

NBER WORKING PAPER SERIES

CLIMATIC CONSTRAINTS ON AGGREGATE ECONOMIC OUTPUT

Marshall Burke
Vincent Tanutama

Working Paper 25779
<http://www.nber.org/papers/w25779>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2019

We thank Solomon Hsiang and Miyuki Hino for helpful comments. All errors are our own. We thank Stanford University for research support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Marshall Burke and Vincent Tanutama. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Climatic Constraints on Aggregate Economic Output
Marshall Burke and Vincent Tanutama
NBER Working Paper No. 25779
April 2019
JEL No. O13,Q5,Q54

ABSTRACT

Efficient responses to climate change require accurate estimates of both aggregate damages and where and to whom they occur. While specific case studies and simulations have suggested that climate change disproportionately affects the poor, large-scale direct evidence of the magnitude and origins of this disparity is lacking. Similarly, evidence on aggregate damages, which is a central input into the evaluation of mitigation policy, often relies on country-level data whose accuracy has been questioned. Here we assemble longitudinal data on economic output from over 11,000 districts across 37 countries, including previously nondigitized sources in multiple languages, to assess both the aggregate and distributional impacts of warming temperatures. We find that local-level growth in aggregate output responds non-linearly to temperature across all regions, with output peaking at cooler temperatures ($<10^{\circ}\text{C}$) than estimated in earlier country analyses and declining steeply thereafter. Long difference estimates of the impact of longer-term (decadal) trends in temperature on income are larger than estimates from an annual panel model, providing additional evidence for growth effects. Impacts of a given temperature exposure do not vary meaningfully between rich and poor regions, but exposure to damaging temperatures is much more common in poor regions. These results indicate that additional warming will exacerbate inequality, particularly across countries, and that economic development alone will be unlikely to reduce damages, as commonly hypothesized. We estimate that since 2000, warming has already cost both the US and the EU at least \$4 trillion in lost output, and tropical countries are $>5\%$ poorer than they would have been without this warming.

Marshall Burke
Department of Earth System Science
Stanford University
Stanford, CA 94305
and NBER
mburke@stanford.edu

Vincent Tanutama
Center on Food Security and the Environment
Stanford University
473 Via Ortega
Stanford, CA 94305
tanutama@stanford.edu

Introduction

Global policy debates about climate change mitigation center on the aggregate economic costs and benefits of particular mitigation action, as well as to whom these costs and benefits accrue. Quantifying the magnitude and distribution of these costs and benefits is now a focus of a large literature, much of which has sought to empirically estimate how various societal outcomes might be affected by a changing climate. Recent efforts combine measurements of the climate system with statistical methods designed to isolate the impacts of observed changes in climate from other factors that might also affect social or economic outcomes of interest. These studies provide strong causal evidence that key productive elements of economies, from agriculture to labor productivity to human health, respond non-linearly to changes in climate.¹

Inferring both aggregate economic costs and their distribution from these multiple foundational damages, however, is often challenging. For instance, a complete accounting of aggregate costs requires both that all relevant economic sectors can be enumerated and evaluated, and also that any interaction in impacts between sectors is accurately quantified. Meeting these requirements is demanding even in data-rich environments,² but is likely implausible in much of the world where key sectoral data are lacking and interactions are difficult to constrain.

An alternate approach uses country-level data on economic aggregates (e.g. gross domestic product) to study total damages.^{3,4} While this approach has the advantage of directly observing the aggregated net effect of myriad interacting sectors that have simultaneously responded to a change in climate, it faces its own set of important challenges. First, national accounts data on economic output are thought to be both noisy and potentially biased for many countries in the world, with alternative data sources – and even different vintages of the same data source – often showing low levels of agreement.^{5,6} Second, even if data are unbiased, measuring the response of economic output to temperature at the country level is challenging because many different producing units within a country could be exposed to opposing temperature anomalies within a given year. In fact, in large economies such as the EU countries or the US, temperature anomalies within a given year are often strongly negatively correlated across the largest producing regions (e.g. US West coast versus East coast; Fig S1a-b). Averaging these opposing anomalies to construct a country-wide annual temperature estimates reduces the temperature signal and could greatly amplify noise in a country-level analysis (Fig S1c-e; Methods).

Third, given the large variation in income within countries, a country-level analysis allows only partial insight into the distributional effects of warming. In particular, it remains difficult to understand whether any apparent differences in the effect of warming between rich and poor regions are due to differences in underlying exposures to harmful temperatures, or to differences in responsiveness at a given level of temperature exposure.⁷ Distinguishing these effects is at the core of understanding how both past and future economic development could shape the size and distribution of climate damages. Finally, given limited country-level data, it has remained difficult to understand whether a given increase in temperature affects the level or the growth rate of output, a difference with again fundamental implications for the overall magnitude

of future damages.^{3,4,8}

Here we assemble a novel district-level panel dataset on climate and GDP across 37 countries and multiple decades, using economic data from public databases and public non-digitized archives (SI Section 1) and matching them to average temperature and precipitation in each district-year using three separate climate datasets. This large dataset on over 11,000 districts allows us to use multiple sources of variation to quantify the persistent effects of warming on economic output, and to assess how these effects differ between and within countries. Our results provide a new foundation for assessment of the local and global economic impact of past and future warming.

We collect GDP data at the most local level available in existing archives, which is typically at the second-level administrative division of each country. We denote these units “districts” and their parent level-1 administrative level units as “states”. Data are harmonized to constant boundaries over the sample period and merged with estimates of temperature and precipitation in each district-year from three separate climate datasets,^{9–11} yielding 154 thousand district-year observations across Brazil, China, Colombia, India, Indonesia, United States, and the EU countries, with temporal coverage varying by country (Fig 1a). To assess data quality in the economic data, we follow earlier country-level work¹² and compare whether growth rates in subnational GDP correspond to growth rates in nighttime lights in the same location, the latter an independent measure of economic activity. We recover average nightlights-GDP relationships in our district data almost identical to these previous country-level estimates (Fig S2), providing some confidence in quality of our district-level economic data. The range of temperature exposures and income levels in our dataset broadly reflect the global country-level distribution (Fig 1b), and our assembled regions represent most of the largest economies globally as well as many of those judged to have the highest national statistical capacities (Fig 1c).

To estimate the causal impact of variation in temperature on economic output, we follow a large literature^{1,3,4} and use fixed effects regression to isolate interannual temperature variation from all other time-invariant factors or common time-varying factors that could be correlated with both average temperature and economic output. In particular, we estimate:

$$y_{istd} = f(T_{ist}) + \lambda_1 P_{ist} + \lambda_2 P_{ist}^2 + \rho_d + \alpha_i + \eta_{st} + \varepsilon_{istd} \quad (1)$$

where i indicates districts, s indicates the admin-1 region each district falls into, t indicates years, and d indicates data source for GDP data. The outcome y_{istd} is the first difference in log per capita GDP (i.e. the growth rate). $f(T_{ist})$ is a non-linear function of temperature; our main specification is quadratic, i.e. $f(T_{ist}) = \beta_1 T_{ist} + \beta_2 T_{ist}^2$, and we test robustness to more flexible polynomials and splines (Fig 2). We control for all time invariant differences between districts with district fixed effects α_i ; these account, for instance, for any fixed geographic, social, economic, or institutional differences between districts. We control for common time-varying factors with a set of state-by-year fixed effects η_{st} , which control non-parametrically for any trending variables at the state level. ϕ_d is a data-source fixed effect, which additionally accounts for

any changes in dataset within country over time that is not already picked up by the state-year FE. Given these fixed effects, our estimates of β_1 and β_2 come from comparing whether economic growth in a given district (as measured by a specific dataset) was above or below normal in a year that was warmer than average for that district, after accounting for any shocks to either growth or temperature common to districts in the surrounding state. Standard errors are clustered at the district level.

Results

We find that district level economic production is concave in temperature exposure, with a negative slope throughout nearly all the observed temperature distribution and increasingly steep at warmer temperatures (Fig 2). Results are robust to choice of temperature dataset (Fig 2b, Table S3), to choice of fixed effects or inclusion of time trends (Fig 2b and Fig S3a), to using more flexible functional forms for temperature (Fig 2b), to dropping one country at a time (Fig S3b), to alternate approaches of excluding or trimming large growth outliers (Fig S3c-e), and to accounting for spatial lags of the dependent variable (Fig S4). Responses when models are estimated separately in each country are largely consistent with the pooled global response (Fig S5), with disagreements largest in countries where the climate datasets show the poorest correlation with each other (Fig S6), suggesting these differences could be driven by climate data quality. A model that uses the distribution of daily temperatures instead of annual average temperature to measure temperature exposure recovers a response consistent in shape with our main response, although point estimates on narrow temperature bins are somewhat noisier (Fig S7). Our results are broadly consistent with subnational estimates from US districts¹³ and state-level estimates from a global sample.¹⁴

While estimated short-run effects of additional warming are almost identical to earlier country-level estimates for warmer regions,⁴ district level estimates do not provide strong evidence of positive effects in cooler regions as was suggested in country data. Applying our estimates globally suggests that nearly all regions, including most wealthy regions, would experience reduced economic growth for any additional warming above their current 2001-2015 average temperature (Fig 2c).

Conventional wisdom and some evidence^{7,15-18} suggests that poorer individuals and regions could be more negatively affected by warming temperatures, but existing country-level analyses disagree on the source of this disparity.^{3,4} In particular, disparities in the impact of warming could occur because the poor are more vulnerable to a given amount of warming, perhaps because they more often work in climate-sensitive industries such as agriculture, or because they have fewer resources (e.g. savings) to buffer against negative shocks. Alternatively, rich and poor could have the same underlying vulnerability to warming described by a common non-linear response function, but could have different baseline temperature exposures that result in different marginal damages when temperatures rise (Fig 3a-b). In the former case, economic development on its own would be expected to reduce future climate damages in poor countries, while in the latter case it would not. Thus distinguishing exposure from vulnerability as sources of disparity in impacts is central to understanding the magnitude and distribution of climate impacts.⁷

We find that global disparities in the aggregate economic impacts of warming are more likely to be driven by

differences in temperature exposures rather than in underlying vulnerabilities. In both our 37-country sample (Fig 3c) and in a separate non-overlapping sample of developing countries for which we have cross-sectional household-level wealth data (Fig S9a), we find strong differences in temperature exposure by income level, with poorer districts being disproportionately very cold or (particularly) very hot. However, this relationship is mainly driven by differences in exposure between countries rather than within them; average cross-sectional relationships within countries are still concave but much weaker (Fig 3d and Fig S9b). This weaker within-country gradient is consistent with drivers of economic convergence between rich and poor regions – for instance, government transfer programs, or the mobility of factors of production such as labor, capital, or ideas – being stronger within countries than between them.¹⁹

To assess differences in vulnerability, we assign each district in our data to its quintile in the global income distribution, and study whether temperature-growth responses to differ by quintile (SI Section 2). At most exposures, estimated responses in each quintile are not statistically different than the pooled global response, although results for individual quintiles are noisier and somewhat sensitive to how quintiles are defined (Fig S10-S11). As another approach, we restrict the analysis to the US, where our dataset consists of 380 US metropolitan areas endowed with a wide range of average temperatures but all falling into the top income quintile in our global sample. In this sample of very wealthy cities, the estimated response function is again noisier but nearly indistinguishable from the pooled global response.

Why might income not be effective in moderating economic responses to temperature? While our data do not allow specific insight into the channels that mediate the economic responses we measure, a host of recent micro-level evidence suggests that temperature can negatively impact key economic inputs even in highly industrialized and climate-protected environments. These include impacts on cognitive function,^{20,21} labor and total factor productivity,^{22,23} and various aspects of health, including mental health.²⁴ While certain protections (e.g. air conditioning) have been shown to moderate impacts for some outcomes,²⁵ they do not appear effective in insulating broader economic output from the negative effects of hot temperatures.

To understand whether our estimated responses actually reflect persistent effects of temperature on output (“growth effects”), or conversely whether regions ‘catch up’ after a temporary increase in temperature (“level effects”), we take two approaches. Following pioneering earlier work,³ we first estimate distributed lag models that allow a given temperature shock to affect output in both current and ensuing years (see SI). Estimates from these models suggest similarly shaped but substantially steeper responses as compared to panels with just contemporaneous temperature (Fig S8, Table S4), albeit with a slightly warmer implied optimal temperature.

Our second approach also builds on earlier literature²⁶ and uses the large number of cross-sectional units in our data to isolate the influence of longer-term temperature trends on output. Simple plots of the raw data suggest that, in most countries, districts that have warmed the fastest since 2000 have grown on average much more slowly than districts that have warmed the slowest (Fig 4a-1). For instance, districts in the top decile of warming in India, Indonesia, Brazil, and China have per capita GDP (GDPc) that have grown a remarkable

56, 30, 12, and 37% less on average since 2000 as compared to districts in the bottom decile of warming. The raw decile comparison for the US does not show this relationship, but this comparison is made off of a very small number of metropolitan areas (14 in each decile); comparisons using all US data suggests faster warming districts did grow more slowly (Fig 4g), but estimates there are somewhat noisy as well.

To quantify these effects of longer-term warming across the full sample, for each district we regress the time series of GDPc and temperature variables on a time trend, and then estimate a long-differences regression of trend in GDPc versus trend in temperature (and trend in temperature squared, for non-linear versions). That is, we estimate:

$$Y_{istd} = \alpha_{isd} + \lambda_{isd}year_t + \varepsilon_{istd} \quad (2)$$

$$T_{ist} = \alpha_{is} + \phi_{is}year_t + \varepsilon_{ist} \quad (3)$$

$$T_{ist}^2 = \alpha_{is} + \phi_{is2}year_t + \varepsilon_{ist} \quad (4)$$

where Y_{istd} is log GDPc and T_{ist} is annual temperature in district i in admin-1 region s and year t and income dataset d . To estimate these trends, we restrict ourselves to districts with at least 10 years of data (which is >10,300 districts, or 93% of our full sample). We then estimate the long differences regression:

$$\hat{\lambda}_{isd} = \beta_1 \hat{\phi}_{is} + \beta_2 \hat{\phi}_{is2} + \eta_s + \rho_d + \epsilon_{isd} \quad (5)$$

where s indexes admin-1 regions/states and d indexes unique dataset used by each district i . To account for time-trending unobservables at the state level, and for the fact that GDP data are often derived from different datasets within the same country, we include both state fixed effects η_s , which remove state-specific average trends in either income or temperature, as well as dataset fixed effects ρ_d , which account for any dataset-specific differences in income trends. The model also includes trends in precipitation and precipitation squared, estimated as for temperature in equations 3 and 4. The identifying assumption in equation 5 is that temperature trends at the district level are uncorrelated with other factors that affect trends in economic activity, once state-level trends are accounted for. One concern is that emissions of pollutants could be correlated with both local economic activity and temperature. While this is not true for long-lived greenhouse gases (e.g. CO₂), which are rapidly mixed in the atmosphere, it could be true for short-lived pollutants such as particulate matter, which are emitted locally due to increased economic activity and also can block sunlight which leads to cooling. In this latter case, economic activity could generate cooling, leading us to spuriously attributed higher output to cooler temperatures. To address this concern, we process high-resolution satellite-based particulate matter estimates for all our districts back to 2000, using data from ref.²⁷ estimate the trend in PM over time as in equation 3 for temperature, and include this trend in the regression in equation 5.

As in the panel growth model, estimates from equation 5 – the marginal effects of which are shown in (Fig 4m) – suggest a nonlinear relationship between warming and income. Point estimates for the entire sample

are negative, suggesting that districts that warmed faster uniformly grew more slowly. Results are robust to the inclusion of particulate matter trends, and do not appear driven by any one country or the EU, and are robust to inclusion or dropping of growth outliers (Fig S12).

Our long difference estimates are not immediately comparable to our panel estimates: the former regresses change in log GDPc on change in temperature, and the latter regresses change in log GDPc on temperature in levels. To make units comparable, we use estimates from equation 1 and 5 to compute the district-level change in income that each model would estimate to have occurred in response to the warming trend estimated in equation 3. For the long differences model, given a sample length of n years, log GDPc would be estimated to change by $\Delta Y_i^{LD} = (\hat{\beta}_1 \hat{\phi}_{is} + \hat{\beta}_2 \hat{\phi}_{is2}) * n$. For the panel model, we first estimate the detrended temperature series $\tilde{T}_{it} = T_{it} - \hat{\phi}_i year_t$ using the estimates from equation 3 above. We then use the panel response $f(\cdot)$ estimated in equation 1 to estimate the growth decrement (or increment) in each year due to the warming, and then sum these log changes across years to get the total change in log income as estimated by the panel model, i.e. $\Delta Y_i^P = \sum_{t=1}^n f(\tilde{T}_{it}) - f(T_{it})$. We then compare ΔY_i^{LD} to ΔY_i^P . Results from this exercise (shown in Fig 4n) suggest that, if anything, short-run panel estimates *understate* the impact of longer-term warming. These results provide strong evidence that hot years have large and persistent effects on output and that these impacts become increasingly negative at warmer baseline average temperatures.

Finally, we use our data to estimate the impact of recent historical warming on the level and distribution of per-capita GDP, focusing on the subset of countries in our sample with data back to 2000. Using our (substantially) more conservative panel estimates, we simulate how much slower or faster each district would have grown in each year over the 2000-2015 period had temperature stayed at its 1951-2000 average, and cumulate these effects over the period to calculate the increase or decrease in total output (Fig 5a-b). Our exercise is similar to comparison with the long differences estimates above, but here we focus warming above long-run historical average rather than very recent (post-2000) warming; this distinction is relevant given substantial warming that occurred throughout much of the world in the second half of the twentieth century.

We find that recent warming has likely reduced output in most districts in our sample, with substantial negative impacts ($> 10\%$ losses) in most tropical districts as well as in warmer temperate districts (Fig 5e-h). However, given weak average temperature-income gradients within countries, we do not find consistent evidence across our sample that recent warming has amplified within-country inequality (Fig 5i); increases in inequality in some countries (e.g. the US) are offset by modest declines in others (e.g. Indonesia). Cumulative overall impacts are nevertheless large at the country level, with recent warming having cost both the US and the EU $> \$4$ trillion in cumulative lost output (Fig 5j).

While it is unknown whether the response functions we estimate can be accurately extrapolated to other districts or countries not in our sample, the consistency of the estimated response when countries are analyzed independently or held out of the pooled estimation, and the similarity in response across income quintiles, suggests our estimates might be globally relevant. If this is true, then the global impact of additional warming might be more negative than past estimates based on country data,⁴ given no clear evidence in the

subnational data of positive effects in cooler regions. Similarly, because wealthier regions do not appear less vulnerable to additional warming, we find no strong evidence to support the notion that future economic development will protect economies from the impacts of warming – a key assumption in some integrated assessment models.²⁸ While it remains the case that unprecedented future adaptation could lessen climate impacts, for example through the invention of novel climate-protective technologies, such technologies did not appear to emerge during the substantial warming of recent decades.

References

- [1] Carleton, T. A. & Hsiang, S. M. Social and economic impacts of climate. *Science* **353**, aad9837 (2016).
- [2] Hsiang, S. *et al.* Estimating economic damage from climate change in the united states. *Science* **356**, 1362–1369 (2017).
- [3] Dell, M., Jones, B. F. & Olken, B. A. Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* **4**, 66–95 (2012).
- [4] Burke, M., Hsiang, S. M. & Miguel, E. Global non-linear effect of temperature on economic production. *Nature* **527**, 235 (2015).
- [5] Johnson, S., Larson, W., Papageorgiou, C. & Subramanian, A. Is newer better? penn world table revisions and their impact on growth estimates. *Journal of Monetary Economics* **60**, 255–274 (2013).
- [6] Jerven, M. *Poor numbers: how we are misled by African development statistics and what to do about it* (Cornell University Press, 2013).
- [7] Hsiang, S., Oliva, P. & Walker, R. The distribution of environmental damages. *Review of Environmental Economics and Policy* **13**, 83–103 (2019).
- [8] Obama, B. The irreversible momentum of clean energy. *Science* **355**, 126–129 (2017).
- [9] Rohde, R., Muller, R., Jacobsen, R., Perlmutter, S. & Mosher, S. Berkeley earth temperature averaging process. *Geoinformatics & Geostatistics: An Overview* **01** (2013).
- [10] Harris, I., Jones, P. D., Osborn, T. J. & Lister, D. H. Updated high-resolution grids of monthly climatic observations—the cru ts3. 10 dataset. *International journal of climatology* **34**, 623–642 (2014).
- [11] Dee, D. P. *et al.* The era-interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the royal meteorological society* **137**, 553–597 (2011).
- [12] Henderson, J. V., Storeygard, A. & Weil, D. N. Measuring economic growth from outer space. *American economic review* **102**, 994–1028 (2012).


- [13] Deryugina, T. & Hsiang, S. The marginal product of climate. Tech. Rep., National Bureau of Economic Research (2017).
- [14] Kalkuhl, M. & Wenz, L. The impact of climate conditions on economic production. evidence from a global panel of regions (2018).
- [15] Skoufias, E. *The poverty and welfare impacts of climate change: quantifying the effects, identifying the adaptation strategies* (The World Bank, 2012).
- [16] Hallegatte, S. & Rozenberg, J. Climate change through a poverty lens. *Nature Climate Change* **7**, 250 (2017).
- [17] Hallegatte, S. *et al.* *Shock waves: managing the impacts of climate change on poverty* (The World Bank, 2015).
- [18] Winsemius, H. C. *et al.* Disaster risk, climate change, and poverty: assessing the global exposure of poor people to floods and droughts. *Environment and Development Economics* (2018).
- [19] Dell, M., Jones, B. F. & Olken, B. A. Temperature and income: reconciling new cross-sectional and panel estimates. *American Economic Review* **99**, 198–204 (2009).
- [20] Heyes, A. & Saberian, S. Temperature and decisions: evidence from 207,000 court cases. *American Economic Journal: Applied Economics* (2018).
- [21] Goodman, J., Hurwitz, M., Park, J. & Smith, J. Heat and learning. Tech. Rep., National Bureau of Economic Research (2018).
- [22] Seppanen, O., Fisk, W. J. & Lei, Q. Effect of temperature on task performance in office environment. Tech. Rep., Lawrence Berkeley National Laboratory, Berkeley, CA (US) (2006).
- [23] Zhang, P., Deschenes, O., Meng, K. & Zhang, J. Temperature effects on productivity and factor reallocation: Evidence from a half million chinese manufacturing plants. *Journal of Environmental Economics and Management* **88**, 1–17 (2018).
- [24] Burke, M. *et al.* Higher temperatures increase suicide rates in the united states and mexico. *Nature climate change* **8**, 723 (2018).
- [25] Barreca, A., Clay, K., Deschenes, O., Greenstone, M. & Shapiro, J. S. Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy* **124**, 105–159 (2016).
- [26] Burke, M. & Emerick, K. Adaptation to climate change: Evidence from us agriculture. *American Economic Journal: Economic Policy* **8**, 106–40 (2016).

- [27] Van Donkelaar, A. *et al.* Global estimates of fine particulate matter using a combined geophysical-statistical method with information from satellites, models, and monitors. *Environmental science & technology* **50**, 3762–3772 (2016).
- [28] Diaz, D. & Moore, F. Quantifying the economic risks of climate change. *Nature Climate Change* **7**, 774 (2017).


Figures

Figure 1: **Summary of subnational GDP data.** **a.** We compile available GDP per capita data from 11,189 districts across Brazil, China, Colombia, the EU, India, Indonesia, and United States, resulting in 154,244 district-year observations. Colors indicate number of years of data available per district (median=15). **b.** Distributions of annual average temperatures and **c.** GDP per capita across district-years in study sample resemble those across country-years in the global sample. **d.** The seven study countries/regions (blue dots) are some of those with largest total GDP and overall statistical capacity as measured by the World Bank.


a



b



c



d




Figure 2: **District-level GDP growth responds non-linearly to temperature.** **a.** Quadratic response of growth in GDP per capita to annual temperature fluctuations, controlling for linear and quadratic precipitation, district and state-year fixed effects, as well as data source fixed effects (n=154,244 district-year observations). Blue lines show 1000 bootstraps. **b.** Results are robust to higher order polynomials and cubic splines, and to alternate weather data. The main effect from (a) is shown in black, with the light blue region the 95% confidence interval. The country-level temperature response from Burke, Hsiang and Miguel 2015 (BHM) is shown by the dark blue line. **c.** Based on the temperature response function in (a), most regions in the world would experience reduced economic growth for an additional 1°C increase in temperature above the 2001-2015 average temperature. Our temperature-output response is estimated on data from the countries outlined in black.

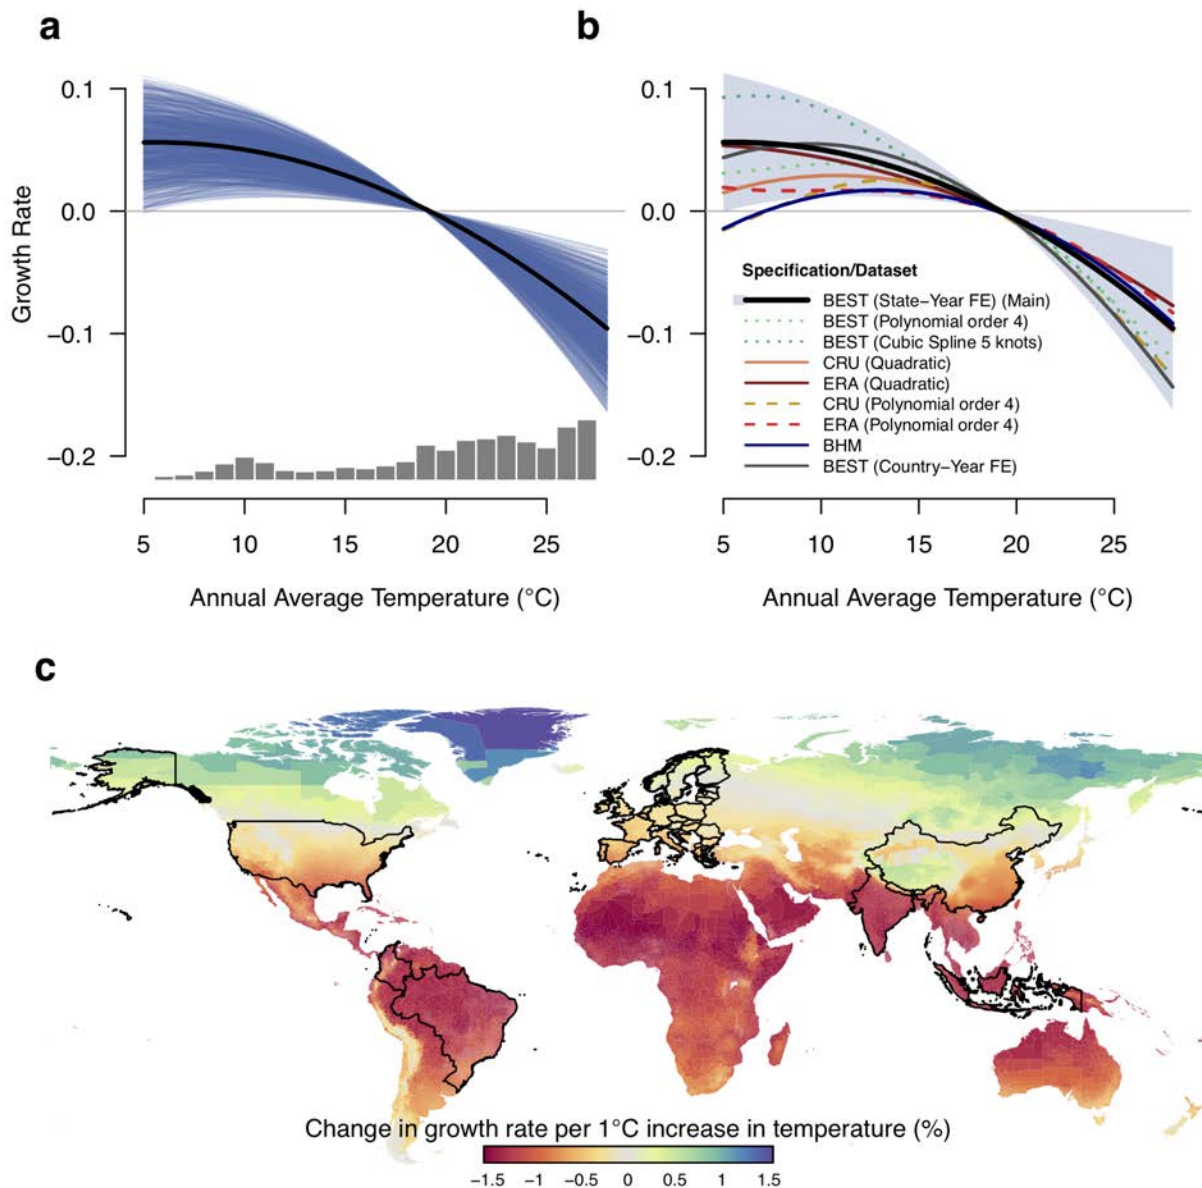


Figure 3: Understanding distributional impacts of warming. Differences in the impacts of warming temperature between rich and poor regions could occur due to differential vulnerability at the same level of exposure, as shown schematically in (a), and/or to similar vulnerability but differing exposures (b). In cross-sectional data, temperature exposure is highly correlated with income across all districts in our global sample (c), but much less correlated across districts within the same country (d). Density of underlying data are depicted in grey heatmap, polynomial fits to the data are given by the blue lines. e. Quadratic responses of growth in GDP per capita to annual temperature fluctuations for each income quintile. Quintiles are assigned relative to the global distribution when districts enter the dataset and remain fixed throughout. Histograms at base show temperature exposure in each income quintile. At most exposures, estimated responses in each quintile are not statistically different than the pooled global response shown in black. f. The temperature response function estimated separately for US metro areas, all of which fall in the top quintile, is noisier but nearly indistinguishable from the pooled global response.




Figure 4: Long difference estimates of the response of income to temperature. **a-f.** The evolution of GDP per capita over time in each country, comparing districts that warmed the most since 2000 in each country (red lines, top decile) versus districts that warmed the least (blue lines, bottom decile). Thin lines are individual districts, thick lines are the average over districts in each decile. Sample is restricted to districts with average temperature $>15^{\circ}\text{C}$, the temperature at which panel estimates suggest the marginal effect of warming is robustly negative (See Fig S8a). Values in all districts are indexed to one in the base year in each country. Values reported at the right of each plot are the percent difference in income in the last year between slowest and fastest warming districts. **e-h.** Scatter plots of the annual average change in income versus the annual average change in temperature over the study period in each country; each point is a district, with points colored as a function of baseline average temperature. The slope of the estimated relationship and its standard error are shown at the top of each plot. Annual average increases are estimated by regressing either GDP/cap or temperature on a time trend. **m.** Marginal effects estimates from a long differences regression of annual average change in per capita GDP on annual average change in temperature, estimated on the pooled sample of all districts (equation 5). Blue lines indicate 1000 bootstraps. **n.** Estimated change in GDP/cap due to observed warming since 2000, using the panel growth model (x-axis) or the long-differences levels model shown in (m). Grey line is 45-degree line. See text for details on calculation. Long differences estimates are almost uniformly larger in absolute value than panel estimates, suggesting that short-run panel estimates understate the longer-term impacts of warming.


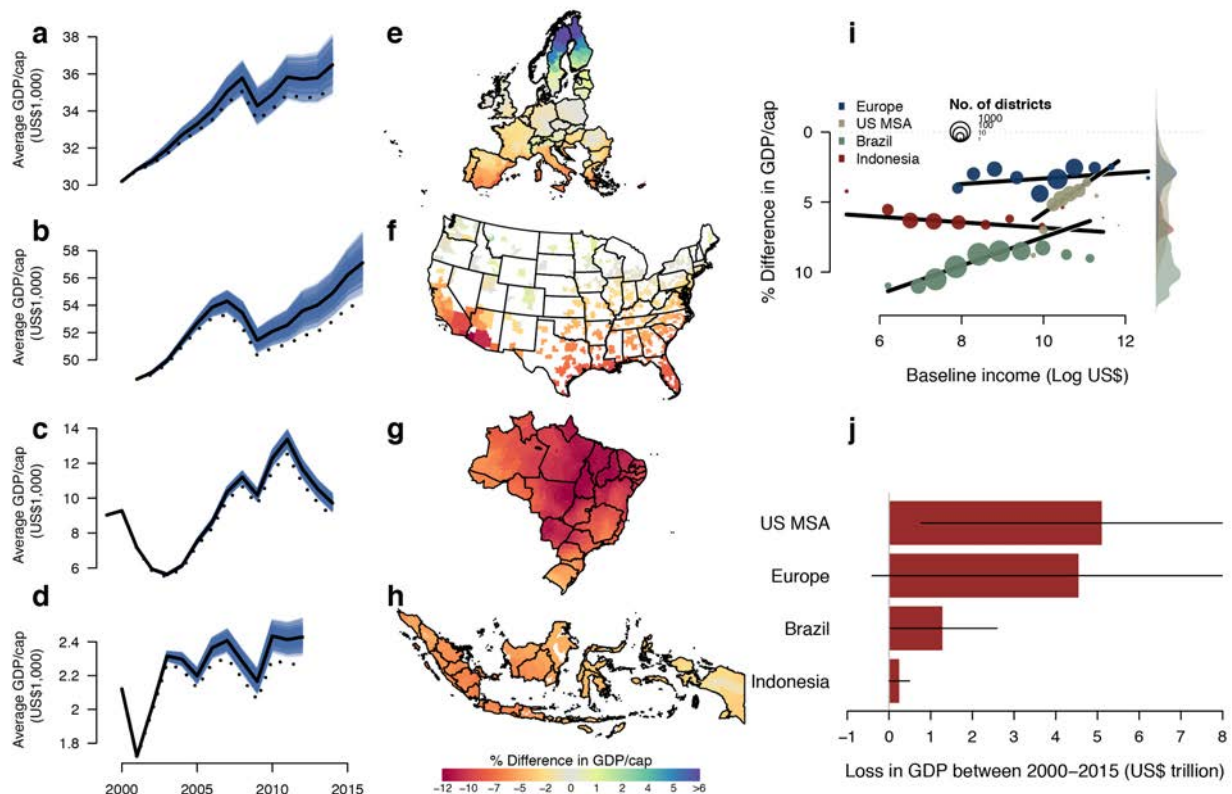


Figure 5: **The effect of recent warming on GDP per capita is spatially heterogeneous, with large inequality across countries and modest inequality within countries.** We simulated the counterfactual evolution of GDP per capita if temperatures had remained at 1951-2000 baseline averages, instead of observed values, during the 2000-2015 period. Comparison between observed GDP/cap (dotted) and counterfactual GDP/cap (solid lines) plotted for **a** EU countries, **b** US metro areas, **c** Brazil and **d** Indonesia. Blue lines show 1000 estimates from bootstrapping the underlying panel estimates. **e-h**: district-level differences between counterfactual and observed GDP per capita in 2015. **i** The difference between counterfactual and observed GDP per capita in 2015 is plotted against baseline income; bubbles represent average percent change in GDP across districts within a region at a given baseline income. Differences in impacts are larger between than within countries, with larger percentage losses in poorer countries. **j** Cumulative losses of GDP over the 15 year period due to warming above 1951-2000 baseline. Whiskers show 95% confidence intervals across 1000 bootstraps.



Supplementary Information

1 Data Construction

1.1 GDP Data

We gathered subnational data on GDP per capita or GDP and population from national statistical agencies which publicly provide these data in any formats, digitized or non-digitized. We obtained data at the most granular level available, usually at the second-level administrative division (“districts”) of each country. First we started by accessing the statistical agency website of every country in the world in alphabetical order. Then we selected the GDP per capita data that were available in all districts in the majority of states of a country, have growth rates within 97.5% of those in all available sample (Fig S3d-e), and pass the nightlights validation test (Section 1.1.1). Despite national variations in the naming convention of administrative units, we use the all-encompassing terms “state” and “district” for the first and second-level administrative divisions of each country. Table S1 lists the data source for each country in our dataset, the number of available district-year observations, and if we processed the data originally contained in non-digitized documents and spreadsheets. In the discussion below, we refer to the European Union as a country.

We fluidly use the district definition for three countries. China’s statistical agencies provide more granular GDP per capita data at the third-level administrative division. We use these data and assign the provinces in which districts reside as their parent states. United States only has four-year prototype GDP per capita data for the universe of districts,²⁹ so the GDP per capita of Metropolitan Statistical Areas (MSAs) serve as US’ most disaggregated subnational accounting. We assign the first state mentioned in the name of interstate MSA as its parent state. Lastly, we follow European Union Statistics’ second-level divisions (NUTS 3) which match imperfectly with member countries’ official second-level administrative units.

Except for GDP per capita data from US, Europe and Indonesia, data from all other countries were assembled from multiple sources. All districts in Brazil experienced a switch in data source because data were contained in two separate files. Similarly, Colombia experienced a switch in data source from CEDE Panel Municipal to DANE. China had used different measures of national accounting before 1989/1990: “社会总产值” (Total Social Output Value) or “工农业总产值” (Gross Value of Industrial and Agricultural Output). To account for changes in growth rates that were caused by the idiosyncratic nature of dataset switches and measurement changes, we include data source fixed effects in our regression analyses. When we collapsed the panel data into a cross-section of districts in the long-difference analysis, the data source fixed effects is the concatenation of all data sources across years used in each district.

As district level growth rates, particularly in China, have extreme outliers (Fig S3d-e), our baseline model trims the 2.5% tails in the growth distribution of our global sample. We attempted alternative approaches to

Table S1: Summary of GDP data sources used in our dataset

Data Source	Number of district-Years	Number of years	Digitization	Data Source	Number of district-Years	Number of years	Digitization
Brazil				China			
Brazil IBGE 1999-2011 ³⁰	64694	12	No	Shanghai Statistical Yearbook ⁵⁴	31	31	Yes
Brazil IBGE 2010-2014 ³¹	16164	3	No	Shanxi Statistical Yearbook ⁵⁵	1667	16	Yes
China				Shanxi Statistical Yearbook (Gross National Product) ⁵⁵	93	1	Yes
Anhui Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ³²	10	1	Yes	Sichuan Statistical Yearbook ⁵⁶	2122	14	Yes
Beijing Statistical Yearbook ³³	33	33	Yes	Tianjin Statistical Yearbook ⁵⁷	293	21	Yes
Chongqing Statistical Yearbook ³⁴	548	23	Yes	Tianjin Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ⁵⁷	7	2	Yes
Chongqing Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ³⁴	9	2	Yes	Xinjiang Statistical Yearbook ⁵⁸	1447	19	Yes
Fujian Statistical Yearbook ³⁵	1082	22	Yes	Xinjiang Statistical Yearbook (Gross National Product) ⁵⁸	152	2	Yes
Fujian Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ³⁵	94	2	Yes	Tibet Autonomous Region Statistical Yearbook (Dual Sector: Gross Domestic Product in Agriculture and Industry) ⁵⁹	114	3	Yes
Gansu Statistical Yearbook ³⁶	1077	17	Yes	Yunnan Statistical Yearbook ⁶⁰	1445	17	Yes
Gansu Statistical Yearbook (Net National Income) ³⁶	64	2	Yes	Zhejiang Statistical Yearbook ⁶¹	751	15	Yes
Guangdong Statistical Yearbook ³⁷	1206	21	Yes	Colombia			
Guangdong Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ³⁷	104	2	Yes	CEDE Panel Municipal 2000-2010 ⁶²	9356	9	No
Guangxi Statistical Yearbook ³⁸	1553	21	Yes	Colombia DANE 2011-2015 ⁶³	4012	4	No
Guizhou Statistical Yearbook ³⁹	915	14	Yes	Europe			
Hainan Statistical Yearbook ⁴⁰	265	17	Yes	EU Statistics ⁶⁴	19247	15	No
Hainan Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ⁴⁰	7	1	Yes	Indonesia			
Hainan Statistical Yearbook (Gross National Income) ⁴⁰	38	4	Yes	World Bank INDO-DAPOER, based on Indonesia Statistics ⁶⁵	5099	12	No
Hebei Statistical Yearbook ⁴¹	1267	11	Yes	India			
Heilongjiang Statistical Yearbook ⁴²	1262	23	Yes	Govt. of Andhra Pradesh ⁶⁶	167	9	Yes
Heilongjiang Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ⁴²	35	1	Yes	Govt. of Assam ⁶⁷	23	1	Yes
Henan Statistical Yearbook ⁴³	1475	16	Yes	Govt. of Bihar ⁶⁸	221	6	Yes
Hubei Statistical Yearbook ⁴⁴	156	2	Yes	Govt. of India Planning Commission ⁶⁹	3258	9	No
Hunan Statistical Yearbook ⁴⁵	1403	14	Yes	Govt. of Jharkhand ⁷⁰	51	3	Yes
Inner Mongolia Statistical Yearbook ⁴⁶	235	3	Yes	Govt. of Karnataka ⁷¹⁻⁷⁴	183	8	Yes
Jiangsu Statistical Yearbook ⁴⁷	905	20	Yes	Govt. of Kerala ^{75,76}	14	1	Yes
Jiangxi Statistical Yearbook ⁴⁸	1003	15	Yes	Govt. of Madhya Pradesh ⁷⁷	225	5	Yes
Jiangxi Statistical Yearbook (Dual Sector: Total Social Output in Agriculture and Industry) ⁴⁸	197	3	Yes	Govt. of Maharashtra ^{78,79}	102	3	Yes
Jiangxi Statistical Yearbook (Gross National Product) ⁴⁸	121	2	Yes	Govt. of Odisha ⁸⁰	180	6	Yes
Jilin Statistical Yearbook ⁴⁹	726	21	Yes	Govt. of Punjab ^{81,82}	32	2	Yes
Liaoning Statistical Yearbook ⁵⁰	155	20	Yes	Govt. of Punjab, India Statistical Portal ^{81,82}	34	2	Yes
Ningxia Statistical Yearbook ⁵¹	292	21	Yes	Govt. of Rajasthan ⁸³	63	2	Yes
Qinghai Statistical Yearbook ⁵²	385	14	Yes	Govt. of Rajasthan, India Statistical Portal ⁸³	122	4	Yes
Shandong Statistical Yearbook ⁵³	414	25	Yes	Govt. of Tamil Nadu ⁸⁴	60	2	Yes
				Govt. of Uttarakhand ⁸⁵	64	5	Yes
				United States			
				US BEA ⁸⁶	5715	15	No

dealing with growth outliers, such as trimming the 0.5%, 1% tails, or top- and bottom-coding observations at said percentiles, and results were not very sensitive to this choice (Fig S3c).

Finally, in order to spatially merge per capita GDP growth rate with weather data, we performed string merging of districts and their corresponding shapefiles on names or alternative names in case of district name changes. In order to achieve the most balanced panel of districts over time, districts that split or merged are aggregated to their largest unions and these unions are kept constant over the sample period. These districts are usually observable in the data by having missing data before they split from parent districts or after they merged with other districts. We obtained a list of territorial/name changes from Statoids⁹³ and the Chinese Government.⁹⁴ Table S2 reports data source of shapefiles. Any remaining errors in our merging will likely amplify noise in our assignment of temperature to districts, attenuating regression estimates.

Table S2: List of shapefile sources used in our dataset

COUNTRY	SOURCE
Brazil	IBGE ⁸⁷
China	All China Marketing Research ⁸⁸
Colombia	GADM ⁸⁹
Europe	EU Statistics ⁹⁰
Indonesia	GADM ⁸⁹
India	ML InfoMap ⁹¹
US MSA	TIGER/Line ⁹²

1.1.1 Data Validation

As one check on data quality, we follow Henderson et al¹² and compare our trimmed subnational GDP data with data on satellite-derived night-time lights. We obtained district night-time lights observations by taking area-weighted average of the digital numbers of stable light pixels⁹⁶ within a district. Similar to Henderson et al, we run a regression which absorbs time-invariant and time-trending observables that affect economic productivity within a state in a country:

$$Y_{istd} = \beta_1 I_{ist} + \mathbf{X}'_{ist} \gamma + \alpha_i + \rho_d + \eta_{st} + \varepsilon_{ist}$$

where Y_{ist} measures the natural logarithm of GDP of district i in state s in year t with dataset d , I_{ist} is log average digital number of night-time lights, and \mathbf{X}_{ist} is a vector of controls as used in Henderson et al to account for the spatial distribution of night-time lights (number of pixels with maximum digital number (top-coded), minimum digital number (unlit), and the Gini index of light distribution). Figure S2 plots the β_1 coefficients from pooling all countries into one regression as well as from running a separate regression for each individual country. Ecuador did not pass the nightlights validation test ($\beta_1 = -0.229$) and has been removed from the final dataset.

1.2 Weather Data

Our main annual average temperature data are derived from Berkeley Earth Surface Temperature (BEST)⁹ gridded data at 1° resolution. We use two additional temperature datasets at different resolutions, namely European Centre for Medium-Range Weather Forecasts' ERA-Interim¹¹ at 0.75° resolution and University of East Anglia's CRU TS 4.01¹⁰ at 0.5° resolution, as checks of robustness to our main results. Because the temporal resolution of the three temperature datasets is monthly, we first took an area-weighted average of temperature pixels within each district and then took an average across twelve months to get the annual temperature values. Only in India's GDP per capita dataset did subnational accounting follow a fiscal year (April 1-March 31) and so annual average temperatures for Indian districts are computed from April to

March, while all other districts follow January to December annual average temperature. Our precipitation data come from University of East Anglia’s CRU TS 4.01 data.¹⁰ Finally, as a robustness to our temperature measurement, we also used number of °C days that fell into a given set of intervals (or “bins”) in a year to measure temperature variation, the data for which came from daily BEST gridded data.⁹

1.3 Simulating the consequences of aggregation

To better understand and illustrate the gains from having sub-national granular data on both output and temperature, we show how patterns of spatial correlation in local temperature anomalies can lead to noisy inference in country level data. In particular, after taking out temperature trends, historical temperature anomalies across the largest metropolitan in the US are often negatively correlated, and the same is true in many other regions (Fig S1a-b). To see the implication of these negative correlations for country-aggregated data, consider two years which at the country level had the same output and same average temperature: one year in which temperature and output were normal in both the East and West, and another in which the temperature (output) was high (low) in the East and low (high) in the West. This second year, while showing strong local responses to temperature, would show no aggregate response at the country level, and would not contribute variation to our understanding of temperature’s effects on output (nor, importantly, would it bias our understanding; it would simply not provide information).

To quantitatively explore these implications, we run the following simulation. We define the true response of growth to temperature at the district level as $y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \psi_t + \epsilon_{it}$, using the detrended observed T_{it} time series in each district i and setting $\beta_1 = 0.02, \beta_2 = -0.001$. $\psi_t \sim N(0, 2)$ is a common growth shock in each year and $\epsilon_{it} \sim N(0, 2)$ is a district specific shock. We then use these constructed outcomes and observed temperatures to estimate a district fixed effects regression

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \alpha_i + \delta_t + \epsilon_{it} \quad (6)$$

then save the coefficient estimates on temperature. We then generate growth aggregates $y_t = \sum_i y_{it}$, and temperature aggregates $T_t = \sum_i T_{it}$ and $T_t^2 = \sum_i T_{it}^2$. We then run the time series regression at the country level:

$$y_t = \gamma_1 T_t + \gamma_2 T_t^2 + \theta * year_t + \varepsilon_t \quad (7)$$

where $\theta * year_t$ is a country time trend. We again save the coefficient estimates on temperature. We repeat this process 1000 times.

In Figure S1c, black lines give responses estimated using the district data, and red lines give the estimated response from the country time series. Panels (d) and (e) show distribution of parameter estimates across bootstraps for the district panel model (black lines) and the aggregated country model (red lines). Aggregated country-level estimates are unbiased but much noisier than district estimates and can get the sign of the

true effect wrong. In essence, given the structure of temperature anomalies in our data, aggregation can substantially reduce the variation that can be learned from. Our district-level data are able to leverage variation that country aggregates cannot.

2 Estimation

Our main panel fixed effects regression is given in the main text (Equation 1). To study quintile-specific responses, we interact the temperature response function with quintile dummies:

$$y_{istd} = \sum_j \mathbf{D}_j * [f(T_{ist}) + f(P_{ist})] + \rho_d + \alpha_i + \eta_{st} + \varepsilon_{istd} \quad (8)$$

where \mathbf{D}_j is an indicator for whether district i falls into the j -th quintile.

To study whether effects of temperature on growth are persistent, or whether there is “catch up” in the next period, we follow refs^{3,4} and estimate a distributed lag model and include up to five years of temperature or precipitation lags:

$$y_{istd} = \sum_{L=0}^5 [f(T_{is(t-L)}) + f(P_{is(t-L)})] + \rho_d + \alpha_i + \eta_{st} + \varepsilon_{istd} \quad (9)$$

with persistence assessed by adding up the contemporaneous and lagged coefficients. That is, with $f(T_{ist}) = \beta_1 T_{ist} + \beta_2 T_{ist}^2$, then the marginal effect of a degree of warming is:

$$\frac{\partial y_{istd}}{\partial T_{ist}} = \sum_{L=0}^5 \beta_{1,L} + 2 \sum_{L=0}^5 \beta_{2,L} T_{is(t-L)d} \quad (10)$$

We also estimate a long-differences regression as described in the main text.

3 Calculating the cost of recent warming

To understand the cost of recent warming, we combine our panel estimates from equation 1 of the marginal impact of warming with how much annual temperatures in each district-year differed from their long-term (1951-2000) average. Using our panel estimates, we simulate how much slower or faster each district would have grown in each year over the 2000-2015 period had temperature stayed at its 1951-2000 average, and cumulate these effects over the period to calculate the increase or decrease in total output (Fig 5a-b) We emphasize that our estimates are agnostic to the cause of recent warming, and do not necessarily represent the impact of recent *anthropogenic* warming.

To calculate the cumulative cost of warming for each district we first compute the annual cost of warming from 1951-2000 baseline level:

$$c_{it} = (\beta_1 T_{it} + \beta_2 T_{it}^2) - (\beta_1 \bar{T}_i + \beta_2 \bar{T}_i^2) \quad (11)$$

where T_{it} is the observed temperature of district i in year t , \bar{T}_i is its average temperature over baseline years of 1951-2000, and β_1 and β_2 are our coefficient estimates for the effects of temperature on growth rate from Equation 1. Thus absent warming, district i 's counterfactual growth rate and income in year t could be

$$\begin{aligned} \hat{y}_{it} &= y_{it} + c_{it} \\ \hat{Y}_{it} &= Y_{it-1} + \hat{y}_{it} \end{aligned} \quad (12)$$

where hatted terms are counterfactual and Y indicates GDP per capita in log terms, and y_{it} is the observed growth rate. The cumulative loss in GDP for district i is then

$$\text{LOSS}_i = \sum_{t \in \Theta} N_{it} (e^{\hat{Y}_{it}} - e^{Y_{it}}) \quad (13)$$

where N_{it} is the population of district i in year t and the sum is computed over all available years Θ in our sample. We focus on the four countries/regions with data back to 2001: the US, EU, Brazil, and Indonesia.

Figure S1: **Patterns of spatial correlation in local temperature anomalies lead to noisy inference in country-aggregated data.** **a.** After taking out temperature trends, historical temperature anomalies across the largest metropolitan in the US are often negatively correlated. Each dot represents the temperature anomaly in two metro areas in a year. **b.** Distribution of pairwise temperature anomalies across all districts within our study sample, for each region. Temperature anomalies are fairly highly correlated in some countries/regions (e.g. Indonesia, India), but often poorly correlated in others (e.g. US, Brazil, China). **c-e.** Simulated effects of aggregation on the estimated effect of temperature on economic growth, using data from the US+EU (see SI Section 1.3). Black lines give response estimated using the district data (1000 resamples), relative to a year at 20C, red lines give the estimated response after first averaging growth, temperature, and temperature squared across districts within a given year and then estimating temperature response using the country time series. **(d)** and **(e)** show distribution of parameter estimates across bootstraps for the district panel model (black lines) and the aggregated country model (red lines). Aggregated country-level estimates are unbiased but much noisier than district estimates and can get the sign of the true effect wrong.




Figure S2: **Validation of district GDP data against night-time lights.** We follow Henderson et al (2012)¹² and estimate panel models of the relationship between log GDP and log night-time lights at the district level, and report the estimated elasticities for each country. All models include district FE and state-year FE, and (as in Henderson et al) additionally control for top-coded and unlit grid cells or the spatial distribution of night-time lights within a region as denoted with the lighter shaded bars. Results in each country are similar to the global country-level results in Henderson et al, suggesting our local-level data are picking up meaningful variation in economic activity.




Figure S3: **Robustness of main results.** **a** Inclusion of linear or quadratic country-specific time trends leads to much steeper responses of growth to temperature. **b** The pooled response does not appear to be dominated by individual countries; dropping each country sequentially from the sample and re-estimating the response across remaining countries leads to similar estimated responses (grey area is the bootstrapped confidence interval on the full sample). **c** Results are similar under alternative approaches to dealing with outliers in the growth distribution, including either dropping top and bottom 0.5%, 1%, or 2.5% tails of observations each country's growth distribution, or similarly top- or bottom-coding observations at the same percentiles (e.g. setting anything below the 0.5th percentile to the value of the 0.5th percentile). Retaining all outliers leads to a much steeper estimated response (green line). **d** Raw distribution of district growth observations in each country, and **e** distributions after dropping top and bottom 2.5% of observations.

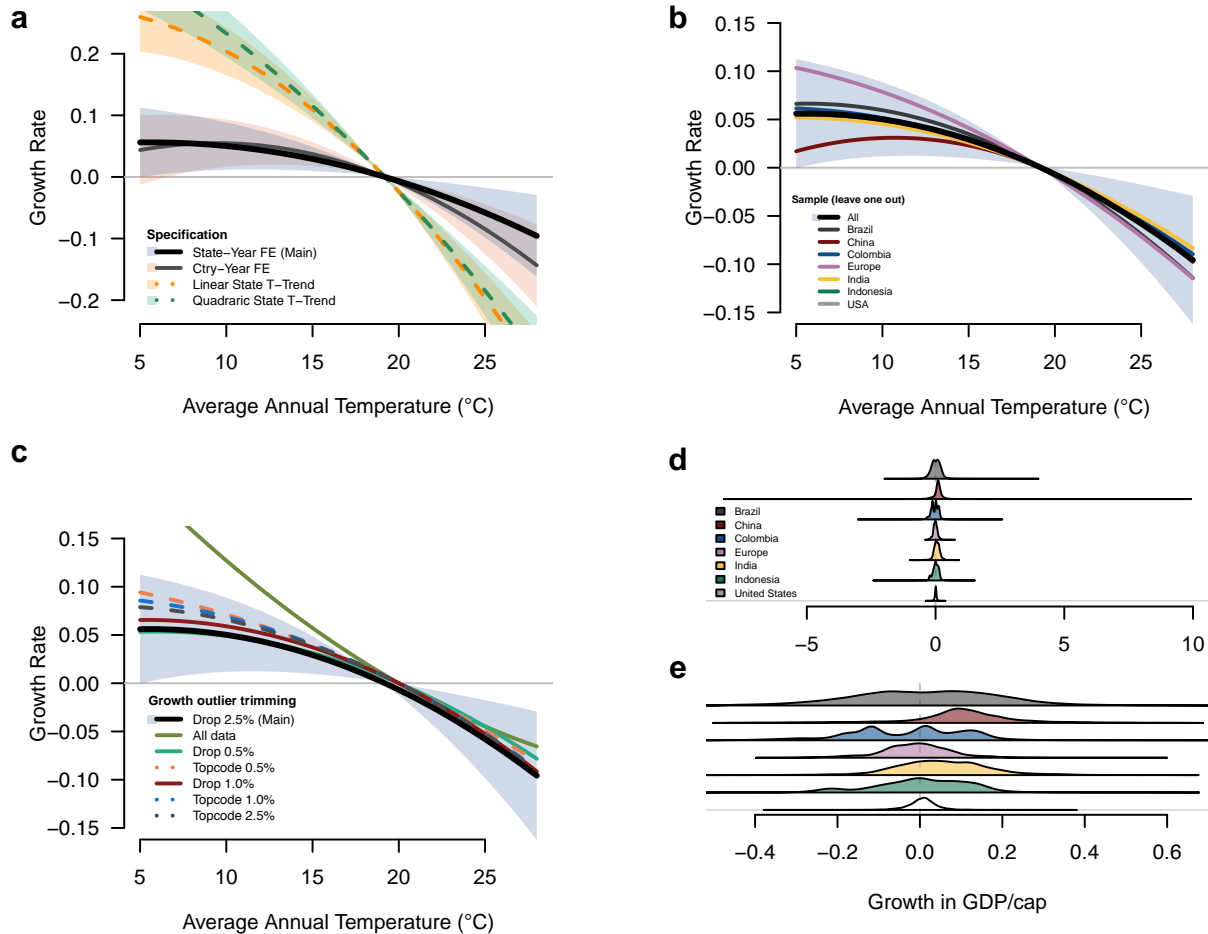


Figure S4: **Robustness of main results to inclusion of spatial lags in growth rate and temperature.** **a.** Change in the baseline estimate upon inclusion of spatial lags of the dependent variable (growth rates) of neighboring districts within 100, 300 and 500 kilometers in the baseline regression. **a** Change in the baseline estimate upon inclusion of spatial lags of temperature in neighboring districts. Estimates on the temperature lags themselves (not pictured) are similar to own-district estimates – i.e. concave and steeply declining at hot temperatures – suggesting spillover effects.

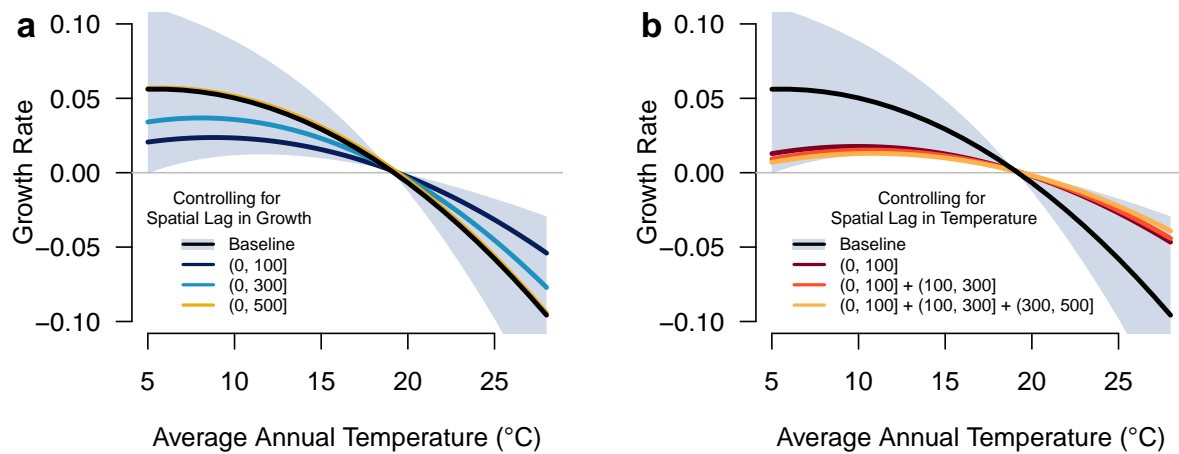


Figure S5: **Estimated response in individual countries/regions.** Each row shows estimated response of growth to temperature in each country (+ EU region) across three temperature datasets: BEST (left column), CRU (middle), and ERA (right). Histograms at bottom of each panel show temperature exposure over the sample period in each country. Lines represent different model specifications as listed in legend at bottom; black line is our baseline model with district and state-by-year FE. Largest disagreements in the response of growth to temperature across climate datasets occur in countries where climate datasets are poorly correlated with each other (see Fig S6

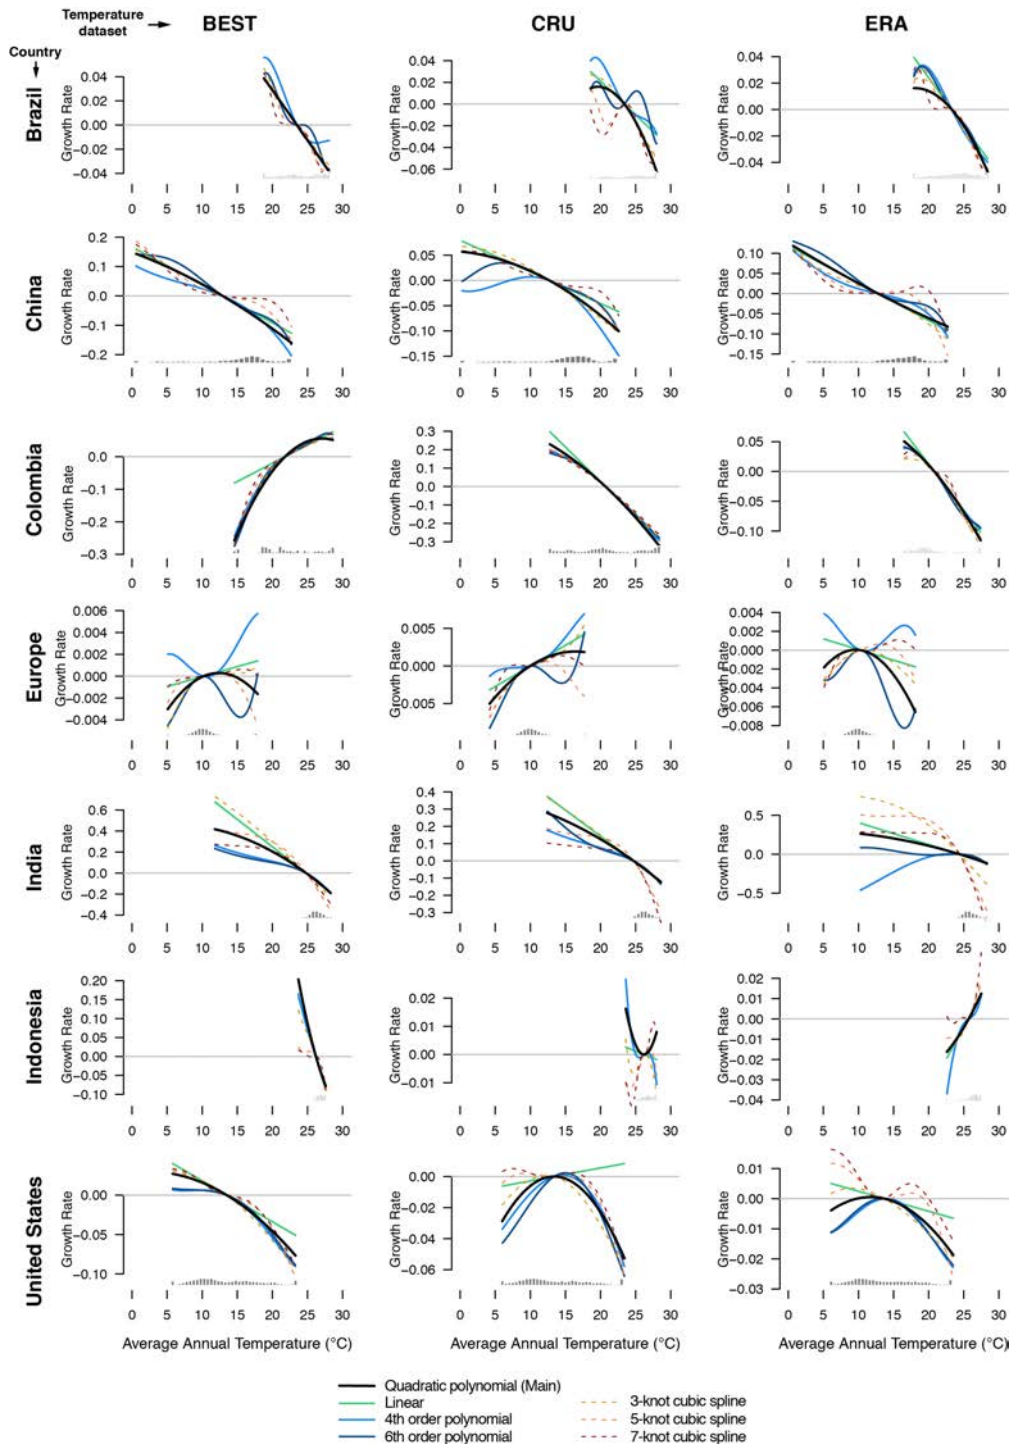


Figure S6: **Correlation in temperature anomalies across temperature datasets.** Values indicate pairwise correlation in temperature values between our three temperature datasets (BEST, CRU, ERA), after removal of district and state-by-year FE. Countries with low correlation across climate datasets also have much more variation in estimated temperature impacts (Fig S5).




Figure S7: **Robustness of main results to temperature measured in degree-days** We study the effects of interannual temperature changes on the changes in growth rate, where temperature is measured as the number of days a district experiences temperatures within -10 and -7 degrees, -7 and -4 degrees, and so forth, which we call “temperature bins”. The bins are closed in the lower temperature bound and open in the upper bound. Blue shaded area indicates confidence interval around the connected point estimate of each temperature bin.




Figure S8: **Persistent effects of temperature on output.** To understand whether growth rates rebound after a transitory, hotter than normal year, we compare our baseline model with no temperature lags (**a** our main model), with a distributed lag model that includes contemporaneous temperature as well as five year lags in temperature (**b**). Dark blue lines are bootstrapped estimates of the marginal effects from these regressions (1000 resamples), and black lines the main effect.

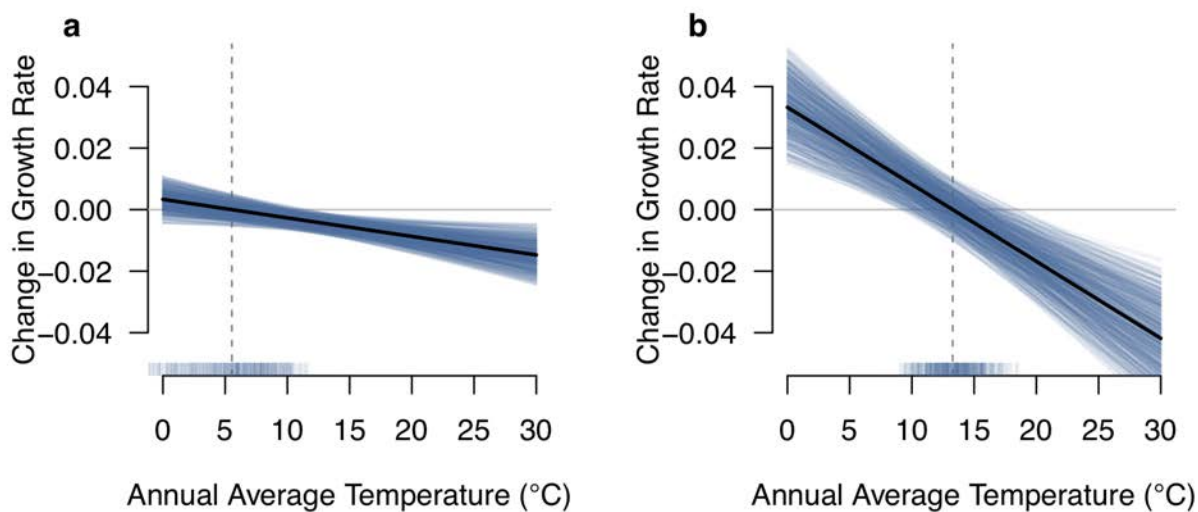


Figure S9: **Cross-sectional relationship between household wealth and temperature is negative across full sample but flat within countries.** To confirm that our exposure disparity in Figure 3 is representative, we assemble geo-referenced cross-sectional data on household asset wealth from 36 developing countries (non overlapping with our sample countries in the main analysis) from the Demographic and Health Surveys, and match these wealth measures to average temperature exposure at each household's location. Households with less asset wealth are much more likely to be exposed to warmer average temperatures when looking across all households in the global sample (a), but this disparity is much weaker when looking only within countries (b).




Figure S10: **Robustness of quintile response.** We explore the robustness of our income-specific response functions shown in Fig 3 to alternate ways of defining the quintiles or specifying time controls (Section 2). The response for each quintile is shown in colored lines from dark red (poorest) to bright orange (richest). **(a)** is our main specification in which income quintile is assigned to a district relative to the global sample when it enters our dataset and is fixed throughout. **(b)** is the specification in which income quintile may change over time, moving depending on districts' relative position in a given year's distribution. We control for time-trending and time-invariant observables that may affect temperature and growth rate separately in each quintile by including quintile-year fixed effects in **(a)** and **(b)**, whereas **(c)** and **(d)** are the fixed and moving quintile specifications without the quintile-year fixed effects. All responses are fairly similar except for the poorest quintiles in **(c)**.

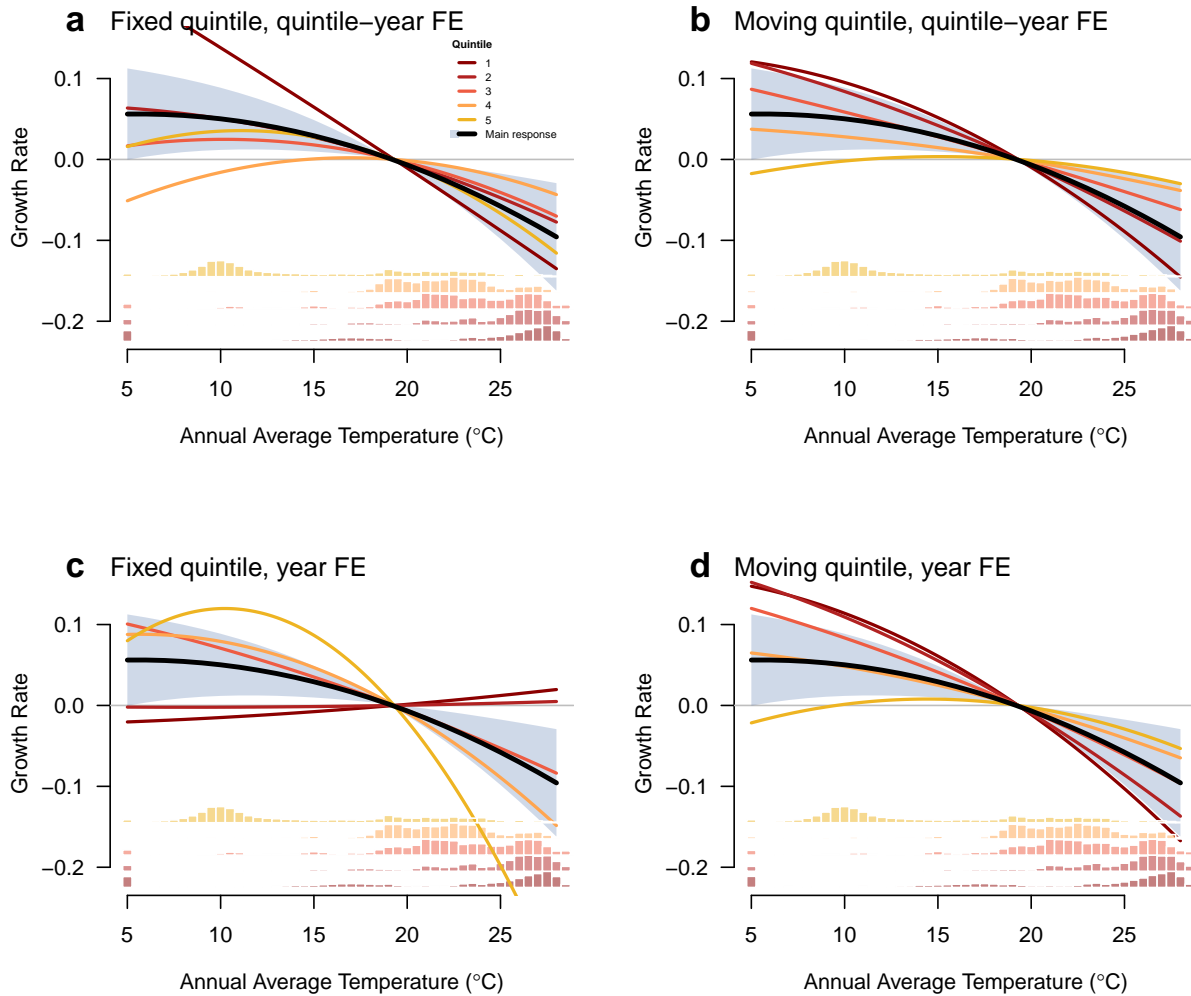


Figure S11: **Associated marginal effects from each quintile response.** For each of the quintile specification in Figure S10, we compute the marginal effect of change in growth rate per °C change at a given temperature. **(a)** is our main specification using fixed quintile and quintile-year fixed effects. **(b)** is the specification with moving quintile and quintile-year fixed effects. **(c)** and **(d)** are the fixed and moving quintile specifications without quintile-year fixed effects.

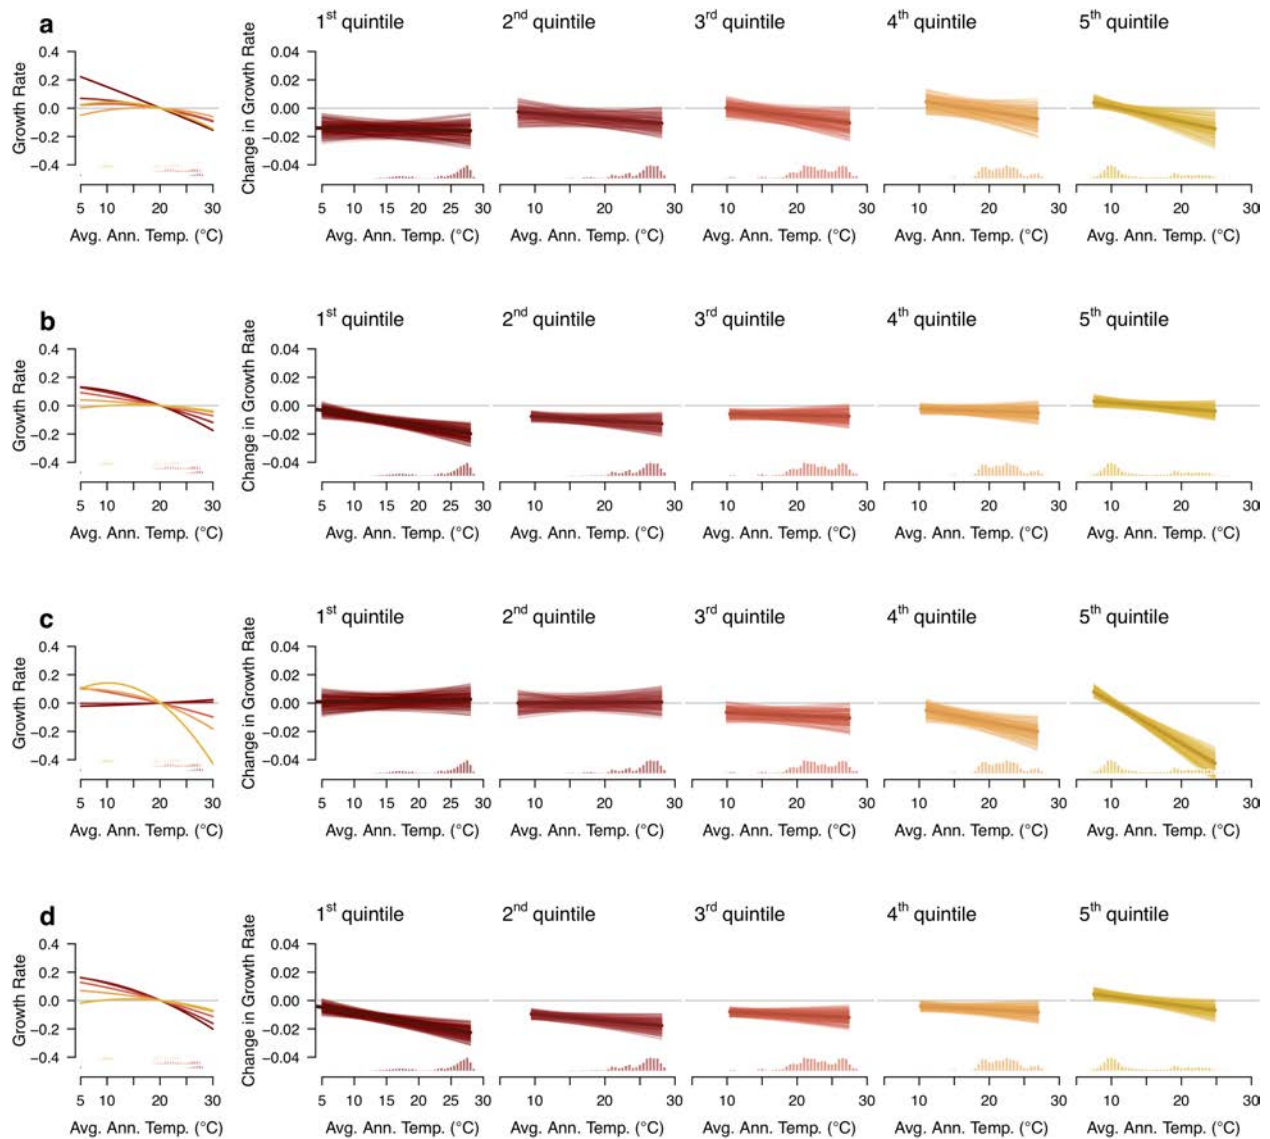


Figure S12: **Robustness of long difference estimates.** Estimates of Equation 5 are robust to exclusion of individual countries (a), exclusion of outliers in long-run growth rates (b), and are unchanged by the inclusion of trends in particulate matter (c).




Table S3: **Regression estimates of the effect of temperature on growth rates**, including results for three climate datasets and various ways of specifying the time controls.

	DEPENDENT VARIABLE: GROWTH RATE			
	(1)	(2)	(3)	(4)
<i>Panel A: BEST Weather Dataset</i>				
Temperature (°C)	0.0033392 (0.0047915)	0.0108371*** (0.0025752)	0.0005566 (0.0022809)	-0.0126719*** (0.0022981)
Temperature ² (°C ²)	-0.0003013** (0.0001408)	-0.0005752*** (0.0000710)	-0.0007785*** (0.0000705)	-0.0004336*** (0.0000716)
Observations	154,117	154,117	154,117	154,117
R ²	0.636	0.562	0.437	0.470
Adjusted R ²	0.584	0.526	0.390	0.424
<i>Panel B: CRU Weather Dataset</i>				
Temperature (°C)	0.0091594** (0.0045649)	0.0090729*** (0.0025191)	-0.0034178 (0.0022799)	-0.0145274*** (0.0023162)
Temperature ² (°C ²)	-0.0004267*** (0.0001367)	-0.0003441*** (0.0000685)	-0.0004833*** (0.0000703)	-0.0002738*** (0.0000722)
Observations	154,229	154,229	154,229	154,229
R ²	0.636	0.562	0.436	0.470
Adjusted R ²	0.584	0.526	0.389	0.423
<i>Panel C: ERA Weather Dataset</i>				
Temperature (°C)	0.0011248 (0.0038087)	0.0070695*** (0.0022956)	-0.0110069*** (0.0020809)	-0.0206612*** (0.0020801)
Temperature ² (°C ²)	-0.0002064** (0.0000942)	-0.0003895*** (0.0000570)	-0.0000936 (0.0000581)	0.0000854 (0.0000583)
Observations	154,227	154,227	154,227	154,227
R ²	0.636	0.562	0.436	0.470
Adjusted R ²	0.584	0.526	0.389	0.423
Time FE	State-Year	Country-Year	Year	Year ²
State Time-trend			✓	✓
district FE	✓	✓	✓	✓
Data source FE	✓	✓	✓	✓
Precip. control	✓	✓	✓	✓

Notes: Table reports the coefficients of the regressions of district GDP per capita growth rate on quadratic annual average temperature. Our baseline specification is reported in panel A using BEST weather temperature on column 1. The included time controls in each specification are shown at the bottom of the table. Column 3 controls for state time-trend, with year fixed effects. Column 4 controls for quadratic state time-trend. Asterisks indicate the level of statistical significance: *p<0.1; **p<0.05; ***p<0.01

Table S4: **Marginal effects of a +1C increase in temperature on the growth rate**, estimated at different temperatures for distributed lag models with zero, one, and five lags. The zero-lag model is our baseline model shown in Figure 2.

Temperature (°C)	DEPENDENT VARIABLE: CHANGE IN GROWTH RATE PER °C TEMPERATURE INCREASE		
	0 Lag	1 Lag	5 Lag
5	0.00033 [-0.00532, 0.00579]	0.00369 [-0.00391, 0.0114]	0.02074*** [0.00703, 0.03531]
10	-0.00269 [-0.00684, 0.00134]	-0.00142 [-0.00723, 0.00433]	0.00822 [-0.00236, 0.01936]
15	-0.0057*** [-0.00992, -0.00153]	-0.00652** [-0.01264, -0.00064]	-0.0043 [-0.01538, 0.00704]
20	-0.00871*** [-0.01449, -0.00292]	-0.01163*** [-0.01999, -0.00368]	-0.01682** [-0.03178, -0.0019]
25	-0.01172*** [-0.0197, -0.00366]	-0.01674*** [-0.0282, -0.00586]	-0.02934*** [-0.04977, -0.00925]
30	-0.01474*** [-0.02516, -0.00417]	-0.02184*** [-0.03674, -0.0077]	-0.04186*** [-0.0684, -0.01596]
35	-0.01775*** [-0.03071, -0.00456]	-0.02695*** [-0.04544, -0.00939]	-0.05438*** [-0.08733, -0.02238]

Notes: Asterisks indicate the level of statistical significance: *p<0.1; **p<0.05; ***p<0.01. 95% confidence interval around each elasticity estimate is reported.

Additional references

- [29] Panek, S., Rodriguez, R. & Baumgardner, F. Research spotlight new county-level gross domestic product. *Survey of Current Business* **99**, 1–105 (2019).
- [30] Instituto Brasileiro de Geografia e Estatística (IBGE), Coordenação de Contas Nacionais. Produto interno bruto dos municípios 2a edição: ano de referência 2002. Methodological Report 29 (2008). URL ftp://ftp.ibge.gov.br/Pib_Municipios/.
- [31] Instituto Brasileiro de Geografia e Estatística (IBGE) Coordenação de Contas Nacionais. Produto interno bruto dos municípios 3a edição: ano de referência 2010. Methodological Report 29 (2016). URL ftp://ftp.ibge.gov.br/Pib_Municipios/.

- [32] 安徽省统计局编 (Anhui Bureau of Statistics). 安徽统计年鉴 (Anhui Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [33] 北京市统计局编 (Beijing Bureau of Statistics). 北京统计年鉴 (Beijing Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [34] 重庆市统计局编 (Chongqing Bureau of Statistics). 重庆统计年鉴 (Chongqing Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [35] 福建省统计局编 (Fujian Bureau of Statistics). 福建统计年鉴 (Fujian Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [36] 甘肃省统计局编 (Gansu Bureau of Statistics). 甘肃统计年鉴 (Gansu Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [37] 广东省统计局编 (Guangdong Bureau of Statistics). 广东统计年鉴 (Guangdong Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [38] 广西壮族自治区统计局编 (Guangxi Bureau of Statistics). 广西统计年鉴 (Guangxi Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [39] 贵州省统计局编 (Guizhou Bureau of Statistics). 贵州统计年鉴 (Guizhou Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [40] 海南省统计局编 (Hainan Bureau of Statistics). 海南统计年鉴 (Hainan Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [41] 河北省统计局编 (Hebei Bureau of Statistics). 河北统计年鉴 (Hebei Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [42] 黑龙江省统计局编 (Heilongjiang Bureau of Statistics). 黑龙江统计年鉴 (Heilongjiang Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [43] 河南省统计局编 (Henan Bureau of Statistics). 河南统计年鉴 (Henan Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [44] 湖北省统计局编 (Hubei Bureau of Statistics). 湖北统计年鉴 (Hubei Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [45] 湖南省统计局编 (Hunan Bureau of Statistics). 湖南统计年鉴 (Hunan Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [46] 内蒙古自治区统计局编 (Inner Mongolia Bureau of Statistics). 内蒙古统计年鉴 (Inner Mongolia Statistical Yearbook). URL <http://tongji.cnki.net/>.

- [47] 江苏省统计局编 (Jiangsu Bureau of Statistics). 江苏统计年鉴 (Jiangsu Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [48] 江西省统计局编 (Jiangxi Bureau of Statistics). 江西统计年鉴 (Jiangxi Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [49] 吉林省统计局编 (Jilin Bureau of Statistics). 吉林统计年鉴 (Jilin Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [50] 辽宁省统计局编 (Liaoning Bureau of Statistics). 辽宁统计年鉴 (Liaoning Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [51] 宁夏回族自治区统计局编 (Ningxia Bureau of Statistics). 宁夏统计年鉴 (Ningxia Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [52] 青海省统计局编 (Qinghai Bureau of Statistics). 青海统计年鉴 (Qinghai Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [53] 山东省统计局编 (Shandong Bureau of Statistics). 山东统计年鉴 (Shandong Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [54] 上海市统计局编 (Shanghai Bureau of Statistics). 上海统计年鉴 (Shanghai Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [55] 山西省统计局编 (Shanxi Bureau of Statistics). 山西统计年鉴 (Shanxi Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [56] 四川省统计局编 (Sichuan Bureau of Statistics). 四川统计年鉴 (Sichuan Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [57] 天津市统计局编 (Tianjin Bureau of Statistics). 天津统计年鉴 (Tianjin Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [58] 新疆维吾尔自治区统计局编 (Xinjiang Uygur Bureau of Statistics). 新疆统计年鉴 (Xinjiang Uygur Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [59] 西藏自治区统计局编 (Tibet Autonomous Region Bureau of Statistics). 西藏统计年鉴 (Tibet Autonomous Region Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [60] 云南省统计局编 (Yunnan Bureau of Statistics). 云南统计年鉴 (Yunnan Statistical Yearbook). URL <http://tongji.cnki.net/>.
- [61] 浙江省统计局编 (Zhejiang Bureau of Statistics). 浙江统计年鉴 (Zhejiang Statistical Yearbook). URL <http://tongji.cnki.net/>.

- [62] Acevedo, K. M. & Bornacelly, I. CEDE Municipal Panel. CEDE Documents, Universidad de los Andes - CEDE (2014). URL <https://datoscede.uniandes.edu.co/datoscede/microdatos-detalle.php/263/panel-municipal-del-cede/>.
- [63] Departamento Administrativo Nacional de Estadística (DANE). Indicador de importancia económica municipal. URL <https://www.dane.gov.co/index.php/estadisticas-por-tema/cuentas-nacionales/cuentas-nacionales-departamentales/indicador-de-importancia-economica-municipal>.
- [64] Statistical Office of the European Communities. EUROSTAT: GDP at regional level. URL https://ec.europa.eu/eurostat/statistics-explained/index.php/GDP_at_regional_level.
- [65] World Bank Group. The Indonesia Database for Policy and Economic Research (INDO-DAPOER). URL <https://datacatalog.worldbank.org/dataset/indonesia-database-policy-and-economic-research>.
- [66] Govt. of Andhra Pradesh Planning Department. Socio economic survey 2015-16. Tech. Rep. (2016). URL http://www.indiaenvironmentportal.org.in/files/file/AP-Socio_Eco.-2015-16_Final.pdf.
- [67] Govt. of Assam Directorate of Economics and Statistics. Economic survey, assam, 2011-12. Tech. Rep. (2012). URL https://des.assam.gov.in/sites/default/files/swf_utility_folder/departments/ecostat_medhassu_in_oid_3/portlet/level_1/files/economic_survey2011-12-min.pdf.
- [68] Govt. of Bihar Department of Planning and Development Directorate of Economics and Statistics. State domestic product 2004-05 to 2010-11 (adv.) district domestic product 2004-05 to 2007-08 bihar. Tech. Rep. (2008). URL [http://bspb.bih.nic.in/Publication/DDP%20Book-2004-05%20to%202007-08\(BIHAR\).pdf](http://bspb.bih.nic.in/Publication/DDP%20Book-2004-05%20to%202007-08(BIHAR).pdf).
- [69] Strengthening state plan for human development (ssphd). Tech. Rep. (2014).
- [70] Estimates of gross/net state domestic product and district domestic product of jharkhand. Tech. Rep. (2009).
- [71] State and district domestic product of karnataka, 2009-10. Tech. Rep. (2010).
- [72] State and district domestic product of karnataka, 2010-11. Tech. Rep. (2011).
- [73] State and district domestic product of karnataka, 2014-15. Tech. Rep. (2015).
- [74] State and district domestic product of karnataka, 2015-16. Tech. Rep. (2016).
- [75] Gross domestic product of kerala and india from 2004-05 to 2012-13. Tech. Rep. (2013).

- [76] Gross domestic product of kerala and india (new series) from 2011-12 to 2015-16. Tech. Rep. (2017).
- [77] Economic survey of madhya pradesh 2013-14 (2014).
- [78] Economic survey of maharashtra 2014-15. Tech. Rep. (2015).
- [79] Economic survey of maharashtra 2016-17. Tech. Rep. (2017).
- [80] Odisha economic survey 2014-15. Tech. Rep. (2015).
- [81] District domestic product of punjab (2009-10 to 2010-11). Tech. Rep. (2011).
- [82] District domestic product of punjab (2011-12 to 2013-14). Tech. Rep. (2014).
- [83] Estimates of district domestic product of rajasthan 2004-05 to 2011-12. Tech. Rep. (2012).
- [84] Statistical handbook of tamil nadu 2017. Tech. Rep. (2017).
- [85] Estimates of district domestic product of uttarakhand 2004-05 to 2013-14. Tech. Rep. (2014).
- [86] U.S. Bureau of Economic Analysis. Gross domestic product by metropolitan area, 2017. News Release September 18 2018. URL https://www.bea.gov/system/files/2018-09/gdp_metro0918_0.pdf.
- [87] Instituto Brasileiro de Geografia e Estatística (IBGE). Regioes Geograficas 2017 (2017). URL <ftp://geoftp.ibge.gov.br/>.
- [88] All China Marketing Research Co). China 1982,1990,2000,2010 county population census data with GIS maps (2013).
- [89] Global Administrative Areas. GADM database of Global Administrative Areas, version 2.0 (2012). URL <http://www.gadm.org/>.
- [90] Statistical Office of the European Communities). NUTS GISCO 2013 (2013). URL https://ec.europa.eu/eurostat/cache/GISCO/distribution/v2/nuts/shp/NUTS_RG_01M_2013_3035.shp.zip.
- [91] ML InfoMap.
- [92] Walker, K. tigris: An R Package to Access and Work with Geographic Data from the US Census Bureau. *The R Journal* **8**, 231–242 (2016).
- [93] Law, G. *Administrative Subdivisions of Countries: A Comprehensive World Reference, 1900 Through 1998* (McFarland Publishing, 1999). URL <http://statoids.com/>.
- [94] The Central People’s Government of the People’s Republic of China. (evolution of administrative divisions of the people’s republic of china) (2006). URL http://www.gov.cn/test/2006-02/27/content_212020.htm.

- [95] Henderson, J. V., Storeygard, A. & Weil, D. N. Measuring economic growth from outer space. *American economic review* **102**, 994–1028 (2012).
- [96] NOAA's National Geophysical Data Center). Version 4 DMSP-OLS Nighttime Lights Time Series (2016). URL <https://ngdc.noaa.gov/eog/dmsp/downloadV4composites.html>.
- [97] Rohde, R., Muller, R., Jacobsen, R., Perlmutter, S. & Mosher, S. Berkeley earth temperature averaging process. *Geoinformatics & Geostatistics: An Overview* **01** (2013).
- [98] Dee, D. P. *et al.* The era-interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the royal meteorological society* **137**, 553–597 (2011).
- [99] Harris, I., Jones, P. D., Osborn, T. J. & Lister, D. H. Updated high-resolution grids of monthly climatic observations—the cru ts3. 10 dataset. *International journal of climatology* **34**, 623–642 (2014).