as between countries, there would be an even larger role for adaptation within countries than between countries.⁹

C. The Omitted Variable Interpretation

⁸ For example, in developed countries (US, Japan, Europe) Barro and Sala-I-Martin estimate within-country convergence coefficients of approximately 0.02 – 0.03.

⁹ For example, prices can offset productivity shocks, such as those due to temperature. To the extent that markets are more integrated within countries than across countries, the price adaptation mechanism may offset the effects of temperature differences more completely within countries.

A typical objection to the cross-country temperature-income relationship is that it may be driven by omitted variables. However, the findings of DJO (2008) suggest a substantial, causative effect of temperature on growth for poor countries, and the above analysis shows how these growth effects can be reconciled with the cross-sectional evidence. One may then ask: is there still no role for omitted variables in the cross-section? In fact, the same framework above allows one to assess the role of such omitted variables; mathematically, omitted variables are analogous to the ρ adaptation parameter. To see this, we can write the growth process as

$$(5) \quad \frac{\log y_i(t)}{dt} = g + \gamma T_i(t) + \boldsymbol{\theta} \mathbf{Z}_i + \varphi (\log y_*(t) - \log y_i(t)) \text{ for } t \geq 0$$

where \mathbf{Z}_i is a vector of omitted variables that influence growth and also happen to be correlated with average temperature, \bar{T}_i .

However, for omitted variables to reconcile the cross-section and panel estimates without any role for adaptation, the omitted variables would need to have strongly *positive* effects on growth in high temperature countries. That is, very hot countries (such as the Saharan countries Chad, Mauritania, and Niger) would need to have characteristics that are making them grow faster than they otherwise would. Cases where this omitted variable story seems plausible are the Persian Gulf states, which are extremely hot but grow through large oil resources. However, even if we drop these states the world cross-sectional coefficient β rises only to -0.097, and the implied adaptation coefficient ρ is still 0.0036, so omitted variables would still need to be very positive in other hot countries to reconcile the data without some adaptation being present.¹⁰

From the perspective of future climate change, the omitted variable interpretation of the cross-section suggests worse effects of future warming than the adaptation interpretation of the

¹⁰ It is also possible that omitted variables are more substantial in a cross-country setting than a within country-setting. This could help reconcile the milder income-temperature relationship within countries with the sharper relationship across countries without relying on different adaptation or convergence rates.

cross-section. With omitted variables, the long-run effect of warming on the income distribution is γ/φ , which is substantially more negative than the long-run effect under adaptation, which is $(\gamma + \rho)/\varphi$. DJO (2008) emphasize an adaptation view and thus provide a lower-bound type of analysis of the future impacts of climate change.

III. Conclusion

This paper provides new evidence on the relationship between temperature and income. Using sub-national data from 12 countries in the Americas, we show that the negative cross-sectional relationship between temperature and income exists within countries, as well as across countries. We then provide a theoretical framework for reconciling the substantial, negative association between temperature and income in cross-section with the even stronger short-run effects of temperature shown in panel models. The theoretical framework suggests that half of the negative short-term effects of temperature are offset in the long run through adaptation.

References

- Acemoglu, Daron and Melissa Dell. "Beyond the Neoclassical Growth Model: Technology, Institutions and Within-Country Differences," *American Economic Journal: Macroeconomics*, forthcoming.
- Acemoglu, Daron, Simon Johnson, and James Robinson. "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," *Quarterly Journal of Economics* 117, 1231-1294, 2000.
- Albouy, David. "What Are Cities Worth: Land Rents, Local Productivity, and the Value of Amenities," *mimeo*, University of Michigan Department of Economics.
- Barro, Robert and Xavier Sala-i-Martin. *Economic Growth*. McGraw-Hill, 1995.
- Caselli, Francisco, Gerardo Esquivel, and Fernando Lefort . "A New Look at Cross-Country

- Growth Empirics,” *Journal of Economic Growth* 1(3), 363-389, 1996.
- Conley, Timothy. “GMM Estimation with Cross Sectional Dependence.” *Journal of Econometrics*, 92(1), 1-45, 1999.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. “Climate Shocks and Economic Growth: Evidence from the Last Half Century,” NBER Working Paper 14132, 2008.
- Higgins, Matthew, Daniel Levy, and Andrew Young. “Growth and Convergence across the United States: Evidence from County-Level Data,” *The Review of Economics and Statistics* 88(4), 671-681, 2006.
- Hijmans, R.J. and Cameron, S.E. and Parra, J.L. and Jones, P.G. and Jarvis, A. and others. “Very high resolution interpolated climate surfaces for global land areas,” *International Journal of Climatology* 25(15), 1965-1978, 2005.
- Huntington, Ellsworth. *Civilization and Climate*, New Haven, CT: Yale University Press, 1915.
- Montesquieu, Charles de. *The Spirit of Laws*, 1750.
- Nordhaus, William. “Geography and Macroeconomics: New Data and Findings,” *Proceedings of the National Academy of Science* 103, 3510-3517, 2006.
- Rodrik, Dani, Arvind Subramanian, and Francesco Trebbi, “Institutions Rule: The Primacy of Institutions Over Geography and Integration in Economic Development,” *Journal of Economic Growth* 9, 131-165.
- Sachs, Jeffrey D. “Institutions Don’t Rule: Direct Effects of Geography on Per-Capita Income,” National Bureau of Economic Research working paper No. 9490, 2003.

Table 1: Data Summary

Country	Source	Year	Obs.	No.	Income		Temp.	
				Munic.	Mean	SD	Mean	SD
Bolivia	Encuesta de Hogares	2002	8,166	106	7,256	2,486	14	6.5
Brazil	Population Census	2000	3,517,842	1,517	15,462	6,525	20	3
El Salvador	Encuesta de Propositos Multiples	2006	22,937	64	10,955	3,227	23	1.3
Guatemala	Encuesta Nacional de Condiciones de Vida	2000	11,440	226	10,190	5,683	18	4.5
Honduras	Encuesta de Condiciones de Vida	2004	13,236	98	6,121	3,300	22	2.5
Mexico	Population Census	2000	2,735,333	2,442	18,628	9,103	16	4
Nicaragua	Encuesta Nacional de Hogares sobre Medicion de Nivel de Vida	2005	12,847	136	8,615	4,477	25	1.9
Panama	Population Census	2000	94,928	30	19,499	7,522	26	1.3
Paraguay	Encuesta Integrada de Hogares	2001	6,867	175	12,237	5,964	21	0.8
Peru	Encuesta Nacional de Hogares	2001	22,207	609	11,082	7,363	18	6.4
U.S.	Population Census	2000	7,401,157	2,071	67,865	19,143	12	4.7
Venezuela	Population Census	2001	677,524	219	14,848	3,141	29	4.0

Table 2: Temperature and Income

	Dependent variable is:				
	Log per capita GDP (PWT)		Log labor income		
	(1)	(2)	(3)	(4)	(5)
Temperature	-0.085 (0.017) [0.017]	-0.089 (0.083) [0.072]	-0.085 (0.007) [0.004]	-0.012 (0.005) [0.004]	-0.019 (0.015) [0.009]
Precipitation	0.000 (0.016) [0.015]	0.019 (0.041) [0.047]	-0.003 (0.001) [0.001]	0.000 (0.001) [0.001]	0.002 (0.001) [0.001]
Elevation, slope, coast	no	no	yes	yes	yes
Country F.E.	no	no	no	yes	yes
State F.E.	no	no	no	no	yes
R-squared	0.23	0.21	0.61	0.82	0.88
Number of clusters			260	260	260
Number of observations	134	12	7684	7684	7684

Notes: The dependent variable in columns (1) – (2) is the log of GDP per capita in 2000 (Heston et al., 2006) and in columns (3) through (5) is the log of mean municipality labor income (Acemoglu and Dell, 2009). Columns (3) through (5) are weighted by the number of observations in the municipality. Robust standard errors, clustered by state in columns (3) through (5), are reported in parentheses, and Conley standard errors are reported in brackets.