

Supplementary Materials

This Supplementary Materials Appendix contains fourteen sections, with additional treatment of the following topics:

Framework

- A.1 Aggregation across productive units (Equation 1)
- A.2 Deriving Figure 1
- A.3 Growth effects

Data and methods

- B.1 Data
- B.2 Empirical approach

Robustness and generalizability of the main result

- C.1 Robustness to model specification, samples, and sources of data
- C.2 Growth versus level effects
- C.3 Studying heterogeneous responses
- C.4 Comparison with Dell Jones Olken 2012

Climate projections and their uncertainty

- D.1 Building impact projections
- D.2 Projected impacts and robustness to alternative specifications
- D.3 Constructing the damage function
- D.4 Shape of the damage function
- D.5 Damage function uncertainty

A Framework

A.1 Aggregation across productive units (Equation 1)

Numerous micro-econometric studies have identified the effects of momentary temperature on small units of analysis, such as the effect of variation in hourly temperature on county-level agricultural yields^{1,2} or the effect of variation in daily temperatures on the productivity of individual workers^{3,4}. In these cases, the effect of temperature on these scales is highly nonlinear (Figure 1), which has led some researchers to suggest that the macroeconomy should exhibit similarly nonlinear responses^{5,6}. However, the units analyzed in micro-econometric studies are vanishingly small, in terms of economic value, relative to the scale of a macroeconomy. For example, the output of an individual worker is negligible relative to the output of an entire country. Further, the periods of time analyzed in micro-econometric studies are typically short, ranging from an hour to a day in most cases and sometimes weeks or months—periods that are much shorter than the annual timescale over which macroeconomic data are typically aggregated*. Here we use a simple model to consider how highly nonlinear changes

*In some cases,^{1,7} prior micro studies recover the net effect of a short-lived (e.g. daily) temperature event using annual data. In these cases, the “daily response” is actually the cumulative response to a daily event, integrated over a period substantially longer than a day. This approach is particularly important in cases where short-lived events have

in productivity on short time scales and across many small elements in the economy might aggregate and be reflected in macroeconomic responses over longer periods of time.

We partition a macroeconomy into “industries” indexed by i , with all individual units of production within each industry assumed to respond identically to temperature. Each industry could thus be highly specific—for example, one industry could be all maize farms using a given technology to produce a specific variety of maize. Production in each industry occurs at numerous small locations in space, indexed by ℓ ; countries, indexed by \mathcal{L} , are large collections of locations. The incremental moments in time that micro studies have analyzed (e.g. hours) are indexed by t and longer periods of time composed of many sequential moments (e.g. years) are indexed by τ .

We follow the notation of Deryugina and Hsiang (2014)⁷ and describe capital K_i and labor L_i in each industry as having respective productivities A_i^K and A_i^L that are functions of instantaneous temperature $T_{\ell t}$ experienced at a location ℓ and time t . The total quantity of capital and labor allocated to industry i could also potentially change with temperature. The price of a unit of output is p and α is a constant in this stylized production function. For a subunit of the economy at a location ℓ at time t using technologies described by i , total production $Y_{i\ell t}$ is then

$$Y_{i\ell t}(T_{\ell t}) = p_i (A_i^K(T_{\ell t})K_{i\ell t}(T_{\ell t}))^\alpha (A_i^L(T_{\ell t})L_{i\ell t}(T_{\ell t}))^{1-\alpha}. \quad (2)$$

For simplicity, we assume that capital and labor are not rapidly reallocated across locations in response to temperature changes[†]. Changes in the total allocation of time individuals allocate to labor is known to change with temperature³, however this response can be easily described by changes to labor productivity A_i^L since we observe empirically that labor is not reallocated across different industries in response to temperature^{3,7}. Note that in a competitive equilibrium $\frac{K_{i\ell t}}{L_{i\ell t}} = \frac{\alpha}{1-\alpha}$, such that capital labor ratios are fixed and output scales linearly with the total quantity of capital and labor allocated to i (constant returns to scale).

For notational convenience, we define $U_{i\ell t} = p_i K_{i\ell t}^\alpha L_{i\ell t}^{1-\alpha}$ as a scalar measure of resources applied to i at location ℓ at time t . We think of U_i as describing the number of modular units of production allocated to industry i (e.g. firms). These assumptions and notations allow us to simplify Equation 2 to

$$\begin{aligned} Y_{i\ell t}(T_{\ell t}) &= \underbrace{(A_i^K(T_{\ell t})^\alpha A_i^L(T_{\ell t})^{1-\alpha})}_{f_i(T_{\ell t})} p_i K_{i\ell t}^\alpha L_{i\ell t}^{1-\alpha} \\ &= f_i(T_{\ell t}) U_{i\ell t} \end{aligned} \quad (3)$$

where $f_i(T_{\ell t})$ is a function describing how overall productivity in industry i responds to instantaneous temperatures. Note that for clarity, we have assumed here that the economy is additively separable across industries and locations, with firms behaving as atomistic producers. It is likely that large scale climatic changes generate emergent impacts on firms beyond what an atomistic firm might experience in response to an isolated change in their individual climate exposure, since cross-firm spillovers might be substantial and novel price responses might emerge when climatic events are correlated in either

delayed impacts^{8,9}.

[†]This assumption is consistent with findings in Deryugina and Hsiang (2014) from the United States.

time or space. For example, climate-induced interruptions in a firm's supply chain might amplify the economic impact of a climatic event that the firm is itself exposed to. If these effects are substantial and cross national boundaries, then our empirical approach below is likely to understate the overall economic impact of large scale climatic changes since we focus on country-level changes as if they occur in isolation.

To form a measure of aggregate output, such as Gross Domestic Product (GDP), we must sum across all industries i and integrate production across all locations in a country and all moments in time within the period of observation. Thus total output in country \mathcal{L} during year τ is then:

$$\begin{aligned} Y_{\mathcal{L}\tau} &= \sum_i Y_{i\mathcal{L}\tau} \\ &= \sum_i \int_{t \in \tau} \int_{\ell \in \mathcal{L}} f_i(T_{\ell t}) U_{i\ell t} d\ell dt. \end{aligned} \quad (4)$$

The spatial and temporal distribution of units $U_{i\ell t}$, as well as the spatial distribution of atmospheric temperatures, will determine what temperatures $T_{\ell t}$ individual units are exposed to. Within country \mathcal{L} and period τ , we can integrate the number of points in time when individual productive units are exposed to a momentary local temperature $T_{i\ell t}$ to construct a marginal distribution function[‡] summarizing temperature exposure within industry i . Let the shape of this marginal distribution function be described by $g_i(\cdot)$ which is mean zero and can be shifted by the location parameter $\bar{T}_{\mathcal{L}\tau}$, defined as average temperature in country \mathcal{L} during period τ . Thus $g_i(T - \bar{T}_{\mathcal{L}\tau})$ looks like a histogram of the temperatures that units U_i are exposed to within a large region and interval of time. For simplicity, here we assume $g_i(\cdot)$ does not change in shape across countries or years, although the location parameter $\bar{T}_{\mathcal{L}\tau}$ may change. In the real world, the shape of $g_i(\cdot)$ may change based on changes in the within-country and within-year distribution of temperatures that productive units are exposed to.

We note that $g_i(\cdot)$ has two important properties. First, for a single industry, the total quantity or “mass” of productive units M_i is the integral of $g_i(\cdot)$ over all possible temperatures

$$M_i = \int_{-\infty}^{\infty} g_i(T - \bar{T}_{\mathcal{L}\tau}) dT = \int_{t \in \tau} \int_{\ell \in \mathcal{L}} U_{i\ell t} d\ell dt. \quad (5)$$

Second, the shape of $g_i(\cdot)$ reflects the distribution of productive units across space and time such that

$$\int_{-\infty}^x g_i(T - \bar{T}_{\mathcal{L}\tau}) dT = \int_{t \in \tau} \int_{\ell \in \mathcal{L}} U_{i\ell t} \mathbf{1}[T_{\ell t} < x] d\ell dt \quad (6)$$

for $x \in (-\infty, \infty)$.

We can now write total production at the aggregate level in terms of average temperature $\bar{T}_{\mathcal{L}\tau}$,

[‡]Note that this marginal distribution is not a *marginal probability distribution* because the total number of units at each temperature are not normalized by the total number of units. i.e. this marginal distribution is more analogous to a histogram measuring frequencies rather than a histogram measuring probabilities.

measured at the aggregate level, and $g_i(\cdot)$

$$\begin{aligned} Y(\bar{T}_{\mathcal{L}\tau}) &= \sum_i Y_i(\bar{T}_{\mathcal{L}\tau}) \\ &= \sum_i \int_{t \in \tau} \int_{\ell \in \mathcal{L}} f_i(T_{\ell t}) U_{i\ell t} d\ell dt \\ &= \sum_i \int_{-\infty}^{\infty} f_i(T) g_i(T - \bar{T}_{\mathcal{L}\tau}) dT \end{aligned} \quad (7)$$

which no longer requires we know detailed information on the spatial and temporal distribution of $U_{i\ell t}$. If the shape of $g_i(\cdot)$ is relatively unchanged across periods τ , then $\bar{T}_{\mathcal{L}\tau}$ is a sufficient statistic for temperature exposure at the aggregate level. As shown in Figure 1, changing annual average temperature $\bar{T}_{\mathcal{L}\tau}$ shifts the distribution of temperature exposure for individual micro-level units.

Essentially, we have changed variables by collapsing the joint spatial and temporal distribution of temperatures and micro-level productive units into the marginal distribution $g_i(\cdot)$ and a location parameter $\bar{T}_{\mathcal{L}\tau}$, which is a country's annual average temperature.

A.2 Deriving Figure 1

Figure 1 depicts how the application of Equation 7 (i.e. Equation 1 in the main text) to previously derived micro-level response functions generates a macro-level response. Prior work^{1–5} has shown that basic units of the economy, such as crops and labor, have a response to momentary temperature that is highly nonlinear and well-approximated by a piecewise-linear function similar to Figure 1d (recall Figure 1a-c). In general, the productivity of basic units in the economy is either flat or slightly increasing at lower temperatures, and then declines steeply with temperature above a critical temperature threshold. These responses are the function $f_i(T)$ in Equation 7. Thus, to develop a sense of how macro-level responses to temperature should look, we assume the micro-level function $f_i(\cdot)$ is piecewise linear with kink at the critical instantaneous temperature $T = \tilde{T}$:

$$f(T) = \begin{cases} c_1 + b_1 T & \text{if } T < \tilde{T} \\ c_2 + b_2 T & \text{if } T \geq \tilde{T} \end{cases} \quad (8)$$

where where slope terms b_1 and b_2 and intercept terms c_1 and c_2 satisfy

$$\begin{aligned} c_1 + b_1 \tilde{T} &= c_2 + b_2 \tilde{T} \\ b_1 > 0, \quad b_2 < 0, \quad -b_2 > b_1. \end{aligned} \quad (9)$$

These conditions ensure $f_i(\cdot)$ is continuous and non-differentiable due to a “kink” at the critical temperature \tilde{T} , with a downward slope above \tilde{T} that is steeper than the upward slope below \tilde{T} . Figure 1d displays these properties.

For ease of comparability across countries and industries of different economic size, we normalize total production in an industry by the the total mass of productive units M_i and focus on $\frac{Y_i}{M_i}$. When examining data, we implement an analogous normalization by focusing on GDP per capita[§]. We are

[§]The standard GDP per capita normalization does not account for capital, however we note that, at least in our

interested in how aggregate productivity changes with each country's average annual temperature, so we differentiate this normalized measure $\frac{Y_i}{M_i}$ with respect to $\bar{T}_{\mathcal{L}\tau}$ while substituting from Equation 7

$$\begin{aligned} \frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left(\frac{Y_i}{M_i} \right) &= \frac{1}{M_i} \frac{\partial Y_i}{\partial \bar{T}_{\mathcal{L}\tau}} \\ &= \frac{1}{M_i} \frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left[\int_{-\infty}^{\infty} f_i(T) g_i(T - \bar{T}_{\mathcal{L}\tau}) dT \right] \\ &= \frac{1}{M_i} \frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left[\int_{-\infty}^{\tilde{T}} f_i(T) g_i(T - \bar{T}_{\mathcal{L}\tau}) dT + \int_{\tilde{T}}^{\infty} f_i(T) g_i(T - \bar{T}_{\mathcal{L}\tau}) dT \right] \end{aligned} \quad (10)$$

where the final equality is simply separating the integral into a portion below the critical temperature and a portion above the critical temperature. These integrals appear difficult to differentiate because the shape of $g_i(\cdot)$ is unknown, but a change of variables clarifies the derivative by making the shift parameter $\bar{T}_{\mathcal{L}\tau}$ an argument of $f_i(\cdot)$ instead. Define a new variable $T' = T - \bar{T}_{\mathcal{L}\tau}$. We analogously define \tilde{T}' such that $\tilde{T} = \tilde{T}' + \bar{T}_{\mathcal{L}\tau}$. Substituting T' and \tilde{T}' into Equation 10 and noting the linearity of $f_i(\cdot)$ within the range of each integral we have

$$\begin{aligned} \frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left(\frac{Y_i}{M_i} \right) &= \frac{1}{M_i} \frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left[\int_{-\infty}^{\tilde{T}'+\bar{T}_{\mathcal{L}\tau}} f_i(T' + \bar{T}_{\mathcal{L}\tau}) g_i(T') dT' + \int_{\tilde{T}'+\bar{T}_{\mathcal{L}\tau}}^{\infty} f_i(T' + \bar{T}_{\mathcal{L}\tau}) g_i(T') dT' \right] \\ &= \frac{1}{M_i} \left[\int_{-\infty}^{\tilde{T}'+\bar{T}_{\mathcal{L}\tau}} b_1 g_i(T') dT' + \int_{\tilde{T}'+\bar{T}_{\mathcal{L}\tau}}^{\infty} b_2 g_i(T') dT' \right] \\ &= b_1 m_{i1}(\bar{T}_{\mathcal{L}\tau}) + b_2 m_{i2}(\bar{T}_{\mathcal{L}\tau}) \end{aligned} \quad (11)$$

where $m_i(\bar{T}_{\mathcal{L}\tau})$ is defined as the fraction of productive unit-hours exposed to T below \tilde{T} for a given national mean temperature $\bar{T}_{\mathcal{L}\tau}$. Specifically,

$$\begin{aligned} m_{i1}(\bar{T}_{\mathcal{L}\tau}) &= \frac{\int_{-\infty}^{\tilde{T}'+\bar{T}_{\mathcal{L}\tau}} g_i(T') dT'}{M_i} = \frac{\int_{-\infty}^{\tilde{T}} g_i(T - \bar{T}_{\mathcal{L}\tau}) dT}{\int_{-\infty}^{\infty} g_i(T - \bar{T}_{\mathcal{L}\tau}) dT} \\ m_{i2}(\bar{T}_{\mathcal{L}\tau}) &= 1 - m_{i1}(\bar{T}_{\mathcal{L}\tau}) \end{aligned}$$

and m_1 and m_2 are illustrated as two shaded masses in Figure 1e.

Thus, the country level aggregated response $\frac{Y_i}{M_i}$ can be computed by integrating the weighted average in Equation 11 with respect to average annual temperature $\bar{T}_{\mathcal{L}\tau}$, where the weights m_{i1} and m_{i2} are determined by the shape of the distribution of productive units $g_i(\cdot)$ and its location relative to the critical temperature \tilde{T} . This integration is depicted in Figure 1f. The aggregated productivity function recovered from this integration will be smooth if $g_i(\cdot)$ is continuous and contains no point-masses, even though there is a known sharp kink in the micro-level response function $f_i(\cdot)$. The distribution $g_i(\cdot)$ "smoothes" over this kink, causing the the effect of average warming (increasing $\bar{T}_{\mathcal{L}\tau}$) to be much less abrupt on the macro-economy than local and instantaneous warming is on micro-level productive units. A broader distribution function $g_i(\cdot)$, either caused by a wider dispersion of

simple model, capital-labor ratios are constant in equilibrium so this normalization would still be a valid approach to comparing countries of different size.

productive assets across space and/or longer periods of observation, will cause more smoothing.

A key characteristic of interest regarding the shape of $\frac{Y_i}{M_i}$ is the temperature at which the peak of the function is located. This turning point occurs at the temperature where the derivative of the response is zero:

$$\frac{\partial}{\partial \bar{T}_{\mathcal{L}\tau}} \left(\frac{Y_i}{M_i} \right) = b_1 m_{i1}(\bar{T}_{\mathcal{L}\tau}) + b_2 m_{i2}(\bar{T}_{\mathcal{L}\tau}) = 0$$

which implies that the turning point temperature $\bar{T}_{\mathcal{L}\tau}^*$ has the property:

$$\frac{m_{i2}(\bar{T}_{\mathcal{L}\tau}^*)}{m_{i1}(\bar{T}_{\mathcal{L}\tau}^*)} = \frac{b_1}{-b_2} < 1$$

where the inequality comes from the conditions in Equation 9. This implies

$$\bar{T}_{\mathcal{L}\tau}^* < \tilde{T}$$

if $g_i(\cdot)$ is roughly symmetric or negatively skewed.[¶] Note that the greater the difference between $|b_1|$ and $|b_2|$, the lower the temperature at which the peak occurs. Also, the greater the dispersion in $g_i(\cdot)$, the lower the value of $\bar{T}_{\mathcal{L}\tau}^*$.

For multiple sectors, total production $\frac{Y}{M} = \sum_i \frac{Y_i}{M_i}$ will be concave in annual average temperature because it is the weighted sum of several concave functions of annual average temperature.

We note that this result reconciles a long-standing debate about whether degree days or seasonal averages are better measures of exposure to climate change,¹⁰ since we have shown that the two measures are closely related mathematically (and identical under certain assumptions). The kinked micro-level response reported by degree day studies^{1,11} has a direct mapping (Equation 7) to the seasonal average response that is smoother and roughly quadratic.¹⁰ The quality of fit may differ between these two approaches depending on spatial and temporal autocorrelation of the outcome variable within averaging periods and regions, as well as the extent to which $g_i(\cdot)$ is actually the same across different units of analysis. Nonetheless, at their core the two approaches are not fundamentally different and may produce findings that are mutually consistent even though the temperature response functions recovered by the two approaches differ.

A.3 Growth effects

Is it plausible that temporary productivity losses caused by temperature can translate into enduring productivity shocks? That is, could we expect that economic *growth* is affected by temporary productivity changes? The micro-level responses to temperature in prior analyses (discussed above) generally characterize temporary changes in productivity. However, as discussed in detail below, our main analysis estimates the nonlinear effect of temperature on GDP per capita growth, in part because log GDP per capita is known to have a unit root and thus requires first differencing for proper inference.¹² As detailed in Section C.2 below, this transformation of the dependent variable does not

[¶]This condition may not hold if $g_i(\cdot)$ is strongly positively skewed, depending on the values b_1 and b_2 .

itself imply that output is depressed relative to trend in the long run; getting at this question requires examining the cumulative effect of lagged independent variables, which we implement in Section C.2. (As stated in the main text and detailed in Section C.2, our results are somewhat ambiguous as to whether the nonlinear effect of temperature generates a permanent or temporary loss of output relative to trend.) In this section we point out how, in a simple model, temporary productivity losses generated through a mechanism like Equation 2 could generate growth effects if savings rates do not change with temperature to compensate and offset temporary productivity losses.

Let total productivity be smooth, twice differentiable, and concave with respect to \bar{T} , as derived above in Section A.1-A.2 (we drop the $i\mathcal{L}\tau$ subscripts here). If year-to-year changes in capital stocks are modest, then we can apply Taylor's theorem and linearize output with respect to total productive units M , with local slope $\psi(\bar{T})\gamma$:

$$Y = \psi(\bar{T})\gamma M \quad (12)$$

where $\frac{\partial^2}{\partial T^2} < 0$. Thus, the temporary change in temperature affects output similarly to a change in total factor productivity, akin to the approach in Nordhaus and Boyer (2000)¹³.

Following seminal work by Solow (1956)¹⁴, we assume a fraction δ of M depreciates each period and a fraction s of output is saved and re-invested in augmenting M . The equation of motion for M is then

$$\frac{\partial M}{\partial t} = sY - \delta M$$

where M is measured in units of U , which recall accounts for both capital and labor contributions to production. Substituting from Equation 12:

$$\begin{aligned} \frac{\partial M}{\partial t} &= s\psi(\bar{T})\gamma M - \delta M \\ &= \underbrace{(s\psi(\bar{T})\gamma - \delta)}_{\text{net growth if } > 0} M. \end{aligned}$$

Thus the productive stock of a country M will grow on net if savings rates, average temperature, and the initial stock M (which affects γ) allow investment to outpace depreciation.^{||}

Importantly, if productivity is reduced in a given period due to temperature, this has an impact on output in subsequent periods because it reduces investment in M . To see this, note that for a given initial stock M_{t-1} just prior to a temperature realization \bar{T}_{t-1} , then output in the subsequent period will be

$$\begin{aligned} Y_t &= \psi(\bar{T}_t)\gamma M_t \\ &= \psi(\bar{T}_t)\gamma (M_{t-1} + \Delta M_{t-1 \rightarrow t}) \\ &= \psi(\bar{T}_t)\gamma (M_{t-1} + s\psi(\bar{T}_{t-1})\gamma M_{t-1} - \delta M_{t-1}). \end{aligned} \quad (13)$$

Differentiating Equation 13 with respect to the prior year's temperature \bar{T}_{t-1} , we see that current

^{||}Here we assume that the climate does not affect depreciation, although recent evidence suggest this may be an important direction for future work.^{15,16}

output is influenced by temperature changes in the prior year

$$\frac{\partial Y_t}{\partial T_{t-1}} = s\psi(\bar{T}_t)\gamma^2 M_{t-1} \left(\frac{\partial}{\partial T} \Big|_{\bar{T}_{t-1}} \right) \quad (14)$$

because output is reduced through lowered reinvestment during a less productive prior period. Here, the savings rate s is assumed fixed as in the standard Solow model, although relaxing this assumption might alter this result. While effects may diminish with time if savings compensate for productivity losses,** i.e. $\frac{\partial s}{\partial Y} < 0$, empirically it has been found that individuals tend to prefer to smooth consumption $(1-s)Y$ and allow savings to fluctuate in response to climate-related productivity shocks^{18–20}, i.e. $\frac{\partial s}{\partial Y} > 0$. Nevertheless, it is possible that some amount of compensatory savings may occur in later periods causing production to eventually return to trend.²¹ Whether this occurs is ultimately an empirical question, which we consider in Section C.2 below.

B Data and methods

B.1 Data

Our main source of data on per capita GDP is the World Bank's World Development Indicators²², which cover the years 1960–2012 for all countries in the world, although data for only a subset of years are available for some countries. To study effects in agricultural and non-agricultural sectors, we use World Bank data on value added for different sectors, also available for most countries over the same period. As robustness we re-estimate our main results with income data from the Penn World Tables version 8.0²³. Our source for temperature and precipitation data is the University of Delaware reconstruction assembled by Matsuura and Willmot²⁴, which contains 0.5 degree gridded monthly average temperature and total precipitation data for all land areas over the period 1900–2010, as interpolated from station data. We aggregate the 0.5 degree grid cell estimates to the country-year level, weighting by population density in the year 2000 using data from the Gridded Population of the World²⁵. Our full dataset contains 6584 country-year observations between the years 1960–2010.

B.2 Empirical approach

Using a 51-year longitudinal sample of countries around the world, we take first differences of the natural log of annual real (inflation-adjusted) gross domestic product per capita Y . These first differences (ΔY) can be interpreted as per-period growth rates in income. We deconvolve the factors that might affect these changes in income with a simple and general model:

$$\Delta Y_{it} = h(T_{it}) + \lambda_1 P_{it} + \lambda_2 P_{it}^2 + \mu_i + \nu_t + \theta_1 t + \theta_2 t^2 + \varepsilon_{it} \quad (15)$$

where countries are indexed by i and years by t . Note that these definitions, used through the remainder of the SOM, differ from the definitions for i and t used in Sections A.1–A.3. All time-invariant factors that influence countries' average growth rates, such as history, culture or topography, are accounted

**For a related model of non-savings adaptive compensation, see Dell et al (2009)¹⁷.

for by the country-specific constant terms μ_i (fixed effects). Abrupt global events, such as shocks to energy markets or global recessions, are captured by the year fixed effects ν_t . Gradual changes to individual countries' growth rates that may be driven by slowly changing factors within a country, such as demographic shifts, trade liberalization, and evolving political institutions, are accounted for by the flexible country-specific time trends $\theta_{i1}t + \theta_{i2}t^2$. Because our dependent variable is the derivative of income, quadratic country-specific time trends permit growth rates to evolve nonlinearly over time, allowing us to account for country-specific cubic polynomials in income levels (by integrating first differences).^{††}

In this framework, each country is allowed its own level and nonlinear trend in growth, and the impact of temperature on growth is identified from within-country deviations from this trend. Controlling for trends and convergence in incomes²¹ using country-specific trends outperforms auto-regressive models²⁶. We explicitly account for the effect of *precipitation* P and *precipitation-squared* P^2 because idiosyncratic changes in local annual temperatures tend to be correlated with changes in precipitation²⁷.

Importantly, our technique to controlling for both time-invariant and time-varying influences is more reliable than only controlling for observed variables (e.g. regression on explicit covariates, such as demographic or political variables) because it flexibly accounts for both observed and unobserved controls, it is robust to mismeasurement of controls, and it allows these controls to differentially influence different countries²⁸. For example, even without explicitly modeling the effect of demographic trends, our model accounts for the fact that nonlinear demographic trends may be different in different countries, with measurement errors that differ between countries, and the *effect* of demographic trends on income may also differ across countries. Furthermore, because many traditional “control” variables are themselves likely affected by climatic events, including them may generate bias in our estimates of interest, an issue known as “bad control”²⁹. This issue is discussed in detail by ref [28].

Our focus is the potentially non-linear relationship between annual temperatures T_{it} and income growth described by $h(T_{it})$. We begin by estimating $h(T_{it})$ as a simple quadratic (i.e. $h(T_{it}) = \beta_1 T_{it} + \beta_2 T_{it}^2$), exploring more flexible functional forms such as higher order polynomials and restricted cubic splines for robustness.

Although the motivating micro evidence presented in Figure 1a-c focuses on the effect of temperature on the *level* of output, we explicitly model the effect of temperature on *growth* because measures of GDP within a country exhibit such high levels of serial correlation ($\rho = 0.999$) that they are indistinguishable from a random walk, i.e. they have a unit root. Because unit root processes are non-stationary, regressions using unit root outcomes often generate spurious results and traditional test statistics fail.³⁰ Accounting for country-specific trends in levels does not alleviate this concern. First differencing income, i.e. using growth as an outcome, is the standard approach in this context.³¹ After first differencing and accounting for year effects as well as country-specific quadratic trends in growth, serial correlation in the outcome is much less problematic ($\rho = 0.125$). Because some serial correlation persists in the outcome, even after first differencing, we non-parametrically adjust our standard error estimates to account for arbitrary patterns of autocorrelation between residuals within each

^{††}Results look similar if we include only a linear time trend $\theta_i t$, but an F-test strongly rejects the null that the country-specific quadratic time trends are jointly zero ($p < 0.01$), and so we retain the quadratic trends in our main specification.

country.³²

If temperature temporarily affects productivity, as the micro evidence presented in Figure 1a-c suggests, this will still appear as a change in the current growth rate as estimated in Equation 15. It is also possible that temperature affects the growth rate in economies that exhibit unit-root-like behavior, such that these effects permanently alter income trajectories. This could occur for (at least) two important reasons. First, as discussed in section A.3, level effects in one period can affect output in the next period by affecting the growth rate of the capital stock, an effect that could be amplified if productive capital must be diverted to invest in costly adaptation measures³³. Second, changes in temperature could directly affect the rate of technological change (the basis of growth in standard economic growth models), for instance if warmer temperatures adversely affect the cognition needed for innovation.^{4,34}

If temperature permanently affects income through growth effects, this has larger consequences for the projected long run effects of climate change, as small changes in annual growth rates cumulated over a long period can have large effects on total GDP and living standards^{15,34}. Thus determining the structure of this temperature effect (growth vs. levels) is important for optimal policy.¹⁵ Following ref [35], we introduce lagged independent variables to Equation 15 to examine this issue explicitly (see Section C.2 below), finding evidence for a mixture of growth and level effects consistent with prior work. Notably, the qualitative nature of our mean climate projections do not change under these alternative models, in large part because our mean estimates are consistent with a growth effect that permanently affect income.

It is often wrongly claimed that the empirical approach in Equation 15, by using variation in annual temperatures over time, implicitly assumes that economies cannot adapt. In a model with country fixed effects and a linear temperature term, it is true that all identifying variation comes from within-country variation in temperature over time (what is commonly considered “weather”), variation that is potentially hard for economic agents to anticipate and adapt to. However, in a fixed-effects model with higher-order temperature terms, both within-country and cross-country variation are used to identify the effects of temperature. Countries with higher average temperature are permitted to have a different response to within-country temperature changes. Thus, using both these sources of variation implicitly allows for more historical adaptation to longer-run climate, although the short-run changes in temperature that affect output remain unanticipated. In our projections that utilize these estimates, as a given country warms to a new average temperature, the effect of additional warming is allowed to change in a way that reflects how other countries have been observed to respond at that temperature. For an extended treatment of this important methodological issue, see ref [36], and for a clear short treatment see the SOM of ref [37].

To more clearly demonstrate how the global nonlinear response is identified, Figure ED1 shows how individual response functions for selected countries at different points in the global temperature distribution are aggregated in the global analysis to construct the global non-linear response function. Variation within individual countries identifies the local derivative of $h(T_i)$ in the neighborhood of each country's average temperature. The integral of these locally estimated derivatives is our estimate for $h(T)$.

A critical assumption of this approach is that there exists a global function $h(\cdot)$ on which all

countries lie. We test this assumption by examining subsamples of the data in the main text (Figure 2) and find no evidence that $h(\cdot)$ is dramatically different across various subsamples. Additional comparisons across subsamples are presented below.

Another approach to examining whether $h(\cdot)$ is globally generalizable is to ask whether local slope described by a single country (as depicted in Figure ED1) appears to not lie tangent to the global estimate for $h(\cdot)$. We test for this in Figure ED1h by estimating the local marginal effect $\frac{\partial Y_i}{\partial T_i}$ for each country's linear time series (the local slope). We do this by running a growth regression separately for each country:

$$Y_{it} = \alpha_i + \beta_i T_{it} + \lambda_i P_{it} + \theta_i t + \theta_{i2} t^2 + \varepsilon_{it} \quad (16)$$

and plotting the estimated β_i 's (with confidence intervals) as a function of each country's mean temperature. We then compare these estimated marginal effects, on which we impose no functional form with respect to a country's average temperature, with the derivative of the global response function estimated by Equation 15 shown by the black line. As expected, individual country-specific marginal effect estimates from Equation 16 are noisy because each uses less than 1% of our data, but they exhibit a clear downward sloping relationship with temperature (indicating a nonlinear and concave $h(\cdot)$) and only in 9% of cases (15 out of 166) does the marginal effect of the global response lie outside of the 95% confidence interval for a country's estimated β_i . (Notably roughly half of the cases for which this occurs are major oil-producing countries.) This is fully consistent with our estimated sampling error and provides no evidence against the hypothesis that a global $h(\cdot)$ describes a generalizable response.

A third and related approach to understanding whether the non-linear response observed in Figure 2a is globally generalizable, or whether it is just a composite effect of a negative and linear response in poor, hot countries and no response in rich, cooler countries, is to allow temperature to enter linearly in the regression and then to interact it with both country average temperature as well as country average income. That is, $h(\cdot)$ would be reformulated in the following way:

$$h(T_{it}) = \beta_1 T_{it} + \beta_2 (T_{it} \cdot \bar{T}_i) + \beta_3 (T_{it} \cdot \bar{Y}_i) \quad (17)$$

where \bar{T}_i and \bar{Y}_i are the average temperature and (log) average GDP/capita in country i , respectively. In the absence of the third term (the interaction with country average income), a non-linear and concave temperature response similar to Figure 2a would be indicated by $\beta_1 > 0$ and $\beta_2 < 0$. However, if this differential response was actually being driven by income – and in particular by the fact that poor countries respond differently to temperature and that hot countries tend to be poor – then the inclusion of the $\beta_3 (T_{it} \cdot \bar{Y}_i)$ term should mean that $\beta_2 = 0$, with $\beta_1 < 0$ and $\beta_3 > 0$ (that is, negative effects of temperature at low income levels that attenuate as incomes rise).

Results of re-estimating our main specification but substituting the interactions in Equation 17 for the quadratic are shown in Figure ED1h and Table S1 (with precipitation deviations similarly interacted with country-average precipitation and country-average GDP per capita). We again find strong evidence of a global non-linear temperature response, with $\beta_1 > 0$ and $\beta_2 < 0$ as expected and each statistically significant (as shown in first two rows of Table S1), and we find strong evidence that this non-linear response is driven by average temperature rather than differences in income: introducing the temperature-income interaction does not change how the temperature response varies

with average temperature. The estimated coefficients on the temperature-income interaction are very small, providing further evidence that the non-linearity we describe is due to differences in average temperature rather than income.

Figure ED1h plots the estimated marginal effect of temperature at different points in the distribution of average temperatures, as estimated from regressions with and without the temperature-income interaction included. In both cases, positive marginal effects at low temperatures and negative marginal effects at higher temperatures indicate a globally non-linear response, with point estimates nearly identical whether or not the temperature-income interaction is included (albeit marginally noisier at lower temperatures when income is included).

Supplementary Table S1: **Regression estimates corresponding to Equation 17.** All models include precipitation and precipitation interacted with average precipitation, country fixed effects, quadratic country time trends, and year fixed effects as indicated in the bottom row, with errors clustered at the country level. Temperature is measured in °C, with \bar{T}_i representing the average temperature in country i over the period. Columns 2-3 include temperature and precipitation interacted with country average GDP/capita over the study period (\bar{Y}_i), and column 4 temperature and precipitation interacted with country-average GDP/capita in logs. In both cases we de-mean the income measure such that the temperature effects in the first two rows of the table can be interpreted as the effect evaluated at global average income. Column 3 is the same as Column 2 but substitutes continent-by-year fixed effects for the year fixed effects. The marginal effects plotted in Figure ED1i correspond to the coefficient estimates in the first two rows of columns 1 and 2. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

	(1)	(2)	(3)	(4)
T_{it}	0.0126*** (0.0037)	0.0121** (0.0060)	0.0129** (0.0059)	0.0119* (0.0071)
$T_{it} * \bar{T}_i$	-0.0010*** (0.0002)	-0.0009*** (0.0003)	-0.0009*** (0.0003)	-0.0009** (0.0004)
$T_{it} * \bar{Y}_i$		0.0000 (0.0002)	0.0001 (0.0002)	
$T_{it} * \log(\bar{Y}_i)$				0.0003 (0.0022)
Observations	6584	6584	6584	6584
R squared	0.285	0.286	0.367	0.286
Fixed effects	year	year	cont-by-year	year

C Robustness and generalizability of the main result

C.1 Robustness to model specification, samples, and sources of data

Our main result that the response of growth to temperature is nonlinear and concave with an optimum near 13°C. Here we show that this result is robust to data sample, model specification, and source of data.

Alternative samples Table ED1 provides our benchmark regression estimates of Equation 15 (column 1) and shows how these estimates changes under alternative samples. In column 2 we restrict

the sample to countries with at least 20 years of growth data (as in ref^[35]), in column 3 we drop oil-rich countries (defined as countries where at least 20% of GDP comes from oil; these countries tend to be both rich and hot), in column 4 we drop the US and China from the sample. Removing oil countries has limited impact on the estimated response of total GDP and agricultural GDP, but it alters the estimated effect of temperature on non-agricultural GDP such that the response to high temperatures is substantially steeper (~40%). This suggests that oil-production is likely less sensitive to temperature than other non-agricultural components of these economies.

Specifying controls Columns 4-10 of Table ED1 demonstrate our main result is robust across models that use alternative set of controls, relative to our benchmark model. Column 5 replaces year fixed effects with more flexible continent-by-year fixed effects (to account for continent-specific shocks in a given year that could be correlated with both temperature and growth, such as ENSO³⁸). Column 6 is the same but drops country-specific trends from the model. Column 7 is the benchmark model but drops year fixed effects entirely. Columns 8-10 are the benchmark model plus 1, 3, and 5 lags of the dependent variable (growth) in the model, as is sometimes done in the macroeconomic literature to account for potential time-varying omitted variables, such as partially durable capital investments^{39,40}.

Data source Column 11 estimates our benchmark model using growth data from the Penn World Tables⁴¹ instead of the World Bank's World Development Indicators.

In all cases, the estimated coefficients change little, remaining consistent in size and sign, and are highly statistically significant across specifications. The estimated optimal temperature (bottom row) varies slightly more because it is a nonlinear combination of two uncertain parameters ($\frac{-\beta_1}{2\beta_2}$); nonetheless, it always remains well below the 20-30°C range, as predicted by Figure 1. In general, lower estimates for the optimum generate larger projected losses from climate change because a lower optimum indicates that wealthier, cooler countries are more negatively affected by warming. As quantified below, this suggests that damage estimates using our baseline model are, if anything, conservative.

Specifying functional-form of the temperature response In our main text, we model the effect of temperature on growth using a quadratic polynomial in temperature (i.e. in Equation 15, $h(T) = \beta_1 T_{it} + \beta_2 T_{it}^2$). We use this as our benchmark because it is the most parsimonious form that accurately describes the dominant pattern in the data. Other researchers sometimes prefer restricted cubic splines because of their statistical properties,¹ but these models do not perform better in this context and their analytical complexity makes it more difficult for other researchers to utilize results of that form (e.g. as in the exercise undertaken by ref. ^[15]). Figure ED1j-k uses a variety of increasingly flexible functional forms to estimate the relationship between temperature and growth, using either higher order polynomials (up to 7th order) or restricted cubic splines (up to 7 knots). The more flexible functional forms give response functions extremely similar to our main estimate, suggesting that the inverted-U shaped response function in our main specification is not an artifact of the parsimonious 2nd order polynomial.

C.2 Growth versus level effects

We test whether temperature is affecting the growth or the level of per capita GDP. As discussed above, the difference has important implications for the dynamics of future climate change, as a given effect size in levels will have a much smaller effect on the evolution of GDP over a long period than the same effect size on growth.^{15,35} To test for growth versus level effects, we follow ref [35] and estimate a distributed lag model and then add up effects across years. If temperature affects the level of output but not its growth rate, then the contemporaneous and lagged effects of temperature should have opposite signs: the negative effect of a hot year on output in that year would be followed by a positive effect the following year as the economy rebounds (see Extended Data Fig. ED2a). However, if warmer-than-average temperatures affect the growth rate of output, then lagged effects should be zero or could have the same sign if the effects persist. Adding up the effects across years thus allows for a straightforward test of growth versus level effects. If the sum of contemporaneous and lagged effects shrinks to zero as additional lags are added, then this indicates mainly level effects; if the summed effects are indistinguishable from the contemporaneous effect (or larger in absolute magnitude), this suggests growth effects; something in-between would suggest a combination of growth and level effects.

Importantly, a finding of “only” level effects could still imply a substantial negative impact of climate change. As visualized by the black lines in Extended Data Fig. ED2a, a level effect of a hot year on historical production means that output returns to trend following the hot year, but the loss of output pictured in the hot year is never recouped. If what happens in a historical hot year is then predictive of what will happen as future temperatures warm, then future warming will similarly lead to production losses that are not recouped – and these losses will increase over time as temperatures warm. Nevertheless, it remains the case that a given level effect will have a much smaller effect on long-run GDP than a same-sized effect on the growth rate: a 1% per °C effect on the level of output will mean that an instantaneous increase in temperature of 1°C will lower output 100 years later by 1%, whereas a 1% per °C effect on the growth rate of output implies that an instantaneous 1°C increase will lower output by about 62% 100 years later.³⁴

Because we are estimating non-linear models, we calculate the marginal effects for both the contemporaneous and lagged response functions at each point of the temperature support and add up these marginal effects over time.^{‡‡} The result from this procedure is an estimate for the *cumulative effect on income from one degree of warming*, as a function of a country’s initial temperature. Fig. ED2b displays these results for models with up to 5 annual lags, and Table S2 provides corresponding estimates in table form. While results become noisier as increasing numbers of lags are added, point estimates indicate growth effects at the hot end of the temperature distribution. At the cold end of the temperature distribution, mean cumulative effects reverse sign such that after 3 lags are accounted for, we estimate that the incomes of cold countries are also negatively affected by warming (with substantial uncertainty in this estimate). This suggests that when the dynamic effects of temperature over time are factored in, hotter countries remain worse off with additional warming, and cooler countries might not benefit on net. Importantly, however, as more lags are included, uncertainty in the cumulative effect increases – which is expected because we are adding additional uncertain parameters. This

^{‡‡}That is, in a model with one lag (and ignoring controls), $Y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \beta_3 T_{it-1} + \beta_4 T_{it-1}^2$, and the overall marginal effect on growth at some temperature T^* is then $\hat{\beta}_1 + 2\hat{\beta}_2(T^*) + \hat{\beta}_3 + 2\hat{\beta}_4(T^*)$.

makes it increasingly difficult to reject either the level or growth effect hypotheses. Thus, while we can clearly demonstrate that there is a nonlinear effect of temperature on economic production, we cannot reject the hypothesis that this effect is a true growth effects nor can we reject the hypothesis that it is a temporary level effect.

To account for this uncertainty in our projections (below), we present results where we draw from the full distribution of uncertain parameter estimates (heuristically, resampling from the confidence intervals in Fig. ED2b), using both models with zero lags (“short run” models) and models with five lags (“long run” models). As we show below, adding increasing numbers of lags tends to make the impacts of projected future temperature increases substantially more uncertain as well as *more* negative because marginal cumulative effects become negative for cold countries once more lags are included. This result is consistent with the findings in Dell, Jones, and Olken³⁵, who also find that point estimates on effects of temperature on growth are positive for rich countries in the zero-lag model, but turn negative when increasing numbers of temperature lags are added. Cumulative effects were imprecisely estimated in Dell et al., similarly to our results, and thus were not a focus of their discussion, but are consistent with what we find here.

Supplementary Table S2: **Marginal effect of temperature on growth at different points in the temperature distribution**, based on regression models that include between zero and 5 lags of temperature. For models with at least one lag, the reported estimated is the sum of the contemporaneous and lagged effects, evaluated at the indicated temperature. Numbers in parentheses provide the standard error on each estimate. Estimates correspond to those shown visually in Figure ED2c.

	Lags = 0	Lags = 1	Lags = 3	Lags = 5
5°C	0.0078 (0.0029)	0.0042 (0.0044)	-0.0051 (0.0073)	-0.0047 (0.0067)
10°C	0.0030 (0.0022)	0.0014 (0.0032)	-0.0063 (0.0054)	-0.0057 (0.0053)
15°C	-0.0019 (0.0021)	-0.0013 (0.0027)	-0.0076 (0.0043)	-0.0066 (0.0049)
20°C	-0.0068 (0.0026)	-0.0040 (0.0031)	-0.0088 (0.0046)	-0.0076 (0.0057)
25°C	-0.0116 (0.0034)	-0.0068 (0.0042)	-0.0101 (0.0062)	-0.0085 (0.0072)
30°C	-0.0165 (0.0044)	-0.0095 (0.0055)	-0.0113 (0.0083)	-0.0095 (0.0091)

C.3 Studying heterogeneous responses

Figure ED1g maps the marginal effect of +1°C warming, i.e. the predicted impact on the growth rate for +1°C, as estimated using Equation 15. Effect sizes are large for both hot tropical countries and cooler high latitude countries that are far from the temperature optimum. For instance, an additional +1°C warming is predicted to increase growth rates by >0.5 percentage points in parts of northern Europe, and to decrease growth by >1 percentage points in much of the tropics. A -1 percentage point effect means that a country growing at 2% per year in a “normal” temperature year would grow at 1% per year if the temperature were +1°C hotter. Tropical countries exhibit the largest marginal effect based on our pooled response, but they are also poorer on average. As discussed above, this makes it is less clear whether being poor or being hot makes the effect of temperature larger in these countries. If tropical countries have a larger marginal effect of temperature because they are poor, rather than because they are hot, then our pooled nonlinear model is likely misspecified. Pathbreaking work by Dell et al (2012) suggested that being poor was the critical dimension of countries that determined whether temperature had a large effect on income. Thus, we follow Dell et al. and examine whether being rich or poor alters the nonlinear relationship between temperature and growth, using the same criteria to distinguish rich and poor. In a later section, we reconcile our findings with Dell et al.

To test for differential responsiveness across rich and poor countries and across different sectors of the economy in our data, we interact the temperature and precipitation variables in Equation 15 with an indicator for whether a country’s purchasing-power-parity-adjusted (PPP) per capita income was below the global median in 1980 (the first year that PPP data are available for most of our sample; PPP incomes adjust for price differences across countries, and are important for correctly comparing income levels across countries). That is, with $D_i = 1$ for a country with below-median PPP per capita income in 1980 (and zero otherwise), then the function $h(T_{it})$ in Equation 15 becomes:

$$h(T_{it}) = \beta_1 T_{it} + \beta_2 T_{it}^2 + \beta_3 (T_{it} \times D_i) + \beta_4 (T_{it}^2 \times D_i) \quad (18)$$

with β_1 and β_2 describing the response function for rich countries, and β_3 and β_4 describing adjustments to these parameters that are only applicable to poor countries. If the response of rich and poor countries are different in structure, than the adjustments β_3 and/or β_4 will be significantly different than zero. If these adjustments are not distinguishable from zero, that indicates that the data cannot reject the hypothesis that the response estimated for rich and poor countries are the same.

Table ED2 provides regression results when Equation 18 is substituted into Equation 15. Column 1 provides our benchmark model. We find that both the linear and quadratic effects of temperature on rich countries (β_1, β_2) are significantly different from zero and that adjustments to these parameters for poor countries (β_3, β_4) are not significant. In columns 2-6 we alter our benchmark model by changing the structure of non-parametric controls (similar to Table ED1). Our results are broadly similar whether or not the year fixed effects are allowed to differ between rich and poor countries^{§§} (column 2), whether or not the sample is restricted to countries with at least 20 observations (columns

^{§§}For instance, splitting the sample and running separate panel regressions with year fixed effects for both rich countries and for poor countries implicitly allows the year fixed effects to differ across country groupings. The pooled model that interacts the climate variables with an indicator for whether a country is poor implicitly assumes common year effects across all countries unless these year fixed effects are also interacted with the poor-country indicator. We show in the table that results are similar in either case.

3-4), and whether continent-by-year fixed effects are included in addition to (column 5) or instead of the country time trends (column 6). In all cases the response functions in both rich and poor countries are concave, and in no case are the interaction terms significant, suggesting that we can never reject that rich and poor countries have the same response function to temperature. In columns 2 and 4, statistical uncertainty increases so that estimated parameters for rich countries are no longer significant, but point estimates in these models remain virtually unchanged and the resulting estimates do not reject the benchmark model.^{¶¶} We note that the specification of time controls preferred by Dell et al is the same as the model in column 6, which suggests the nonlinear effect of temperature on rich countries is highly statistically significant and not significantly different from the effect in poor countries.

A different but related question asks whether the *slope* of the response function $h(T)$ is itself statistically different from zero at different temperatures. Testing marginal effects at different temperatures is different than testing for significant parameter estimates (e.g. β_1 and β_2) because our model is nonlinear, so the derivative of $h(T)$ is itself a function of temperature and thus must be evaluated at specific temperatures in order to have meaning. In our benchmark pooled model (Figure 2a in the main text) the derivative of $h(T)$ is clearly different from zero at almost all points except near the optimum, where the derivative is near zero. The estimated derivative of this function is displayed in the “0 lags” panel of Figure ED2c. At low temperatures, warming has positive effects on production, while at high temperatures, warming has clear negative effects.

When our data are broken into subsamples (e.g. rich and poor), statistical uncertainty increases and so the estimated function $h(T)$ exhibits wider confidence intervals, such as in Figure 2b in the main text. In these subsamples, we continue to find that structures in the response (such as the parameter β_2 describing the curvature) are statistically significant even though the function $h(T)$ does not always separate from zero over the range of the data. In these cases, it is also interesting to inspect whether the slope of $h(T)$ is significant at different temperatures for specific subsamples. Importantly, such significance is *not* critical to our main result concerning the nonlinear structure of $h(\cdot)$ —for a quadratic response there must be regions of $h(\cdot)$ where the slope is zero and with any uncertainty, there will likely be a wide range of temperatures with slopes indistinguishable from zero. Nonetheless, understanding the uncertainty of marginal effects is useful for understanding the extent to which rich and poor countries respond similarly to temperature at specific temperatures, and our impact projections incorporate these uncertainties.

Figure ED2 displays the estimated marginal effect of temperature for total growth (panels d-e), agricultural growth (g-h), and non-agricultural growth (j-k) for both rich and poor countries (these are the derivatives of functions in Figure 2b, d, & e). For total growth, the marginal effect on rich countries at cold temperatures is significantly positive while it is significantly negative for poor countries at hot temperatures. A similar pattern holds for agricultural and non-agricultural growth, with the exception of non-agricultural growth in rich countries where high temperatures have a slightly clearer effect than cold temperatures. The structure of these uncertainties primarily reflects the distribution of rich and poor countries: because there are many rich countries at cold temperature and many poor countries at hot temperatures, the structure of $h(\cdot)$ is most precisely estimated in those respective regions of the

^{¶¶}For a complete discussion of cross-model comparisons, see ref [42]

support.

Figure ED2f, i, & l then tests for differences between rich and poor countries in the slope of their response functions at each point in the temperature distribution for our main sample and estimation approach, corresponding to Figure 2b, d, & e. While point estimates suggest that richer countries have a slightly flatter response function for overall GDP and non-agricultural GDP, and somewhat steeper response function for agricultural GDP at higher temperatures (the latter consistent with the finding that less-nutrient limited agricultural systems are actually more sensitive to changes in climate⁴³), at no point in the temperature distribution can we reject the hypothesis that rich and poor countries have the same response function for any of these three outcomes.

The analysis above finds no evidence that rich and poor countries respond differently to temperature, although it also finds many temperatures for which the marginal effect of temperature is not statistically significant in a specific subsample. Interpreting these facts requires care, especially when using pair-wise hypothesis tests. A common error is to dismiss a finding using pooled data as irrelevant for a subsample if an identical estimate using just the subsample is not significant on its own. Such logic is clearly flawed because for any sample of data, a real result will be rendered insignificant if the sample is sufficiently reduced through subsampling. In our setting, we can not know for certain whether rich and poor countries should be pooled or not because we cannot observe the true data generating process and verify that it is in fact the same for rich and poor countries. As stated above, we cannot reject the hypothesis that rich and poor respond identically, but we also cannot reject the hypothesis that the marginal effect of temperature is zero for many temperature values in both rich and poor subsamples using standard critical values. This is particularly relevant for global estimates of the impacts of climate change, where the effect of hot temperature on rich countries are especially influential.

We thus ask whether it is more likely that (A) rich countries (in their reduced subsample) have a negative effect at high temperatures (similar to poor countries) or (B) rich and poor countries have different responses at high temperatures. To consider this, in Figure ED2m-u we plot the p-values for the marginal effects of temperature shown in panels d-l. These p-values represent an estimate of the probability that we would recover an estimated marginal effect as large in magnitude as the estimate we actually recovered at a specific temperature, if the true marginal effect were zero. We find that for all three measures of production, it is more likely that (A) rich countries have negative marginal effects at high temperatures (p-values are smaller) than (B) rich countries have a different marginal effect than poor countries at high temperatures. Nevertheless, since we cannot rule out that rich and poor countries respond differently, or that rich countries do not respond at all, as described in Section D we report impact projections account for the full range of possible responses for both rich and poor countries.

To summarize our various findings on potential heterogeneity:

1. Assuming a single $h(\cdot)$ describes effects for all countries, we find highly significant nonlinearity in the response to temperature.
2. Considering rich countries alone, we find significant nonlinearity in the response to temperature.
3. Considering poor countries alone, we find significant nonlinearity in the response to temperature.

4. We fail to reject the hypothesis that nonlinearity in the response of rich countries (considered alone) mirrors the pooled estimate.
5. We fail to reject the hypothesis that nonlinearity in the response of poor countries (considered alone) mirrors the pooled estimate.
6. We fail to reject the hypothesis that nonlinearity in the response of rich countries (considered alone) mirrors nonlinearity in the response of poor countries (considered alone).
7. Assuming a single $h(\cdot)$ describes effects for all countries, we find highly significant marginal effects of temperature at almost all temperatures.
8. We fail to reject the hypothesis that the marginal effect of temperature for rich countries (considered alone) mirrors the pooled estimate.
9. We fail to reject the hypothesis that the marginal effect of temperature for poor countries (considered alone) mirrors the pooled estimate.
10. We fail to reject the hypothesis that the marginal effect of temperature for rich countries (considered alone) mirrors the marginal effect of temperature for poor countries (considered alone).
11. We find highly significant marginal effects of cold temperatures for rich countries (considered alone) and marginally significant marginal effects of hot temperatures for rich countries (considered alone).
12. We find highly significant marginal effects of hot temperatures for poor countries (considered alone).
13. We fail to reject at conventional confidence levels the hypothesis that the marginal effect of high temperature for rich countries (considered alone) is zero for temperature ranges where pooled estimates are significantly non-zero (although this is unsurprising given the reduced sample and temperature distribution of rich countries).
14. We fail to reject the hypothesis that the marginal effect of low temperature for poor countries (considered alone) is zero for temperature ranges where pooled estimates are significantly non-zero (although this is unsurprising given the reduced sample and temperature distribution of poor countries).
15. We find it is substantially more likely that (A) the marginal effect of high temperatures on rich countries is negative than (B) the marginal effect of these temperatures on rich countries is different from their marginal effect on poor countries.

Taken together, these results provide limited evidence of heterogeneity between rich and poor country responses, and strong evidence that assuming rich and poor countries are similar is more likely correct than assuming zero marginal effect in cases where subsamples do not return significant effects when considered alone. Again, as in Figure ED1i, differential responses to temperature appear driven more by underlying differences in average temperature than underlying differences in average again.

Nevertheless, there is suggestive but statistically insignificant evidence that the aggregate rich-country response function may be flatter than the poor-country response function, and so in constructing the impact projections under climate change (described below), we estimate impacts under both a pooled model where response functions across countries are assumed to be the same, as well as using the interacted model where rich and poor countries are allowed to respond differently to future changes in climate based on the differential response functions shown here. Notably, in all of these projections, we account for the full range of statistical uncertainty in the estimated structure of $h(\cdot)$.

C.4 Comparison with Dell Jones Olken 2012

Seminal earlier work by Dell, Jones, and Olken³⁵ (hereafter DJO) focused on the linear relationship between growth and temperature fluctuations, finding strong negative effects of warmer temperatures on growth in poor countries but not in rich countries. There is a *prima facie* case for DJO's results being consistent with ours: with most poor countries on the downward slope of the response function but rich countries distributed almost symmetrically around the optimum, a linear regression for the effect of temperature would recover a steep negative in poor countries but ambiguous (and closer to zero) slope for rich countries. Thus a globally concave response function is consistent with DJO's result describing a steep negative linear response in poor countries and a flatter average linear response in rich countries.

Here we explore in more detail differences between our results and DJO. Other than differences in the specified functional form of the temperature-growth relationship, there are other modest differences between DJO and our analysis:

- Differences in sample: DJO data sample ended in 2003 and used earlier versions of the World Bank and University of Delaware data, both of which we have updated through 2010. Older versions of World Bank data appear to have included growth data for earlier years for many countries (e.g. for the 1950s); these are not included in the current World Bank database and are thus omitted from our analysis.
- Differences in specification of controls for time-varying omitted variables: DJO preferred specification includes continent-year fixed effects, whereas ours includes year fixed effects and flexible country time trends. Above we demonstrated that use of the DJO model recovered our main result with exceptionally high levels of significance, although we use a model with country-specific quadratic trends in growth because it provides a stationary time series, it accounts for time-varying country-specific factors, and it performs better in terms of prediction²⁶.

Continent-year FE have the advantage of controlling for continent-specific shocks that might be correlated with both temperature and growth in a given year, but the disadvantage of potentially absorbing important variation in temperature if temperature shocks are highly covariate within a continent⁴⁴ (e.g. if a given year is hot in both Kenya and Tanzania). Our main specification, which uses year FE and country trends, will allow a year that is hot in all countries in a given continent (e.g. due to ENSO variation) to inform regression results. Inclusion of flexible country trends ensures that our estimates are not confounded by gradual changes in economic growth rates, such as the relatively recent rapid growth of China since the 1980's, which might be spuriously correlated with recent warming.

We explore how our and DJO's results change under alternate time controls, but allowing for a potentially nonlinear (here quadratic) functional form for temperature (again, DJO's original specification was linear). To make our results most comparable, we restrict our "base" sample to be countries with at least 20 years of growth data, as in DJO, but results are similar if we use the full sample (recall Table ED1). We then gradually vary our data/model to be more similar to the DJO data/model, and visa versa. These various results are shown in Figure ED3a and Table S3. Estimated response functions are similar between studies using the global sample: all estimated functions show a concave relationship between growth and temperature, although as quantified in the bottom row of Table S3, the estimated temperature optimum is uniformly lower in all alternative samples and specifications relative to our baseline specification (shown in the 4th column). Nevertheless, the estimated quadratic temperature response in DJO is estimated less precisely than in our updated data, particularly when continent-year fixed effects are included (columns 1 and 2 of Table S3). As explored in the last two columns of the table, which restrict the sample to the countries and years that we can match between our data and theirs, results become somewhat less precise as DJO's earlier versions of the growth or temperature data are substituted in for our growth or temperature data. If these data have improved over time, and if our seven additional years of data from 2004-2010 are also measured with less error (particularly relative to the data in the 1950s included in DJO but not included in the updated WDI database), then this could explain why our estimates are somewhat more precise than DJOs.

In summary, the non-linear structure of the temperature/growth relationship is more apparent in updated data. However, had DJO used our benchmark model which accounts for important country-specific trends, nonlinear structure would have been apparent. Alternatively, had DJO used their preferred model but had access to either the updated growth or temperature data (this was impossible for them) they would have observed significant nonlinear structure. Finally, our main result is robust to applying DJO's preferred model to the updated datasets.

As described below, the projected impacts of climate change are somewhat *less* negative when using our preferred specification and data as compared to projections that use a quadratic specification with DJO's time-varying controls and/or data (Figure ED3b). Projections that use DJO's original linear specification are less stable across specifications (changing sign depending on whether or not lags are included), but as shown in Figure ED3c-d and discussed in the main text, their linear estimates suggest much more positive (or less negative) impacts relative to comparable non-linear models – indicating the critical importance of accounting for non-linearities in understanding the potential economic implications of climate change.

Supplementary Table S3: **Comparing our results to DJO 2012.** (1) original DJO sample and specification, except quadratic (not linear) in temperature, (2) as in 1 but adding country trends and precipitation, (3) as in 2 but replacing continent-year FE with year FE, (4) baseline result from this study, (5) as in 4 but restricting to DJO sample of countries/years using updated data, (6) as in 5 but with DJOs FE of continent-year FE, poor-year FE, and no country trends, (7) as in 6 but using original growth data from DJO, (8) as in 6 but using original temperature data from DJO. Last line of table give estimated optimum temperature for each model. All models include country FE, and following DJO the standard errors are two-way clustered at the country and year level. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specification:	DJO+quad	(1)+trend	(1)+yearFE+trend	BHM	BHM	DJO+quad	BHM	BHM
Sample:	DJO	DJO	DJO	BHM >20yrs	DJO	BHM >20yrs	BHM >20yrs	BHM >20yrs
Temp. data:	DJO	DJO	DJO	BHM	BHM	BHM	BHM	DJO
Growth data:	DJO	DJO	DJO	BHM	BHM	BHM	DJO	BHM
Temperature	0.0036 (0.0052)	0.0083 (0.0064)	0.0100* (0.0059)	0.0135*** (0.0038)	0.0084** (0.0037)	0.0086** (0.0035)	0.0100** (0.0043)	0.0070* (0.0039)
Temperature squared	-0.0002 (0.0002)	-0.0003 (0.0002)	-0.0004** (0.0002)	-0.0005*** (0.0001)	-0.0004*** (0.0001)	-0.0004*** (0.0001)	-0.0004*** (0.0001)	-0.0003** (0.0001)
Precip.		0.0068 (0.0131)	0.0140 (0.0123)	0.0148 (0.0101)	0.0198* (0.0111)	0.0062 (0.0115)	0.0048 (0.0117)	0.0063 (0.0116)
Precip. squared		-0.0027 (0.0029)	-0.0040 (0.0030)	-0.0049* (0.0026)	-0.0056* (0.0030)	-0.0015 (0.0027)	-0.0010 (0.0028)	-0.0016 (0.0027)
Observations	4924	4924	4967	6477	4705	4705	4796	4705
R squared	0.220	0.294	0.223	0.278	0.218	0.207	0.212	0.206
Optimum	8.90	12.05	12.16	13.39	9.99	12.34	12.75	12.45

D Climate projections and their uncertainty

D.1 Building impact projections

To project the impacts of climate change, we use our historical response functions to generate projected future changes in GDP under various warming scenarios, relative to a world in which temperatures remained fixed at their 1980-2010 average. In particular, the evolution of GDP per capita in country i in year t is given by:

$$GDPcap_{it} = GDPcap_{it-1} \times (1 + \eta_{it} + \delta_{it}) \quad (19)$$

The growth rate absent climate change is given by η_{it} , which we either take from the Shared Socio-economic Pathways (SSPs) or fix to be the average growth rate over 1980-2010 (“base” scenario). The SSPs are “reference pathways describing plausible alternative trends in the evolution of society and ecosystems over a century timescale”⁴⁵, and we focus on the two SSPs (SSP3 and SSP5) that are most consistent with the high emissions needed to generate RCP8.5.^{***} In SSP3, “unmitigated emissions are high due to moderate economic growth, a rapidly growing population, and slow technological change in the energy sector”, while in SSP5, “In the absence of climate policies, energy demand is high and most of this demand is met with carbon-based fuels... Nonetheless, economic development is relatively rapid”.⁴⁵ Both SSP scenarios specify country-specific trajectories for *average income per capita*, using demographic models to account for population changes. As noted, in addition to these two SSPs, we also consider a “baseline” scenario in which each country is assumed to grow in the future at its observed average rate of growth during 1980-2010. In this baseline scenario, we use population projections from the UN World Population Prospects⁴⁶ to compute average income per capita.

The parameter δ_{it} is the predicted additional effect of warming on growth in that year. Denote \bar{T}_i as the average temperature in country i between 1980-2010 (the base period), and T_{it}^+ as the projected temperature in any year after 2010. For our estimated pooled historical response function $h(T)$,

$$\delta_{it} = h(T_{it}^+) - h(\bar{T}_i) \quad (20)$$

For the projections in which the future growth response to temperature is allowed to vary by income, denote the rich-country response function as $h_r(\cdot)$, the poor country response function as $h_p(\cdot)$, and y^* as the income level above which countries’ responses are described with $h_r(\cdot)$ rather than $h_p(\cdot)$ (which we set to the median GDP/cap in the historical period, as in our estimation). Then δ_{it} is calculated as:

$$\delta_{it} = \begin{cases} h_r(T_{it}^+) - h_r(\bar{T}_i) & \text{for } GDPcap_{it-1} > y^* \\ h_p(T_{it}^+) - h_p(\bar{T}_i) & \text{for } GDPcap_{it-1} \leq y^* \end{cases} \quad (21)$$

Based on Equation 18 above, $h_r(T_{it}) = \beta_1 T_{it} + \beta_2 T_{it}^2$ and $h_p(T_{it}) = (\beta_1 + \beta_3)T_{it} + (\beta_2 + \beta_4)T_{it}^2$.

We assume a linear increase in temperature between 2010 (the end of the base period) and the RCP8.5 projected country-specific temperature in 2100, such that for any year $t > 2010$ and a projected

^{***}SSP data were downloaded from <https://secure.iiasa.ac.at/web-apps/ene/SspDb/dsd?Action=htmlpage&page=about> on Feb 10th, 2015. The available data provide projected growth rates at 5-year intervals at the country level, and we linearly interpolate within each 5-year interval to calculate the annual projected growth rates. The SSP also provide growth projections generated by three different research groups; we focus on the projections from the OECD group, as they were the only ones to project growth rates at the country level.

country-specific warming of ΔT_i by 2100 (where ΔT_i is the change between 2010 and 2100), $T_{it}^+ = \Delta T_i \times \frac{t-2010}{2100-2010} + \bar{T}_i$. So for example if $\Delta T_i = +4.3^\circ\text{C}$ (which is the population-weighted global change in average temperature projected in RCP8.5), then in 2111 $T_{i2011}^+ = \frac{1}{90} * 4.3 + \bar{T}_i = 0.048^\circ\text{C} + \bar{T}_i$. With no warming, $T_{it}^+ = \bar{T}_i$ and $\delta_{it} = 0$. Our estimates of country-specific warming ΔT_i are the ensemble mean projected warming for RCP8.5 for each country across all global climate models contributing to CMIP5, calculated by taking a population-weighted average of each climate model grid cell covering a country.^{†††} Projected changes in temperature to 2100 across our sample countries range from 2.7°C to 5.8°C .

We do not know how economic production responds to temperatures that have never been observed historically. Thus, when countries warm beyond the highest observed temperatures in the historical data, we have two options: either we extrapolate the function $h(\cdot)$ beyond the support of historically observed data, or we assume that productivity is equal to the boundary value for all observations beyond the boundary of the support, i.e. $h(T) = h(T_{max})$ for all $T > T_{max}$. We opt for the latter approach because we view it as more conservative, since extrapolation of $h(\cdot)$ causes income to fall even more rapidly at higher temperatures. We cap T_{it}^+ at 30°C , which is the upper bound of the annual average temperatures observed in our sample period. Thus for any year where projected warming will increase average temperatures in a given country beyond 30°C , we fix the effect of that year of additional warming at $\delta = h(30) - h(\bar{T})$. This may cause us to understate losses in countries that eventually exceed 30°C if the “true” response function continues to fall off steeply beyond 30°C , so our projected impacts in the hottest countries are likely conservative. Nevertheless, for most countries in the sample, substantial warming (e.g. $+4^\circ\text{C}$) takes them out of their own historical range of temperature exposure but leaves them well within the observed global distribution of historical temperatures, meaning that we would not be extrapolating out of sample for these countries in any case.

To aggregate national average incomes to global average income, we calculate Gross World Product (GWP) each future year t :

$$GWP_t = \sum_i \omega_{it} \times GDP_{cap_{it}} \quad (22)$$

where ω_{it} is country i 's projected population in year t . This calculation is done separately for warming scenarios and the counterfactual scenario with no warming. As above, country-specific population projections are taken either from the SSPs or from the UN World Population Prospects, consistent with whichever population was used to compute the national projections. Comparing changes in GWP for the warming scenarios to changes in our counterfactual no-warming scenario then gives us projected global losses (or gains) under climate change, an estimate which can then be compared to “damage functions” currently being used by Integrated Assessment Models (IAMs).⁴⁷

The choice of underlying socioeconomic scenario can affect our estimate of global impacts by altering how impacts are weighted across countries. A country that is projected to grow quickly in either income or population in a given scenario (i.e. where η_{it} or ω_{it} are large) will receive higher relative weight in the estimate for global average income loss (in %) than a country predicted to grow more slowly. Thus differences across scenarios in these assumed underlying growth rates could shape estimates of GWP_t ,

^{†††}Data were downloaded from http://climexp.knmi.nl/plot_atlas_form.py on March 2, 2015.

motivating our inclusion of multiple plausible socio-economic scenarios.

Regional projections are constructed similarly, where for region R ,

$$GDP_{Rt} = \sum_{i \in R} \omega_{it} \times GDP_{cap_{it}} \quad (23)$$

We define nine regions based largely on UN designations^{†††}, with “Sub-Saharan Africa” containing UN designations East Africa, Middle Africa, Southern Africa, and Western Africa; “Latin America” containing UN subregions Caribbean, South America, and Central America; “Middle East/North Africa” containing Northern Africa and Western Asia; “Central and East Asia” containing Central Asia and Eastern Asia; “Europe” containing Eastern Europe, Northern Europe, Southern Europe, and Western Europe; and “Oceania” containing Australia and New Zealand, Melanesia, Micronesia, and Polynesia.

To quantify uncertainty in these projections, we block bootstrap the historical response function (1000 times, sampling countries with replacement to account for autocorrelation) and apply the above procedure to each bootstrap. We use the resulting distribution of estimates to characterize projection uncertainty for a given set of assumptions about future temperature change and baseline economic and population growth rates (as just described). In Fig. 4, Fig. 5a, and Figs ED4-ED5, projections are then “visually-weighted”⁴⁸ to illustrate distribution of projections in each future year. This is done by calculating, for each future year 2011-2099, the probability that a given level of impact is observed in the set of 1000 bootstrap runs. The plot is then shaded such that darker areas represent values for which projected impacts were more common. Probabilities are made to sum to one in each year, such that the total “ink” is held constant across each vertical slice (year), and so appears more “spread out” when projections show less agreement.

D.2 Projected impacts and robustness to alternative specifications

To understand the sensitivity of our impact projections to assumptions, we explore how projections change when (i) the temperature response function is estimated using different assumptions and (ii) we adopt different assumptions regarding the future evolution of population and economic activity in the absence of climate change.

We first alter whether the projection accounts for lagged weather effects and whether rich and poor countries are assumed to have different responses (although the data do not reject the assumption that the pooled response is representative, as discussed above). To do this, we generate four separate estimates:

1. “Pooled SR”: assumes a common response across rich and poor countries to temperature changes, and growth in a given year is only affected by temperature in that year (hence “SR” for “short run”). This is the benchmark estimate presented in Figure 2a of the main text.
2. “Pooled LR”: assumes a common response across rich and poor countries to temperature changes, but growth in a given year can be affected by temperature in that year and the previous 5 years (hence “LR” for “long run”). Marginal effects from this model were shown in Figure ED2c.

^{†††}As viewed here <http://unstats.un.org/unsd/methods/m49/m49regin.htm> on Feb 20th 2015

3. “Differentiated SR”: allows rich and poor countries to respond differently to temperature changes, and growth in a given year is only affected by temperature in that year. This is the estimate presented in Figure 2b of the main text. If poor countries grow such that their income surpasses the threshold income used to separate rich and poor countries, they “graduate” to exhibiting the rich country temperature response, as in Equation 21.
4. “Differentiated LR”: allows rich and poor countries to respond differently to temperature changes and for poor countries to “graduate” to the rich country response, but growth in a given year can be affected by temperature in that year and the previous 5 years

Fig. 4b shows global projected impacts using each of these historical response functions under a common underlying socioeconomic scenario (SSP5). Analogous regional estimates (with uncertainty) are shown in Fig. ED5, with the black line in each plot representing the “best guess” of projected impacts (i.e. using our point estimates) and the red shaded area the 5-95% confidence interval, with the color intensity again indicating the likelihood of the projection estimate falling at a given value in a given year.

In models that account for lagged effects or that estimate rich and poor responses separately, projections become more uncertain. This is both because the estimated rich-country response function is slightly flatter than the pooled response function, which slightly alters estimated damages in regions such as North America, but also because the estimated poor-country response function peaks a few degrees °C warmer than the pooled response function, which serves to slightly lower the damages in poor countries as well. However, due to more uncertainty in the income-specific response functions, caused by splitting the sample (in the differentiated model) and accounting for additional uncertain parameters (in the LR model), overall projection uncertainty under these models is larger than under the pooled SR model.

Another important difference is that for LR models, the mean projection is substantially negative for all regions because cold rich regions (e.g. Europe) exhibit net negative impacts in models that allow the growth effects of a single hot year to persist over multiple years. The reason for this is shown in Figure ED2c: as additional lags of temperature are allowed to affect growth in a given year, hotter countries remain worse off with additional warming, and cooler countries no longer unambiguously benefit, with point estimates of marginal effects being negative across the entire temperature distribution. This means that evaluating the impacts of future warming using the “long-run” historical response functions makes economic impacts more negative for countries that are initially cold.

We next consider whether our global projections are sensitive to these modeling assumption while also examining their sensitivity to assumptions about the baseline economic trajectory of countries. The three socioeconomic scenarios we evaluate are:

1. “base”: absent climate change, countries per capita incomes grow every year in the future at the rate they grew on average between 1980-2010. Estimates of future country-level populations are taken from UN projections.⁴⁶
2. SSP3: A “shared socio-economic pathway” with generally slower income growth and less long-run convergence in income levels between poor and rich countries. Under this scenario, the

population-weighted global average growth rate without climate change is 1.2% in 2050 and 1.1% in 2090.

3. SSP5: A “shared socio-economic pathway” with generally higher income growth and faster long-run convergence in income levels between poor and rich countries. Under this scenario, the population-weighted global average growth rate without climate change is 3.6% in 2050 and 2.3% in 2090.

Figure ED4 displays mean projections and uncertainty for the estimated global response under RCP8.5 due to our three different possible underlying growth scenarios and four different response functions. Table ED3 provides the corresponding point estimates and percentiles in the distribution of responses for each scenario.

Similar to the regional projections, mean estimates are uniformly more negative under the LR response functions that allow temperature to have persistent effects on growth. Also similar to the regional projections, estimates are more uncertain – but slightly less negative – when we use the response functions that allow rich and poor countries to respond differently to warming. The underlying socioeconomic scenario that is assumed has little impact on projected changes to GWP.

To understand whether the *nonlinear* response to temperature is important for determining the projected impact of climate change, we compare our results to estimated impacts obtained using coefficients from the *linear* specification in Dell et al 2012 (DJO). The published version of Dell et al 2012 does not contain impact projections, and so we construct them for this purpose; note that these are distinct from the impact projections shown in Figure ED3b, which apply a quadratic functional form to DJO’s data and/or model. As described in Section C.4, DJO estimate the linear effect of temperature on growth, allowing this effect to differ across rich and poor countries. They find large negative effects of temperature on historical growth rates in poor countries ($\sim -1\%$ per $^{\circ}\text{C}$ or larger) and smaller effects in rich countries, with point estimates in rich countries ranging from slightly positive ($+0.26\%$ per $^{\circ}\text{C}$) in the zero-lag model to slightly negative (-0.19% per $^{\circ}\text{C}$) in the 5-lag model, with neither estimate statistically distinguishable from zero.

Using our projection approach, we combine these coefficients from the *linear* DJO model (drawn from their Table 4) with projected temperature changes to estimate regional and global impacts. Results are shown in Figure ED3c-d, and differ from the main impact projections in this study in two main ways. First, the sign of the projected impact using DJO coefficients changes whether one considers the zero-lag or 5-lag model, consistent with the changing sign on the rich-country coefficient. This is in contrast to our estimates, which remain negative in either model. Second, projected regional impacts using DJO coefficients also differ in sign across models, and for most regions remain substantially less negative using DJO coefficients even in the model with 5 lags. This occurs both because our estimated marginal effects at high temperatures are more negative than the poor-country linear effect in the DJO 5-lag model, but also because as countries “graduate” to rich-country status over time, the marginal effect of temperature in DJO reduces by about 80%. Thus, properly accounting for the nonlinear response to temperature has profound impact on the projected economic impact of climate change. Rather than all countries having a sensitivity that declines over time as they become richer, most countries have a sensitivity that increases over time because they become hotter on average. If the nonlinear effect of temperature is not accounted for, then the negative correlation between modern-day

income and average temperature across countries dramatically confounds projections.

Panels d, h, l, & p in Figure ED4 shows projected impacts by 2100 under RCP8.5 by each country's baseline income quintile. Different panels display projections based on the different historical response functions. The “pooled response, SR effect” panel is the same as Fig. 5c in the main text. Echoing the global results, impacts are more uniformly negative across the income distribution when increasing numbers of temperature lags are included. This is because the estimated marginal effect of temperature becomes slightly negative even at cold temperatures in the multi-lag historical models where growth effects are allowed to persist, meaning that even rich countries are on average worse off even with small amounts of warming. Utilizing a response function that is differentiated between rich and poor countries has little effect because these two subsamples both behave similarly to the pooled estimate.

D.3 Constructing the damage function

We estimate a global “damage function” by projecting economic impacts as a function of future changes in global mean temperature. To construct this function, we repeat the projection exercise described above but under alternate amounts of warming in 2100.⁴⁷ To match integrated assessment model (IAM) estimates, temperature changes are calculated “relative to pre-industrial” rather than relative to the present day, as in our above runs. To convert changes in global mean temperature to country-specific estimates needed for our impact calculations, we linearly scale country-specific changes ΔT_i (CMIP5 ensemble mean under RCP8.5) with global mean temperature, correcting for warming that has already occurred.

Specifically, let $\overline{\Delta T^{rcp8.5}}$ be the global average temperature change for RCP8.5 and let ΔT_i be the individual country projected temperature changes (as described above). Our scaling factor is then $\lambda_i = \Delta T_i / \overline{\Delta T^{rcp8.5}}$. For most countries, $\lambda_i > 1$ because land surfaces tend to warm more than the oceans, which have substantial weight in the global averages. Then for an arbitrary global mean temperature increase relative to pre-industrial $\Delta \bar{T}^s$, country-specific temperature changes between 2010-2100 for that scenario are estimated as:

$$\Delta T_i^s = \lambda_i (\Delta \bar{T}^s - 0.8) \quad (24)$$

where 0.8°C is the average warming between “pre-industrial” and present day. We then construct our damage function beginning at +0.8°C relative to pre-industrial (which means a global mean temperature increase of +0°C by 2100 relative to today, i.e. no additional warming beyond what has already occurred). Thus +1°C for the damage function shown in Fig. 5d thus corresponds to +0.2°C warming by 2100 relative to today, +2°C in Fig. 5d corresponds to +1.2°C, etc. Projections are then calculated as above, with assumed linear warming between present day and 2100, and the warming increment for each country in a given year t applied to the level of per capita GDP in the previous year to derive the GDP per capita in that country in year t .

D.4 Shape of the damage function

Why are our estimates of global losses roughly linear in temperature rather than quadratic or exponential as previously theorized? As mentioned in the main text, approximate linearity results from

the broad distribution of initial country temperatures along different parts of a smooth response function, causing the average derivative of the productivity function to change little as countries warm, at least across the projected temperature changes we consider here. The intuition that global economic damages are highly nonlinear because micro-level responses to temperature are highly nonlinear is not correct.

In several cases, the damage function appears curved, but in such a way that marginal damages decline with additional warming. This contrasts even more starkly with the intuition that damages should increase dramatically with warming,^{§§§} so it is worth being clear why this occurs in our analysis. The declining marginal damages in our projections occurs because (i) temperature appears to alter the growth rate of countries, and (ii) cumulative loss of a lowered growth rate increases with the size of the growth penalty, but at a declining rate. Thus, additional warming that leads to a net reduction in global growth rates will generate losses at a declining rate.

To see this, and using the notation from above, a country's per capita income in year t is:

$$GDPcap_{it} = GDPcap_{it-1} \times (1 + \eta_{it} + \delta_{it})$$

where η_{it} is again the assumed growth rate absent climate change and δ_{it} is the additional impact of warming in that year. So starting from initial conditions in 2010, the income of a country integrates to

$$GDPcap_{i\zeta} = GDPcap_{i,2010} \times e^{(\eta_i + \delta_i)\zeta}$$

at time ζ , where we imagine that η_{it} and δ_{it} are fixed across time at η_i and δ_i for simplicity. The percent loss in income, relative to a counterfactual no-climate-change scenario where $\delta_i = 0$ for all t is

$$\begin{aligned} \% \Delta GDPcap_{i\zeta} &= \frac{GDPcap_{i,2010} \times e^{\eta_i \zeta} - GDPcap_{i,2010} \times e^{(\eta_i + \delta_i)\zeta}}{GDPcap_{i,2010} \times e^{\eta_i \zeta}} \\ &= 1 - e^{\delta_i \zeta} \end{aligned} \quad (25)$$

Thus the separation between a climate change income trajectory and a baseline trajectory increases like an exponential function of δ_i , the growth rate penalty imposed by climate change. Because most countries are hotter than the estimated optimum, the δ imposed by warming is negative for most of the global economy. With greater warming, δ becomes increasingly negative. Figure ED6e plots $e^{\delta\zeta}$ for a range of δ . The difference $1 - e^{\delta\zeta}$ increases as δ becomes more negative (i.e. warming), but at a declining rate. Thus the compounding effect of a negative exponential causes the total economic damage projected from warming to increase with warming, but at a declining rate.

D.5 Damage function uncertainty

Figure ED6a shows how the estimated damage function changes when the assumed underlying growth scenario changes and when the assumed structure of $h(\cdot)$ changes (using the variations described above). Overall, the underlying response function represents a larger source of uncertainty than the

^{§§§}Importantly, many researchers argue that damages increase nonlinearly with warming because large-scale irreversible environmental changes might occur that lead to dramatic losses. These effects are not captured by our analysis, and so our results should not be interpreted as overturning this specific intuition.

underlying socioeconomic scenario. Generally, the LR response functions that allowed for persistent effects of one-year warming suggest more negative responses than response functions only allowing for warming in a given year to affect growth in that year, and response functions that allow rich and poor countries to respond differently to future warming (as in Figure 2b) show somewhat smaller losses. All damage functions are negative throughout the range of considered temperature increases, and are either roughly linear in temperature or show slight diminishing effects. To compare our estimates to those from integrated assessment models, Figure ED6b-d then shows the ratio of IAM-estimated damages from three leading IAMs (DICE, FUND, PAGE) to damages estimated in this study under these same sources of uncertainty. At levels of warming below 2°C, our estimates are at least 3 times higher than IAM estimates, typically 5-20 times higher, and sometimes up to 100x higher. At higher levels of warming, our estimates are again at least 2.5 times larger than the highest IAM estimate and typically are much larger.

References

- [1] Schlenker, W. & Roberts, M. Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences* **106**, 15594–15598 (2009).
- [2] Burke, M. & Emerick, K. Adaptation to climate change: Evidence from us agriculture. *SSRN working paper 2144928* (2014).
- [3] Graff Zivin, J. & Neidell, M. Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics* **13** (2014).
- [4] Zivin, J. S. G., Hsiang, S. M. & Neidell, M. J. Temperature and human capital in the short-and long-run. Tech. Rep., National Bureau of Economic Research (2015).
- [5] Hsiang, S. M. Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences* **107**, 15367–15372 (2010).
- [6] Heal, G. & Park, J. Feeling the heat: Temperature, physiology & the wealth of nations. *NBER working paper 19725* (2013).
- [7] Deryugina, T. & Hsiang, S. M. Does the environment still matter? daily temperature and income in the united states. *NBER working paper 20750* (2014).
- [8] Deschênes, O. & Greenstone, M. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics* **3**, 152–185 (2011).
- [9] Burke, M., Hsiang, S. M. & Miguel, E. Climate and conflict. *Annu. Rev. Econ.* (2015).
- [10] Massetti, E., Mendelsohn, R. O. & Chonabayashi, S. Using degree days to value farmland (2014).
- [11] Schlenker, W., Hanemann, W. M. & Fisher, A. C. The impact of global warming on us agriculture: an econometric analysis of optimal growing conditions. *Review of Economics and Statistics* **88**, 113–125 (2006).
- [12] Greene, W. H. *Econometric Analysis, Fifth Edition* (Prentice Hall, 2003).
- [13] Nordhaus, W. & Boyer, J. *Warming the World: Economic Models of Global Warming* (The MIT Press, 2000).
- [14] Solow, R. M. A contribution to the theory of economic growth. *The quarterly journal of economics* 65–94 (1956).

- [15] Moore, F. C. & Dias, D. B. Temperature impacts on economic growth warrant stringent mitigation policy. *Nature Climate Change* **5**, 127–131 (2015).
- [16] Hsiang, S. M., Jina, A. S. *et al.* Geography, depreciation, and growth. *American Economic Review* **105**, 252–56 (2015).
- [17] Dell, M., Jones, B. F. & Olken, B. A. Temperature and income: Reconciling new cross-sectional and panel estimates. *American Economic Review: Papers and Proceedings* **99**, 198–204 (2009).
- [18] Deaton, A. *et al.* Savings in developing countries: theory and review. In *Proceedings of the World Bank annual conference on development economics 1989.*, 61–108 (World Bank, 1990).
- [19] Udry, C. Risk and saving in northern nigeria. *The American Economic Review* 1287–1300 (1995).
- [20] Anttila-Hughes, J. K. & Hsiang, S. M. Destruction, disinvestment, and death: Economic and human losses following environmental disaster. *working paper* (2012).
- [21] Barro, R. J. & Sala-i-Martin, X. *Economic Growth, 2nd edition* (The MIT Press, 2003).
- [22] World Bank Group. *World Development Indicators 2012* (World Bank Publications, 2012).
- [23] Feenstra, R. C., Inklaar, R. & Timmer, M. The next generation of the penn world table. Tech. Rep., National Bureau of Economic Research (2013).
- [24] Matsuura, K. & Willmott, C. J. Terrestrial air temperature and precipitation: Monthly and annual time series (1900 - 2010) v 3.01 (2012).
- [25] CIESIN, I. Gridded population of the world (gpw), version 3.
- [26] Hsiang, S. M. & Jina, A. The causal effect of environmental catastrophe on long run economic growth. *NBER working paper 20352* (2014).
- [27] Auffhammer, M., Hsiang, S. M., Schlenker, W. & Sobel, A. Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy* **7**, 181–198 (2013).
- [28] Hsiang, S. M., Burke, M. & Miguel, E. Quantifying the influence of climate on human conflict. *Science* **341**, 1235367 (2013).
- [29] Angrist, J. D. & Pischke, J.-S. *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, 2008).
- [30] Granger, C. W. & Newbold, P. Spurious regressions in econometrics. *Journal of econometrics* **2**, 111–120 (1974).
- [31] Wooldridge, J. *Econometric analysis of cross section and panel data* (The MIT press, 2002).
- [32] Bertrand, M., Duflo, E. & Mullainathan, S. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* **119**, 249–275 (2004).
- [33] Pindyck, R. S. Climate change policy: What do the models tell us? *Journal of Economic Literature* **51**, 860–872 (2013).
- [34] Dell, M., Jones, B. F. & Olken, B. A. What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature* **52** (2014).
- [35] Dell, M., Jones, B. F. & Olken, B. A. Climate change and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* **4**, 66–95 (2012).
- [36] McIntosh, C. T. & Schlenker, W. Identifying non-linearities in fixed effects models. *UC-San Diego Working Paper* (2006).

- [37] Lobell, D. B., Schlenker, W. & Costa-Roberts, J. Climate trends and global crop production since 1980. *Science* **333**, 616–620 (2011).
- [38] Hsiang, S., Meng, K. & Cane, M. Civil conflicts are associated with the global climate. *Nature* **476**, 438–441 (2011).
- [39] Romer, C. D. & Romer, D. H. The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *The American Economic Review* 763–801 (2010).
- [40] Cerra, V. & Saxena, S. C. Growth dynamics: the myth of economic recovery. *The American Economic Review* **98**, 439–457 (2008).
- [41] Summers, R. & Heston, A. The Penn World Table (Mark 5): An expanded set of international comparisons, 1950–1988. *The Quarterly Journal of Economics* **106**, 327–68 (1991).
- [42] Hsiang, S. M. & Meng, K. C. Reconciling disagreement over climate–conflict results in africa. *Proceedings of the National Academy of Sciences* **111**, 2100–2103 (2014).
- [43] Schlenker, W. & Lobell, D. Robust negative impacts of climate change on African agriculture. *Environmental Research Letters* (2010).
- [44] Fisher, A. C., Hanemann, W. M., Roberts, M. J. & Schlenker, W. The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *The American Economic Review* **102**, 3749–3760 (2012).
- [45] O’Neill, B. C. *et al.* A new scenario framework for climate change research: the concept of shared socioeconomic pathways. *Climatic Change* **122**, 387–400 (2014).
- [46] UN Department of Economic and Social Affairs Population Division. *World population prospects: The 2012 revision* (United Nations, 2013).
- [47] Revesz, R. L. *et al.* Global warming: Improve economic models of climate change. *Nature* **508**, 173–175 (2014).
- [48] Hsiang, S. M. Visually-weighted regression. *SSRN working paper 2265501* (2012).