Labor Productivity and Temperature

Geoffrey Heal, Jisung Park & Nan Zhong

August 25, 2017

Columbia Business School, Harvard University and Xiamen University

Abstract

We develop a theoretical framework suggesting a variety of qualitatively different relationships between temperature and productivity. The exact relationship depends on the importance of income effects and on the way in which income effects and temperatures are distributed across regions.

The connection between temperature and productivity implies a connection between temperature and the rate of growth of productivity. If the capital stock is lower than a certain value then the relationship between temperature and the rate of growth of productivity is qualitatively the same as that between temperature and the level of productivity, implying that the long-run levels of productivity are affected.

We test these ideas on three different data sets. First we use a cross-country data set, and find clear evidence for a $\cap$-shaped relationship between temperature and the level of productivity: an increase in temperature raises productivity in cold countries and lowers it in hot ones. However in hot countries with high levels of air-conditioning, an increase in temperature has no impact on productivity. A change in temperature also has no impact on productivity in "straddling" countries, countries which have hot and cold regions sufficiently different in temperature that we would expect the impacts of a temperature shock to have different signs. With this cross-country data set there is also a $\cap-$shaped relationship between temperature and the rate of growth of productivity, statistically less robust than between temperature and the level of productivity.

Our second data set contains GDP and weather data for 286 Chinese prefecture cities for thirteen years. In this case we find a $\cup$-shaped relationship between temperature and productivity. We also find a $\cup$-shaped relationship between temperature and the rate of growth of productivity.

Finally we work with German data, using Länder-level GDP and weather data over the period 1992 to 2015: in this case we find $\cup$-shaped connection between GDP per capita and temperature, and no clear relationship between temperature and the rate of growth of productivity.

1 Introduction

The effect of temperature on economic performance has been the topic of a growing literature in the last decade (for a recent survey see Heal and Park [2016]). Two concerns motivate this: one is the possible impact of the rising global temperatures associated with climate change, and the other is an interest in understanding the drivers of differences in economic performance between nations, particularly differences between rich (and generally temperate or cold) countries and developing (and generally hot) countries. Dell et al. [2008] (Dell, Jones and Olken: DJO from now on) studied the connections between temperature and performance in a variety of contexts in a series of highly original papers, and found clear statistical connections but did not posit nor test for a specific micro-founded model of the temperature-productivity relationship. Heal and Park [2013] suggested such a model, one in which the connections are based in human physiology: it is a well-established result in physiology that success in task performance depends on temperature. There is an optimal temperature for most tasks, in the low 70s F, and performance falls away on either side of this, as shown in figure 1.1. This figure comes from Seppanen et al. [2006], which summarizes an extensive literature on temperature and task performance.

Figure 1.1: Task Performance vs Temperature

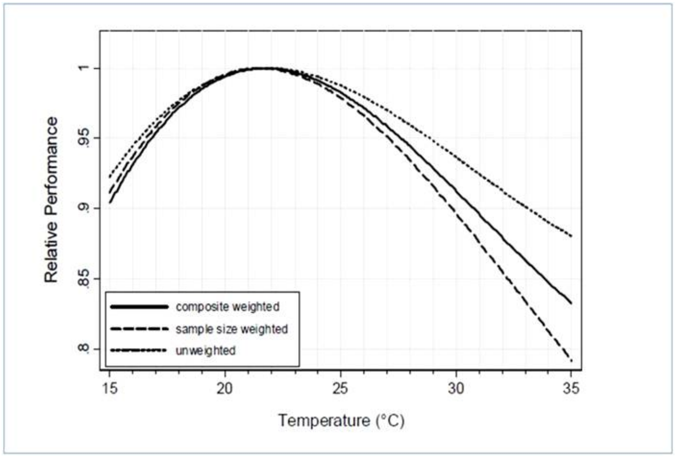

The results summarized in figure 1.1 are based on the study of behavior in a laboratory setting: experimental subjects perform standardized tasks in a carefully controlled environment in which their only concern is the performance of the tasks. To understand whether these results have practical significance, we need to know how people behave in the much more complex setting of a typical economic environment, where an individual is choosing how much labor to offer in response to personal preferences, financial incentives and contractual obligations, and an employer is at the same time choosing how much labor to employ given the wage rate and employee productivity. Can we expect that once employees and employers have made their best choices and a market equilibrium is established, it will still be the case that across economic equilibria productivity rises to a peak with temperature and then falls?

We will show here that this depends on the importance of income effects on the supply of labor. Heal and Park [2013] assumed quasi-linear preferences, which imply that there are no income effects, and showed that in this case productivity in a market situation (as opposed to the laboratory) will generally also be a single-peaked function of temperature, qualitatively similar to figure 1.1. Here we show that allowing for income effects makes matters more far complex and allows a wide range of possible outcomes. We then validate this finding with empirical studies based on cross-country data and data from China, Germany and the US, supplemented by data on stocks of air conditioning equipment in hot countries.

DJO not only suggested a connection between temperature and productivity, but also a possible connection between temperature and the rate of growth of productivity. We pursue this point here and show by embedding our results in a Solow growth model that any short-run connection between temperature and productivity will under certain conditions imply a qualitatively similar connection between temperature and the rate of growth of productivity and so between temperature and productivity in the long run.

2 The Theory of Labor supply, Productivity and Temperature

Figure 2.1 develops the intuition behind our results. It shows a labor supply curve bending backward because of income effects. There are two demand curves, one at high wages corresponding to a cold high-income country and another at lower wages corresponding to a hot, low-income country. In each case we consider the movement of the demand curve in the event of a temperature increase. Following figure 1.1 we assume that a rise in temperature in the cold rich country raises productivity, and we further assume that an increase in labor productivity raises demand for labor at a given wage, moving the demand curve outward (as shown by the arrow in the figure) and so with the labor supply curve shown in figure raising the wage and lowering the amount of labor supplied at equilibrium. Conversely a temperature increase in the hot country lowers productivity (corresponding to the right had side of figure 1.1), lowering the demand for labor at any wage rate, and so lowering the wage rate and the supply of labor. Putting these observations together, we see that an increase in temperature at low temperatures lowers the supply of labor, and also that an increase at high temperatures lowers the labor supply, implying a labor supply vs temperature curve as in figure 2.2, with labor supply monotonically declining in temperature. Making the reasonable assumption that output is an increasing function of labor supply, this implies that output is also a monotonically declining function of temperature. Productivity is generally measured as output

Figure 2.1: Labor Supply and Temperature

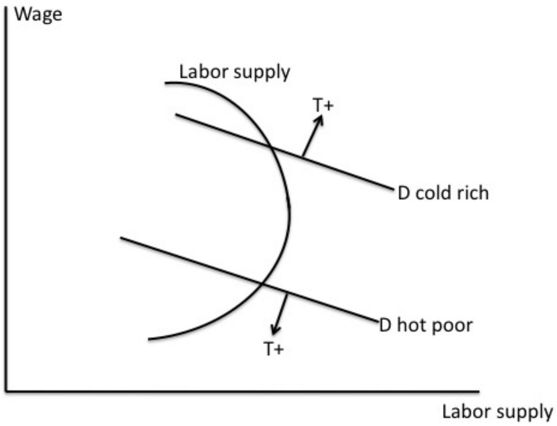

per member of the labor force (not adjusted for hours worked), so the denominator in productivity will be constant independent of temperature and so the productivity-temperature relationship will also have the form of figure 2.2. In fact, as we show next, a wide range of possible outcomes are possible when we take into account all of the factors that determine equilibrium in the labor market.

Equipped with this understanding, we revisit the relationship between productivity and temperature. We do this in several contexts. We use the multi-country panel data set used by Heal and Park and by DJO: we also use a panel data set from 286 Chinese prefectures covering 8 years and a data set covering German Länder since 1992. We find empirical support for several of the wide range of temperature-productivity relationships that the theory suggests are possible.

2.1 A Formal Model of Temperature and Productivity

We begin with a static model of the labor market. Let $Y(T)$ be the total output, T temperature, L the equilibrium labor supply, w the wage rate and $P(T)$ labor productivity as a function of temperature. Clearly total output is given by the amount of labor employed multiplied by the productivity of labor, or $Y(T) = L(w)P(T)$, so that the change in output with respect to

Figure 2.2: Labor Supply Declining in Temperature

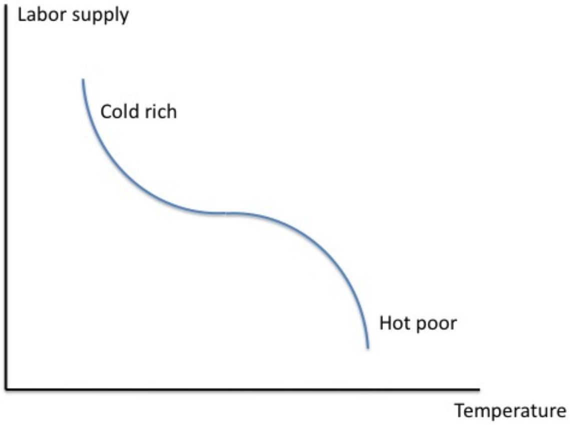

temperature is

$$\frac{\partial Y}{\partial T} = P\left(T\right)\frac{\partial L}{\partial T} + L\left(w\right)\frac{\partial P}{\partial T} \tag{2.1}$$

As labor supply is a function of the wage rate we write

$$\frac{\partial L}{\partial T} = \frac{\partial L}{\partial w}\frac{\partial w}{\partial T}$$

Making the assumption that labor is paid its marginal productivity, we can equate w and P so that $\partial w/\partial T = \partial P/\partial T$ and we can write

$$\frac{\partial Y}{\partial T} = \frac{\partial P}{\partial T}\left\{\frac{\partial L}{\partial w}P + L\right\} \tag{2.2}$$

On the right hand side of this expression both $\partial P/\partial T$ and $\partial L/\partial w$ can change sign as temperature increases (as indicated by figures 1.1 and 2.1). This means that there are multiple possibilities for the behavior of $\partial Y/\partial T$.

1. Forward-sloping labor supply in all regions: If $\partial L/\partial w > 0$ then $sign\,\partial Y/\partial T = sign\,\partial P/\partial T$ and $\partial Y/\partial T$ is positive at low temperatures and negative at high, giving an inverted U shape similar to figure 1.1, as found in the earlier paper by Heal and Park in the study of the cross-country data set.

2. Backward-bending labor supply in all regions: If $\partial L/\partial w < 0$ and is sufficiently negative that $\left\{\frac{\partial L}{\partial w}P + L\right\} < 0$ then $sign\,\partial Y/\partial T = -sign\,\partial P/\partial T$

Figure 2.3: U and Inverted U Relationships Between Productivity and Temperature

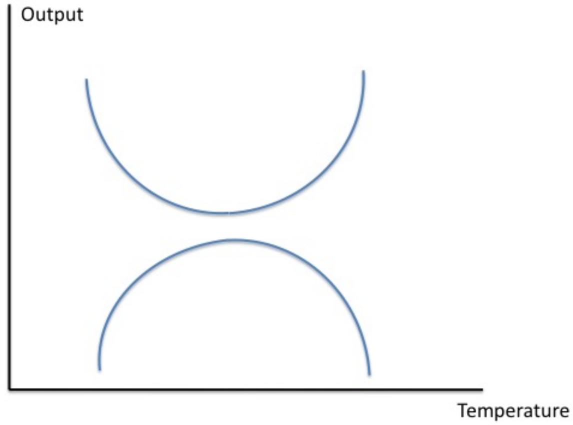

and $\partial Y/\partial T$ is negative at low temperatures and positive at high, giving a function relating temperature and productivity that has a regular U shape. For cases 1 and 2 see see the lower and upper curves respectively in figure 2.3.

3. If $\partial L/\partial w > 0$ at low temperatures and $\partial L/\partial w < 0$ at high temperatures and $\left\{\frac{\partial L}{\partial w}P + L\right\}$ has the same sign as $\partial L/\partial w$ then $\partial Y/\partial T$ is uniformly non-negative, and is zero when $\partial L/\partial w = 0$, with a shape similar to that in figure 2.4.

4. If $\partial L/\partial w < 0$ at low temperatures and $\partial L/\partial w > 0$ at high temperatures and $\left\{\frac{\partial L}{\partial w}P + L\right\}$ has the same sign as $\partial L/\partial w$ then $\partial Y/\partial T$ is uniformly non-positive, zero when $\partial L/\partial w = 0$, with a shape similar to that in figure 2.2: this is the case discussed in the context of figure 2.1.

5. If $\partial L/\partial w$ changes sign but the sign of $\left\{\frac{\partial L}{\partial w}P + L\right\}$ does not change, this is because L is large relative to $\frac{\partial L}{\partial w}P$ in which case $\left\{\frac{\partial L}{\partial w}P + L\right\} > 0$ and we have case 1 again.

This analysis shows us how total output Yvaries in response to changes in temperature T. As noted above, aggregate productivity data measures output

Figure 2.4: Labor Supply Increasing in Temperature

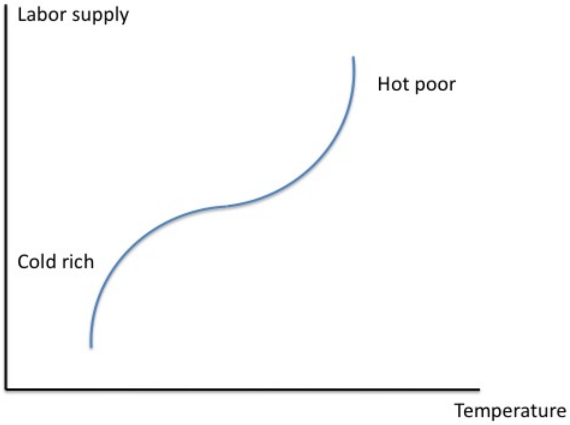

per capita where the numerator is total output Y and the denominator is total population or total labor force: the denominator does not reflect changes in the number of hours worked. So during variations in Y as temperature changes, the denominator in the productivity ratio is constant and hence productivity moves exactly as output Y. So it may show a regular U-shaped relationship with temperature, an inverted U-shaped relationship, or be monotonically increasing or decreasing.

In the empirical works reported in the following sections, we will see that the results from the cross-country panel data set using annual data at the national level are consistent with the inverted U-shape of case 1 above in which $\partial L/\partial w > 0$. This is the result found by Heal and Park [2013], and implies that all countries are on the forward-sloping parts of their labor supply curves, or at least that the data is somehow dominated by countries on the forward-sloping parts.

Our findings for China are consistent with case 2, a regular U-shape as in the top of figure 2.3, or with case 4, a declining curve as in figure 2.2. This means that at low temperatures $\partial L/\partial w$ is sufficiently negative that $\left\{\frac{\partial L}{\partial w}P + L\right\} < 0$ then $sign\,\partial Y/\partial T = -sign\,\partial P/\partial T$ and $\partial Y/\partial T$ is negative at low temperatures, which means that we have a strongly backwards-bending labor supply curve, indicating that the income effects of a change in the effective wage rate outweigh the substitution effects.

This naturally raises the question: when are income effects likely to be important? To answer this question we review briefly the derivation of a labor supply curve. Let an individual's utility be $U(y,L)$ where y is consumption and L is leisure, so U is increasing in both. Income is given by $y=(K-L)w+A$ where w is the wage rate and K is the total number of hours available for both work and leisure. A is the agent's non-labor income. The consumer problem is

$$Max_L U(y,L),\ y+wL=wK+A$$

and we let $S=wK+A$ be the total income the agent could earn if she devoted all her time to work. This is her wealth. The FOCs are

$$U_L-\lambda w=0,\ U_y-\lambda=0$$

From this we get the Marshallian demand function $x_l(p,w,S)$ for leisure and by solving the expenditure minimization problem we get the Hicksian or compensated demands $h_l(p,w,U)$ for leisure. Using the Slutsky equation we can write

$$\frac{\partial x_l}{\partial w}=\frac{\partial h_l}{\partial w}-\left(\frac{\partial x_l}{\partial S}\right)l$$

Note that $x_l(p,w,S)=x_l(p,w,wK+A)$ from which

$$\frac{dx_l^*}{dw}=\frac{\partial x_l}{\partial w}+K\frac{\partial x_l}{\partial S}$$

We use dx_l^*/dw to stand for the derivative noting that x_l depends on w via two of its arguments and this has to be taken into account when differentiating with respect to w.

Substituting this into the Slutsky equation we have

$$\frac{\partial x_l}{\partial w}=\frac{dx_l^*}{dw}-K\frac{\partial x_l}{\partial S}=\frac{\partial h_l}{\partial w}-L\frac{\partial x_l}{\partial S}$$

so

$$\frac{dx_l^*}{dw}=\frac{\partial h_l}{\partial w}+\left(\frac{\partial h_l}{\partial S}\right)(K-L) \tag{2.3}$$

This expresses the total effect of a wage change on the supply of labor. The first term is negative as it is the own price effect: an increase in the wage rate makes leisure more expensive and reduces consumption of leisure (substitution effect), while if leisure is a normal good the second term is positive representing the income effect of a wage change.

Returning to the role of income effects in the impact of temperature changes on the supply of labor, note that in equation (2.3) the income effect $\partial h_l/\partial S$ is multiplied by $(K-L)$, the number of hours worked. So the more hours are worked, the more important the income effect will be in determining the response of labor supply to a wage change and so to temperature change. Casual empiricism suggests that hours worked on China are greatly in excess of those worked in the US or other industrial countries, suggesting that income effects may be more important in China.

To summarize, we have shown that a single-peaked relationship between task performance and temperature, as routinely discovered in laboratory studies, does not necessarily imply a similar relationship between productivity (as normally measured) and temperature. It may lead to a similar single-peaked relationship, but may also lead to the opposite, a U-shaped relationship, and to various other outcomes. What sort of connection we see will depend on the importance of income effects in the regions being studied.

2.2 A Dynamic Model

We now take our static results and apply them in the context of a Solow growth model (Solow [1956]), to investigate their implications for connections between temperature and the rate of growth of productivity. The simplest possible version of the Solow growth model is:

$$Y = F\left(K, L\right), \ \frac{dK}{dt} = sF\left(K, L\right) - \delta K \tag{2.4}$$

where Y, K, L are respectively output, capital stock and labor force: the labor force is assumed to be constant. $s \in (0,1)$ is the savings rate and $\delta \in [0,1]$ the rate of depreciation of capital. We take the labor force to be constant and assume there is no technical progress, so that increases in output per capita come entirely from increases in the capital stock per capita or variations in the efficiency of labor. Let $\hat{K}$ be the stationary solution of differential equation (2.4): it satisfies

$$sF\left(\hat{K}, L\right) = \delta\hat{K} \tag{2.5}$$

Figure 2.5 illustrates this. The main part of the diagram shows $sF\left(K, L\right)$ and δK, while the top panel shows the rate of change of income (and so also of income per capita) as a function of capital stock, positive for $K < \hat{K}$ and negative if $K > \hat{K}$.

For $K < \hat{K}$, $dK/dt > 0$, and vice versa, so that $\hat{K}$ is a stable stationary solution.

Consider first a temperature spike in a hot country leading to a drop in output by lowering output per capita. (We can think of this as a drop in the effective labor supply.) Then $F\left(K, L\right)$ falls. From (2.4) it is clear that dK/dt drops. If the temperature remains high over an interval of time say $[\tau_1, \tau_2]$ then dK/dt remains lower over that interval. Now note that

$$\frac{dY}{dt} = \frac{dF}{dK}\frac{dK}{dt} = F_K\left\{sF\left(K, L\right) - \delta K\right\} \tag{2.6}$$

The impact of a change in the effective amount of labor on the rate of growth of output is

$$\frac{\partial}{\partial L}\left(\frac{dY}{dt}\right) = F_{K,L}\left\{sF\left(K, L\right) - \delta K\right\} + F_K sF_L \tag{2.7}$$

Figure 2.5: Solow Growth Model

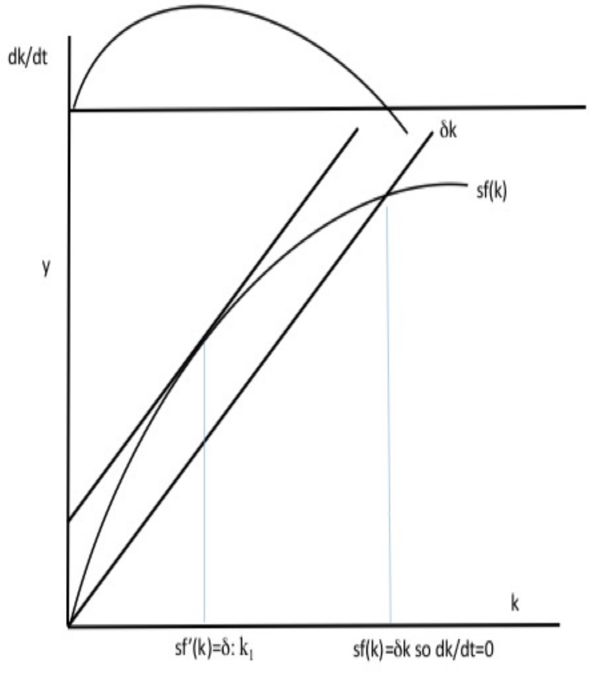

Assuming that $F_{K,L} \geq 0$ so that an increase in the amount of one factor does not lower the marginal product of the other, and assuming also that $K < \hat{K}$, then this is positive so that dY/dt, the rate of growth of output per capita, also falls when there is an increase in temperature and a decrease in effective labor and is also lower over the interval $[\tau_1, \tau_2]$. This would be reflected in a drop in the path of income growth in the upper panel of figure 2.5. Note that when the temperature returns to its normal value, this process is reversed: output per capita rises and so does its rate of growth.

The next point is to investigate whether as dK/dt is lower over $[\tau_1, \tau_2]$ because of the higher temperature then K is also lower than it otherwise would be after this interval: consider K_T for $T > \tau_2$:

$$K_T = \int_0^T \frac{dK}{dt}(\tau)\, d\tau = \int_0^{\tau_1} \frac{dK}{dt}(\tau)\, d\tau + \int_{\tau_1}^{\tau_2} \frac{dK}{dt}(\tau)\, d\tau + \int_{\tau_2}^{T} \frac{dK}{dt}(\tau)\, d\tau \quad (2.8)$$

The second term on the right here is reduced by the high temperature and the first term is unchanged.

With respect to the third term on the right of the integral,

$$\frac{\partial}{dK}\left(\frac{dK}{dt}\right) = sF_K - \delta \quad (2.9)$$

which is positive or negative according as $K \gtreqless K_1$ where $sF_k(K_1) = \delta$ (see figure 2.5).[1] As K is lower after τ_2 then for $K < K_1$, dK/dt is smaller, and for $K > K_1$ it is larger. So for $K < K_1$ the rate of growth of capital per capita and thus of output per capita will remain lower after τ_2. In this case the path of income growth in figure 2.5 will be moved downwards from the time of the temperature spike onwards. So growth and productivity will be lower after the shock and its reversal than they would otherwise have been.[2]

The conclusion here is that a drop in the effective labor supply, for example caused by an increase in temperature that reduces productivity, will lead to lower levels of output per capita and if $K < K_1$ will also lead to lower rates of growth of output per capita, and these will persist over time. So in a hot country with a low capital stock an increase in temperature may lower the rate of growth of output per capita over some period of time. In a country with a higher value of the capital stock ($K > K_1$) it will still be true that a temperature spike lead to a drop in productivity, and to a short-term drop in the rate of growth of productivity over the period for which the temperature is higher, but in the longer term this drop in the rate of growth of productivity may be reversed.

All the above arguments go into reverse if the labor supply is increased, so an increase in temperature that raises effective labor supply (increases productivity) will (for $K < K_1$) increase the rate of growth of productivity too, and for $K > K_1$ there will be a short-term increase in the rate of growth of productivity

---

[1] As $s < 1$, K_1 is less than the "golden rule" value of the capital stock.

[2] Note that the stationary value $\hat{K}$ has not changed so that asymptotically the value of K is as before. By the same reasoning the asymptotic value of income per capita has not changed, but it will be reached more slowly.

which might be reversed over the longer term. So if we observe, across countries or regions, an inverted U relationship between temperature and productivity then we may also observe such a relationship between temperature and the rate of growth of productivity, depending on the endowments of capital: the latter will be a consequence of the former. (This is assuming that $K < \hat{K}$, which is certainly the case if $K < K_1$: otherwise the relationship is reversed in going from levels to rates of change, so an inverted U relating temperature to productivity will translate into a regular U relating temperature to the rate of growth of productivity.)[3]

We can summarize this in the following

Proposition 1. *Let* $K_{\tau_1} < \hat{K}$. *Then a decrease (increase) in the effective labor supply (indicated by output per capita) relative to a pre-existing solution to equation (2.4) over the interval* $[\tau_1, \tau_2]$ *causes a drop (rise) in both the level and the rate of growth of productivity over that interval. If in addition* $K_{\tau_1} < K_1$ *then the level and rate of growth of productivity remain permanently lower (higher) relative to the pre-existing solution.*

What this is telling us is that if the capital stock is "low enough" (below K_1) then a change in temperature that causes an immediate drop in productivity will in addition lead to a drop in the rate of growth of productivity and the level of productivity in the long run. Conversely a temperature change that causes an increase in productivity will lead to the opposite effects. So the short-run effects on productivity discussed in points 1 through 5 of section 2.1 will be replicated in the rate of growth of productivity and in the long-run level of productivity.

There is a clear intuition behind this: a drop in productivity leads to a drop in output and in investment, and this leads to a drop in capital accumulation and so in the rate of growth of output. So short-term temperature spikes can have long-run impacts.

$K < K_1$ is a sufficient condition for this to happen, but is not necessary. But if $K > K_1$ then from equation (2.9) the rate of capital accumulation is increased after a drop in productivity and this could offset the initial decline in accumulation.

3 A Cross-Country Panel

In the remainder of the paper we test the ideas that emerged from the theoretical analysis of the previous section. We begin with a cross-country panel data set, similar to that used by DJO, but augmented by a set of data on the penetration of air-conditioning units in hot countries. We then switch to a more micro level and look at data on GDP per capita and temperature from 286 Chinese prefectures over an 8 year period, and follow this with studies at the county level in the US and at the Land level in Germany. We find evidence for the temperature-productivity relationships shown in figure 2.3, but (at least for this

---

[3] The sign of the first term on the right of equation 2.2 changes from positive to negative in this case, and for small values of $F_K s F_L$, the derivative $\partial/\partial L\,(dY/dT)$ can be negative.

data) not for those in figures 2.2 and 2.4. The cross-country data set suggests an inverted U relationship as in the lower part of figure 2.3, while the Chinese and German data suggests a regular U as in the upper part of figure 2.3.

3.1 Data

3.1.1 Climate Data

Annual average temperature and precipitation data (in degrees centigrade and mm respectively) at the country level are taken (as in Dell et al. [2009]) from Terrestrial Air Temperature and Precipitation: 1900-2006 Gridded Monthly Time Series, Version 1.01 (Matsuura and Willmott [2007]), and are weighted by population.[4] Population weighting ensures that the country average picks up the most economically relevant climate realizations. If, for example, most of a country's population lives in its southern region, one might expect most of its economic activity to take place there as well. In that case, taking a geographic average temperature might be misleading, particularly if that country has sparsely populated areas in extreme climates (e.g. Russia and Siberia, Canada and its arctic areas, the United States and Alaska).

[5]

3.1.2 International Economic Data

We use income data from the UN National Accounts, which records measures for income, population, and capital stock from 1970 to 2011. Real GDP per capita is measured in terms of 2005 USD$ using Laspreyes constant prices. Note that, unlike Dell et al. [2008], our analysis uses economic variables from the UN National Accounts (as opposed to the Penn World Tables or the World Development Indicators). This is motivated by the desire to mitigate the bias introduced by adjustments for purchasing power parity (PPP) embedded in the Penn World Table data, as well as to test whether the temperature shock-per capita income relationship found in DJO can be replicated using different income data. As Deaton and Heston (2010) note, the manner in which Penn World Table income data incorporate successive adjustments to PPP may introduce issues of consistency, especially for non-OECD countries, whose implied growth rates exhibit spurious correlation if taken on short enough time intervals. Like Dell et al. (2012), we drop countries for which either the climate or GDP data do not exist or the panel data does not extend for at least 20 years. This leaves an unbalanced panel of 134 countries, most of which have economic data for the period 1970-2011, and a total of 6,101 observations.

---

[4] We take the daily average temperature, and average it over the year while weighting by population.

[5] Ideally, one would use a less aggregated measure of temperature, for instance, cooling and heating degree days (CDD, HDD). CDD and HDD data, though available at more localized levels in OECD countries, was not readily available for the cross-country dataset used here.

Trade data on air conditioning import value is taken from the United Nations COMTRADE database, a subset of the World Integrated Trade Solution data set, beginning in 1962. The choice of imports as opposed to exports is motivated by the fact that, for any given country, imports are usually recorded with more accuracy than exports since imports generally generate tariff revenues while exports do not. The unit we use to measure air conditioning imports is trade value (in million dollars) of "air conditioning machines". The reason we use this measure instead of a raw quantity measure is because countries do not report the same unit to report trade quantity; some countries report trade quantity in kilograms, other in number of items, others still in pounds. Trade value seems a more consistent measure of actual air conditioning achieved. We take this flow measure of import value and create an index of cumulative air conditioning equipment per capita by year. In 1995, for instance, expenditures on air conditioning equipment (proxied by cumulative imports of air conditioning equipment since 1960) ranged from \$0 per capita (most Sub-Saharan African countries, for example) to \$161 per capita (Kuwait) .

3.2 Statistical Model

Given our model, and the literature on task performance under thermal stress, we expect the underlying relationship between output per capita and temperature to take the following form:

$$y_{it} = f(T_{it}) + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{3.1}$$

where $f(T_{it})$ is some potentially non-linear function of temperature, K_{it} is a vector of "capital stock variables", which in principle may include all country-specific, time-varying contributors to income per capita, θ_i denotes time-invariant country-specific factors such as natural resource endowments or institutions, γ_t represents year-specific common shocks (e.g. global recessions), and ϵ_{it} is a country-year specific error term. A more structurally restrictive version of this equation may assume a single-peaked $\cap-$shaped or conversely a $\cup$-shaped relationship between income and temperature, as suggested in the models of section 2:

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{3.2}$$

In this formulation the coefficients β_1 and β_2 are respectively positive and negative or negative and positive depending on whether we have a $\cap$ or $\cup$-shaped relationship between productivity and temperature. We will also use a non-parametric formulation that allows for either of these relationships and also for montonically increasing or decreasing functions. We also test for the increasing and decreasing cases parametrically by using cubic specifications of the productivity-temperature connection.

The simplest way to estimate this relationship is to run a cross-sectional OLS regression of the following form, where δ_i denotes a country-specific residual:

$$y_i = \alpha + \beta_1 T_i + \beta_2 T_i^2 + \delta_i.$$

Following this basic estimation strategy, Horowitz [2001] finds that a one degree increase in temperature is associated with -8.5% change in GDP per capita.[6]

We confirm that there exists a strongly negative cross-sectional relationship between temperature and income, particularly in countries where population-weighted average temperatures are above 20°. Of course, a key limitation of the existing cross-sectional analyses is that they may miss country-specific factors such as natural resource endowments or institutions. Researchers often point to the starkly different fortunes of North and South Korea as indicative of the crucial role of institutional factors.[7]

It is worth noting, furthermore, that previous studies which emphasize the monotonic cross-sectional relationship between temperature (latitude) and income (growth) may miss a significant component of the relationship, due to the limited number of cold countries in most samples. For example, in our sample there are only 5 countries which have annual average temperatures below 5° Celsius, even though a much larger number of countries have regions with very cold climates. More research is needed to uncover the temperature-income gradient within countries, especially those that have significant cold regions. At the very least, the temperature-income gradient in the cross-section provides us with an upper bound for any contemporaneous impact of temperature on income.[8]

The panel nature of our dataset allows us to control for time-invariant, country-specific unobservables that may influence income per capita: for instance, institutions or natural resource endowments (θ_i), and average climate ($\bar{T_i}$). In addition, we control for country-specific factors that may be changing over time by adding measures of country-specific capital stock directly. Using data from the Penn World Tables, we control for physical capital (log capital stock per capita) and human capital accumulation (in the form of an index).[9]

---

[6]

Dell et al. [2009] and Nordhaus [2006] represent marginal improvements on this regression by using disaggregated data at the municipality and grid-cell levels respectively. Both find strong, statistically significant negative relationships between temperature and income in a cross-section, of slightly smaller magnitude. In Nordhaus' case, the finding is of a strongly single-peaked relationship.

[7] Selection via migration to more favorable climates is also something that cross-sectional correlations cannot account for. Cross-sectional analyses may also be sensitive to period-specific idiosyncrasies. If the data is from a year in which there was a global recession, it is unclear to what extent this globally correlated shock is affecting the underlying relationship.

[8] Selective migration based on the intensity of preferences for climate amenities (or adaptive capacity) notwithstanding.

[9] Both variables are taken from the Penn World Tables, version 8.0 (Heston et al. [November 2012].

One way to think of this is that we are identifying the impact of hotter or colder than average years for a particular country on that country's total output, controlling for all sources of variation in income per capita apart from annual weather fluctuations. By utilizing the "within-group" variation in GDP with respect to temperature, we can interpret an association between temperature fluctuations and income fluctuations as causal. As a number of other studies note (Hsiang et al. [2013], Auffhammer et al. [2013]), such annual fluctuations in weather variables can be considered essentially random.

Thus, our preferred regression framework utilizes country- and year-fixed effects, as well as country-specific trends in physical and human capital accumulation:

$$y_{it} = f(T_{it}) + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{3.3}$$

This empirical specification, while utilizing within-country variation, is not immune to issues of spurious correlation. If variation in temperature is correlated with variation in capital stock variables, we may be attributing too much of the variation in income levels to temperature shocks. We discuss the issue of potential spurious correlation and our attempts to adjust for this in section 3.4.

It is worth noting that our identification strategy relies on the hypothesis that variations in temperature from year to year in a given country (short-term variations, inter-annual variability) lead to the same sort of economic responses as variations in temperature across countries that are maintained over long periods of time (climate variation). In other words, as a country experiences say a 2 degree C hotter than average year, it reacts in the same way as a country that is on average 2 degrees C hotter, conditional on compositional characteristics (agricultural value-added, air-conditioning penetration, etc). Short and long-run responses are, as a matter of simplification, treated as if they are the same: there is only one temperature-income relationship rather than several that depend on the time scale. The various papers by DJO use the same assumption (Dell et al. [2008, 2009]), as does Hsiang [2010]. An alternative is that this is not true, and that countries that are maintained at high temperature over long periods of time can adapt to these in ways that take time and investment and to some degree mitigate the impact of temperature, while countries that experience a temperature shock that is not expected to last do not adapt. In this case we would expect to see more response to short-run (year to year) fluctuations than to long-run differences, and our coefficients could overstate the impact of temperature differences that are maintained over long periods of time. We touch on this in section 3.4.

3.3 Results

We begin by estimating a quadratic relationship between temperature and income per capita. Table 1 presents the coefficients from estimating equation (3.2) above.[10] We allow for the possibility that temperature may affect GDP

---

[10] All regressions contain year and country fixed effects.

Table 1: Productivity vs Quadratic Temperature

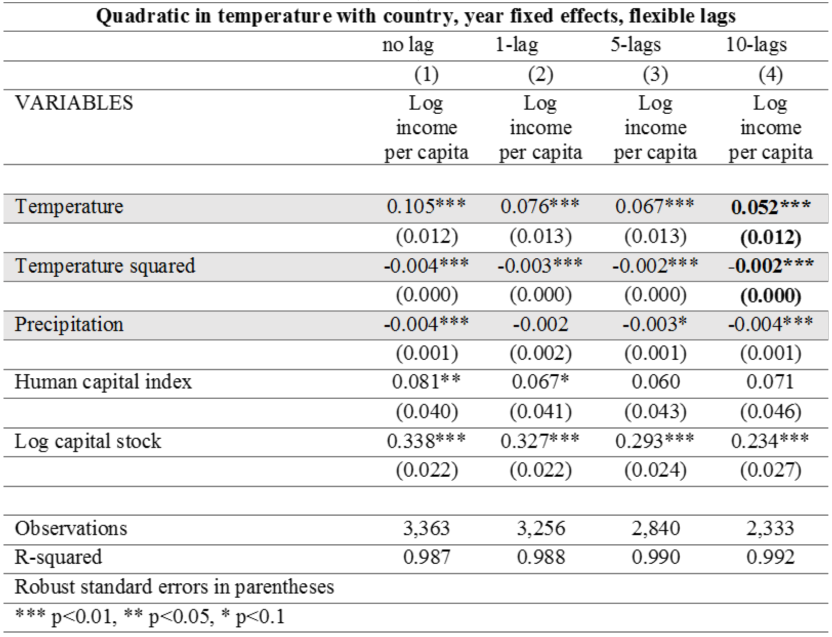

Quadratic in temperature with country, year fixed effects, flexible lags

	no lag	1-lag	5-lags	10-lags
	(1)	(2)	(3)	(4)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita
Temperature	0.105***	0.076***	0.067***	**0.052*** **
	(0.012)	(0.013)	(0.013)	**(0.012)**
Temperature squared	-0.004***	-0.003***	-0.002***	**-0.002*** **
	(0.000)	(0.000)	(0.000)	**(0.000)**
Precipitation	-0.004***	-0.002	-0.003*	-0.004***
	(0.001)	(0.002)	(0.001)	(0.001)
Human capital index	0.081**	0.067*	0.060	0.071
	(0.040)	(0.041)	(0.043)	(0.046)
Log capital stock	0.338***	0.327***	0.293***	0.234***
	(0.022)	(0.022)	(0.024)	(0.027)
Observations	3,363	3,256	2,840	2,333
R-squared	0.987	0.988	0.990	0.992

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

with a time lag, by allowing for 1, 5, and 10 lags. Allowing for lagged impacts controls for the potential for serial correlation in the shocks, due, for example, to ENSO climate cycles, usually with a periodicity of 4-8 years. Allowing for lags also helps us to come closer to isolating the physiological "effective labor supply" channel as separate from other long-lived investment impacts.[11]

Our coefficient of interest, therefore, is the contemporaneous impact of temperature in year t on income in year t. The table suggests a significant, concave relationship temperature (degrees C) and log income per capita, allowing for 0 to 10 lags. Whether or not we allow for lagged effects, the concave relationship persists. The implied "optimal" temperature is in the range of 15° and 20° Celsius across all specifications, consistent with the medical literature.[12]

Next, we consider a more flexible functional relationship between temperature and GDP (3.1), by creating dummies for a range of average temperature bins and allowing for piecewise linear relationships within each bin. We report the results for a 5-bin classification, where countries are classified into "very hot" (average annual temperature above 25°C), "hot" (20-25°C), "temperate"

---

[11] While we do not discuss long-term impacts of climate shocks here, we note that, in principle, a large enough thermal shock could have impacts that persist for a very long time. For example, a heat wave in utero may affect income in one's twenties and thirties.

[12] These ranges are likely shifted downward systematically relative to the optimum implied by lab studies, primarily due to the fact that our data is in annual averages, which counts nighttime temperatures as well as daytime temperatures.

Table 2: 5 Temperature Bins

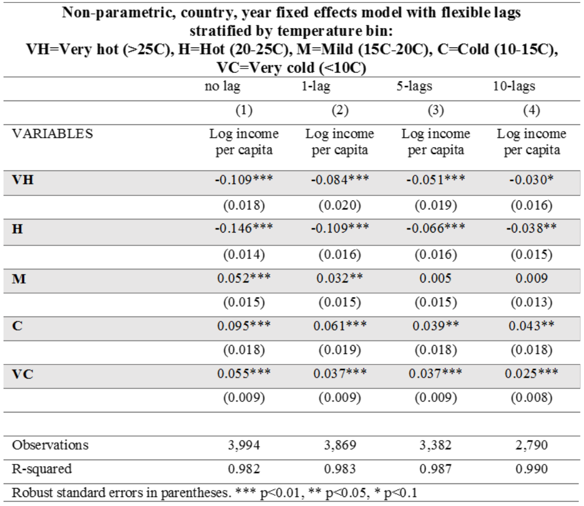

Non-parametric, country, year fixed effects model with flexible lags stratified by temperature bin:
VH=Very hot (>25C), H=Hot (20-25C), M=Mild (15C-20C), C=Cold (10-15C), VC=Very cold (<10C)

	no lag	1-lag	5-lags	10-lags
	(1)	(2)	(3)	(4)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita
VH	-0.109***	-0.084***	-0.051***	-0.030*
	(0.018)	(0.020)	(0.019)	(0.016)
H	-0.146***	-0.109***	-0.066***	-0.038**
	(0.014)	(0.016)	(0.016)	(0.015)
M	0.052***	0.032**	0.005	0.009
	(0.015)	(0.015)	(0.015)	(0.013)
C	0.095***	0.061***	0.039**	0.043**
	(0.018)	(0.019)	(0.018)	(0.018)
VC	0.055***	0.037***	0.037***	0.025***
	(0.009)	(0.009)	(0.009)	(0.008)
Observations	3,994	3,869	3,382	2,790
R-squared	0.982	0.983	0.987	0.990

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

(15-20°C), "cold" (10-15°C), and "very cold" (10°C and below). The results suggest a single-peaked relationship, with the implied peak again occurring somewhere between 15° and 20° Celsius (Table 2). A hotter than average year is associated with lower than average output per capita in countries with average annual temperatures above 20°C (during 1950-2005), while a positive temperature shock of similar magnitude is associated with higher output per capita in cooler countries (average annual temperatures below 20°C). There is higher variance among very hot countries, but the overall pattern of negative effects of heat shocks in warm climates and positive effects of heat shocks in cooler climates is noticeable. This pattern persists across various bin classifications (e.g. three climate bins as opposed to five, see table 3).

The magnitude of temperature-related output fluctuations implied by these regressions is large. Very hot countries such as Thailand, India, and Nigeria suffer negative output shocks on the order of 3-4% per capita GDP per degree Celsius. Very cold countries such as the UK, Canada, Norway, and Sweden have significantly higher output in warmer years (and lower output in colder years). These effect sizes are consistent with the emerging literature, and well within the upper bounds signified by cross-sectional studies. For example, looking at 28 Caribbean countries, Hsiang [2010] finds large contemporaneous impacts of temperature shocks on output which ranges from negligible in some to over -6% per degree C in others. The implication seems to be that a quadratic (concave) relationship between temperature and income per capita is a good approximation of the underlying relationship, controlling for time-invariant factors such as institutions and natural resource endowments.

Table 3: 3 Temperature Bins

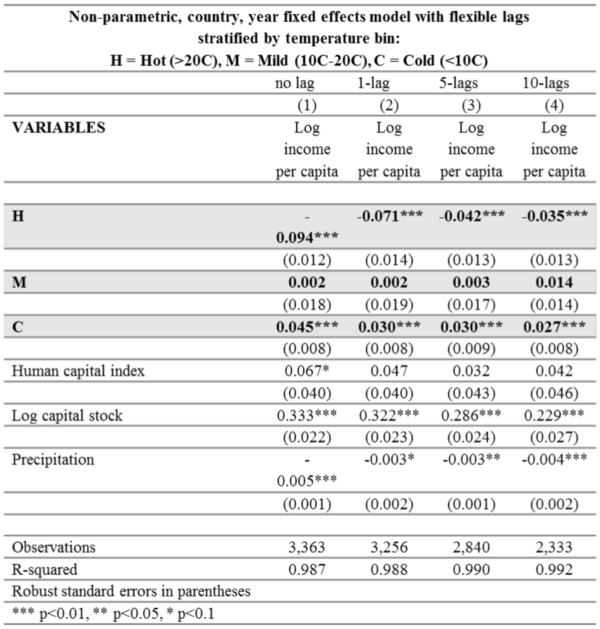

Non-parametric, country, year fixed effects model with flexible lags stratified by temperature bin:
H = Hot (>20C), M = Mild (10C-20C), C = Cold (<10C)

	no lag	1-lag	5-lags	10-lags
	(1)	(2)	(3)	(4)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita
H	**-0.094*** **	**-0.071*** **	**-0.042*** **	**-0.035*** **
	(0.012)	(0.014)	(0.013)	(0.013)
M	**0.002**	**0.002**	**0.003**	**0.014**
	(0.018)	(0.019)	(0.017)	(0.014)
C	**0.045*** **	**0.030*** **	**0.030*** **	**0.027*** **
	(0.008)	(0.008)	(0.009)	(0.008)
Human capital index	0.067*	0.047	0.032	0.042
	(0.040)	(0.040)	(0.043)	(0.046)
Log capital stock	0.333***	0.322***	0.286***	0.229***
	(0.022)	(0.023)	(0.024)	(0.027)
Precipitation	-0.005***	-0.003*	-0.003**	-0.004***
	(0.001)	(0.002)	(0.001)	(0.002)
Observations	3,363	3,256	2,840	2,333
R-squared	0.987	0.988	0.990	0.992

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

3.3.1 Robustness Checks for Omitted Variables and Spurious Correlation

We have established a single-peaked relationship between temperature and output per capita for our cross-country data set, and posited that this arises in part from the physiological factors discussed in earlier sections. However there are of course alternative mechanisms which could lead to this relationship. We know for example that the connection between crop yield and temperature is highly non-linear, with yields increasing in temperature up to a point and then falling rapidly (Schlenker and Roberts [2006]). This suggests that looking across agricultural societies, we could find a single-peaked connection between temperature and output. One would not expect this relationship to persist across industrial countries, but it could be an explanation for our findings for a part of our sample. However, average agricultural value-added as a proportion of GDP in OECD countries is roughly 3% (over the period 1960-2006), and even in many developing economies less than 10%, suggesting that the effects cannot be totally attributable to decreases in agricultural yield.

There is also evidence to believe that there are negative public health aspects of higher temperatures, working through a diverse range of mechanisms such as the spread of disease vectors and the effects of heat stress on mortality. While the focus recently has been on thermal stress at the high end (Deschenes and Greenstone [2007]), it is also the case that very low temperatures lead to increased mortality, and to a range of health stresses too. All of these explanations are consistent with our findings.

Another concern is the potential for spurious correlation arising from secular but heterogeneous time trends in the temperature data. If some countries were warming (cooling) faster than others during the period of interest, we may incorrectly attribute secular changes in GDP to climate fluctuations.

We attempt to control for potential spurious correlation by allowing for country-specific temperature trends (as opposed to global trends in temperature, which are captured by year fixed-effects in the previous regressions). While controlling for country-specific temperature trends reduces the power of the coefficients on temperature markedly, the resulting point estimates remain consistent with a single-peaked relationship between thermal stress and economic productivity.

3.3.2 The Role of Air Conditioning

Additional evidence strengthens the case for physiological impacts as a key causal mechanism. We test for the impact of air conditioning on the temperature-output gradient, by using data on country-specific air conditioning penetration. Insofar as this may buffer the impacts of thermal stress on labor productivity (as opposed to crop failures, for example), we would expect the sensitivity of income shocks to temperature to be lower in areas with higher levels of AC.

We examine whether access to AC attenuates the effect of thermal stress at high temperatures, working with countries that have above average annual temperatures. Because country-specific data on AC penetration per capita is not readily available, we construct a measure of AC penetration per capita by imputing the value of AC equipment imports for each country in our data set. The trade data is taken from the United Nations COMTRADE database, a subset of the World Integrated Trade Solution data set. In 1995, for instance, expenditures on air conditioning equipment (proxied by cumulative imports of air conditioning equipment since 1960) ranged from $0 per capita (most Sub-Saharan African countries, for example) to $161 per capita (Kuwait). Detailed descriptions of air conditioning penetration per capita are presented in the Appendix.

Using this data, we stratify the sample of hot countries (countries that we have labered as “very hot,” “hot,” and “mild”) by AC penetration per capita, dividing it into thirds. Table 4 presents the results for two subsets of countries, the top and bottom thirds by AC penetration allowing for lagged impacts once again. Consistent with the notion that higher levels of AC dampen the impact of thermal stress on productivity, the subset of countries in the top third by AC penetration show no significant impact of temperature on income per capita, whereas the bottom third show a significant negative impact.

Moreover, it seems that this difference is not being driven wholly by the correlation between air conditioning and other unobservables that are correlated with income. While countries with better access to AC tend to be richer on average, there are also relatively hot and poor countries with high air conditioning

Table 4: Impact of AC

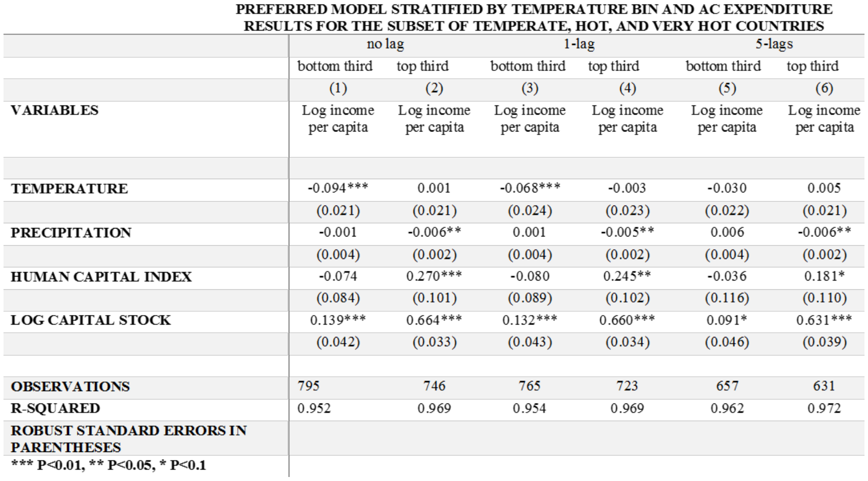

PREFERRED MODEL STRATIFIED BY TEMPERATURE BIN AND AC EXPENDITURE
RESULTS FOR THE SUBSET OF TEMPERATE, HOT, AND VERY HOT COUNTRIES

	no lag		1-lag		5-lags	
	bottom third	top third	bottom third	top third	bottom third	top third
	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita	Log income per capita	Log income per capita
TEMPERATURE	-0.094***	0.001	-0.068***	-0.003	-0.030	0.005
	(0.021)	(0.021)	(0.024)	(0.023)	(0.022)	(0.021)
PRECIPITATION	-0.001	-0.006**	0.001	-0.005**	0.006	-0.006**
	(0.004)	(0.002)	(0.004)	(0.002)	(0.004)	(0.002)
HUMAN CAPITAL INDEX	-0.074	0.270***	-0.080	0.245**	-0.036	0.181*
	(0.084)	(0.101)	(0.089)	(0.102)	(0.116)	(0.110)
LOG CAPITAL STOCK	0.139***	0.664***	0.132***	0.660***	0.091*	0.631***
	(0.042)	(0.033)	(0.043)	(0.034)	(0.046)	(0.039)
OBSERVATIONS	795	746	765	723	657	631
R-SQUARED	0.952	0.969	0.954	0.969	0.962	0.972
ROBUST STANDARD ERRORS IN PARENTHESES						
***** P<0.01, ** P<0.05, * P<0.1**						

penetration (for instance, Libya). It seems that the vulnerability to thermal stress as implied by access to thermoregulatory capital is not simply a function of "poorness" per se. This is an admittedly crude measure, but points us in the right direction for pressing policy-relevant research on climate adaptation.

3.4 Conclusions From Cross-Country Panel

The cross-country panel data shows a clear $\cap-$shaped relationship between temperature and productivity. Temperature has a significant and robust positive effect on productivity in cold countries and a negative one in hot countries. The impact of AC in hot countries is highly significant and suggests that this reflects the physiological factors discussed in the introduction. There remain however some open questions.

Some of the countries in our sample cover many climate zones: for example the US stretches from northern Alaska, which is arctic, to Puerto Rico, which is tropical. These different regions of the US are clearly on opposite sides of the inverted U curve in figure 1.1. The same is true of China, whereas countries like Belgium and Luxembourg are clearly in a single climate zone. It is certainly possible that in a country like the US, about half of economic activity might be to the left of the peak in figure 1.1 and half to the right, so that in aggregate the country shows no response to temperature in spite of the fact that in each region there is a strong response. How does this affect our outcomes and can we correct for any biases that it might introduce? We tested for this as follows. Table 1 implies that countries reach peak productivity at about 13 degrees C, so we divided countries into two groups, those whose temperatures straddle 13 C and the remainder.[13] We then tested the above models on both groups separately:

---

[13] Straddling countries are the following: Afghanistan, Albania, Bulgaria, China, France,

the results are shown in table 5.

This shows clearly that our earlier results are repeated for the non-straddle countries but that there is no systematic and significant relationship between temperature and productivity for the countries that straddle the optimal temperature so that we can expect the impact of temperature shocks to have opposite signs in different regions. Ideally we would like to break all the straddle countries into homogeneous climate regions but data limitations have so far prevented this. However the regionally disaggregated studies that we report next take a step in this direction.

One additional point that we investigated is the robustness of our results to the specification of the time period. A part of our motivation is an interest in the consequences of climate change, yet the temperature shocks we have studied are all short-term and generated by variations in annual weather patterns. Climate evolves over many years, so it would be of interest to understand if the results we have found continue to hold in the study of longer-term change in temperature. To begin to tackle this question we have rerun the studies described above but using both three year moving averages and also using as the time interval three years. So in this latter case our observations are now the three-year averages of temperature and productivity. This reduces our sample size by a factor of three, but it is still large enough to be of interest. Figure 6 shows the results for three year moving averages: those for three year averages are similar.

The results are presented in table 6, and show no change from the previous cases: we present the results for five temperature bins, as in table 2, and once again the coefficients on hot bins are negative, on cold ones positive, and on the intermediate bin zero. The magnitudes of the coefficients are similar to the one-year case.

4 Chinese Prefectures

We next test the ideas from section 2 on disaggregated data from China. The Chinese economic data comes from the China City Statistical Yearbook, which records annual GDP and population on the level of prefectural city. A prefectural city is an administrative division used in China, which ranks below a province and above a county. By the year of 2013, there were 286 prefectural cities in China. In our study, the dependent variable used is annual real per capita GDP per prefectural city, which is measured in terms of 2005 Chinese Yuan using the GDP deflator. China's GDP is contributed by three major sectors. The primary sector includes farming, forestry, animal husbandry and fisheries. The secondary sector includes construction and industry (mining, manufacturing, provision of electricity, water, and gas). The tertiary sector is the service sector. Besides the overall GDP, we also used the real per capita GDP by each sector as the dependent variable.

Greece, Hungary, Iran, Islamic Rep., Italy, Japan, Korea, Dem. Rep. Korea, Lebanon, Romania, Spain, Syrian Arab Republic, Turkey, United States

Table 5: Straddling and Non-Straddling Countries

	"non-straddle"			"straddle"		
	no lag	1-lag	10-lags	no lag	1-lag	10-lags
VARIABLES	log pc income	log pc income	log pc income	log pc income	log pc income	log pc income
Temp	0.091***	0.065***	0.047***	-0.092	-0.11	-0.12
Temp^2	-0.003***	-0.002***	-0.001***	0.002	0.003	0.004
Precip	-0.007***	-0.004***	-0.005***	0.025***	0.021***	0.012**
Temp 1 lag		0.056***	0.032***		-0.063	-0.037
Temp 2 lag		-0.002***	-0.001**		0.001	0
Precip 1 lag		-0.005***	-0.003**		0.019***	0.012**
$\sum temp$ effects		0.121***	0.235***		-0.173	-0.099
$\sum temp^2$effects		-0.004***	-0.006***		0.004	-0.003

Table 6: 3 Year Moving Averages

	With capital	and	human capital	
	no lag	1-lag	5-lags	10-lags
	(5)	(6)	(7)	(8)
VARIABLES	log pc income	log pc income	log pc income	log pc income
hc	0.071*	0.053	0.025	0.048
k	0.335***	0.326***	0.290***	0.235***
temp_VH	-0.173***	-0.210***	-0.150**	-0.118**
temp_H	-0.099***	-0.088***	-0.091***	-0.075**
temp_M	0.006	0.011	0.010	0.064
temp_C	0.076***	0.038	0.063	0.082*
temp_VC	0.072***	0.043***	0.078***	0.075***

The China weather data contains daily maximum temperature and precipitation data, averaged by year and by prefectural city level: spatial averaging was population weighted. The data comes from the European Center for Medium-Range Weather Forecasts (ECMWF),[14] and is a global atmospheric reanalysis product. Our data covers 286 prefecture cities from 2000 to 2013. Using this data set, we have repeated many of the tests described in the previous section. We have tested the relationships between annual GDP per capita and temperature in a panel data set, with and without lagged values of temperature.

---

[14] Details are available here: http://onlinelibrary.wiley.com/doi/10.1002/qj.828/abstract

Table 7: China Quadratic Total GDP and Primary Sector

	Log p.c	GDP			Log p.c.	primary	GDP	
	no lag	1-lag	3-lags	5-lags	no lag	1-lag	3-lags	5-lags
VARIABLES	lpGDP	lpGDP	lpGDP	lpGDP	lpGDP_1st	lpGDP_1st	lpGDP_1st	lpGDP_1st
temp	-0.0725**	-0.0779***	-0.0697***	-0.0562***	-0.121	-0.148**	-0.132***	-0.0682**
temp^2	0.000562*	0.000647***	0.000567***	0.000506***	0.00103	0.00127*	0.00110***	0.000535*
L.temp		-0.0234	-0.0295	-0.0245		-0.0412	-0.0852	-0.0553
L2.temp			-0.0342	-0.0280*			-0.0776	-0.0557
L3.temp			-0.0102	-0.0184			0.0291	-0.00293
L.temp^2		0.000161	0.000236	0.000239		0.000364	0.000735	0.00047
L2.temp^2			0.000297	0.000259*			0.000639	0.000550*
L3.temp^2			0.000106	0.000206			-0.000247	6.30E-05
L4.temp				-0.0102				0.0304
L5.temp				0.0128				0.121*
L4.temp^2				0.000135				-0.000214
L5.temp^2				-7.93E-05				-0.000933*
Observations	3,940	3,655	3,084	2,512	3,938	3,653	3,082	2,510
R-squared	0.989	0.991	0.993	0.994	0.899	0.91	0.939	0.97
$\sum temp$ effects		-0.101**	-0.144***	-0.124		-0.189	-0.266**	-0.0309
$\sum temp^2$ effects		0.000808**	0.00121***	0.00126***		0.00163	0.00222**	0.000472

The results are presented in tables 7, 8 and 9.

Table 8: China Quadratic Secondary & Tertiary Sectors

	Ln p.c	**secondary**	**GDP**		**Ln p.c.**	**tertiary**	**GDP**	
	no lag	1-lag	3-lags	5-lags	no lag	1-lag	3-lags	5-lags
Variables	lpGDP_2nd	lpGDP_2nd	lpGDP_2nd	lpGDP_2nd	lpGDP_3rd	lpGDP_3rd	lpGDP_3rd	lpGDP_3rd
trmp	-0.112***	-0.103***	-0.0989***	-0.0848***	-0.0234	-0.0226	-0.0145	-0.0136
temp^2	0.000977***	0.000929***	0.000892***	0.000847***	3.64E-05	7.56E-05	-1.57E-05	7.50E-05
L.temp		-0.0436	-0.0358	-0.0355		-0.00968	-0.0141	-0.012
L2.temp			-0.0364	-0.0297			-0.0258	-0.032
L3.temp			-0.0545*	-0.0254			0.00708	-0.0175
L.temp^2		0.000393	0.000344	0.000416		-4.20E-05	1.65E-05	3.56E-05
L2.temp^2			0.000379	0.000348			0.000138	0.000143
L3.temp^2			0.000637**	0.000397			-0.000203	2.36E-05
L4.temp				-0.0376				-0.00362
L5.temp				-0.00468				-0.017
L4.temp^2				0.000468				-7.61E-05
L5.temp^2				0.000184				3.22E-05
Observations	3,939	3,654	3,083	2,511	3,939	3,654	3,083	2,511
R-squared	0.978	0.979	0.983	0.986	0.98	0.983	0.987	0.99
$\sum temp$ effects		-0.147***	-0.226***	-0.218***		-0.0322	-0.0474	-0.0957
$\sum temp^2$ effects		0.00132***	0.00225***	0.00266***		0.0000336	-0.0000647	0.000234

The first two of these shows the results of fitting a quadratic form to the data, as was done with the cross-country panel in table 1. Tables 7 and 8 reports results for the whole of GDP and for the GDPs of the primary secondary and tertiary sectors. The results are strikingly different from those in table 1: in the cross-country data the linear term in the quadratic form had a positive coefficient and the quadratic term a negative coefficient: here matters are exactly the opposite. For the tertiary sector (services) the results are not significant, but for the others and for GDP as a whole the relationship between temperature and productivity is a regular U shape, as in the top panel of figure 2.3, and not as before an inverted U.

We repeated the non-parametric approach with the Chinese data and obtained the results in table 9. In this case the results are less sharp: we used three temperature bins labeled H (hot), M (medium) and C (cold). The coefficients on the hot bin are generally positive, but only sometimes significant, while those on the cold bin are negative and significant. The tertiary sector is an exception (as it was in the quadratic case) with the coefficients there suggesting an inverted U.

Table 9: China 3-Bin

	Log p.c GDP		
	no lag	1-lag	5-lags
VARIABLES	lpGDP	lpGDP	lpGDP
tempH	0.000976	0.00603	0.0139*
tempM	-0.00232	0.00248	0.00552
tempC	-0.0192**	-0.0164*	-0.00920*
R-squared	0.989	0.99	0.994
sum temp effects: C		-0.0257	-0.00804
sum temp effects: M		0.000874	0.0364
sum temp effects: H		0.00253	0.0815

4.1 Conclusions on China

The results for Chinese cities are surprising given what we and others have found for the cross-country data set. But they are consistent with the theory of section 1, which shows that in cases where income effects are important then we may find a productivity-temperature connection that is $\cup$-shaped. Equation (2.3) suggests that income effects will be important in situations where people work long hours, and while we do not have data on this as a part of this study, casual empiricism suggests that this fits the case of China.

5 Länder in Germany

Table 10: German Länder: Productivity vs Temperature

	Log per capita GDP			
	no lag	1-lag	3-lags	5-lags
VARIABLES	lpGDP	lpGDP	lpGDP	lpGDP
temp	-0.0868	-0.0275	-0.0493**	-0.0348*
temp^2	0.00092	0.00031*	0.000482**	0.000353***
L.temp		-0.0691	-0.0294	-0.029
L.temp^2		0.000731	0.000324	0.000415*
L2.temp			-0.0333**	-0.0355
L2.temp^2			0.000334**	0.000392*
L3.temp			-0.0390**	-0.0510***
L3.temp^2			0.000381***	0.000431***
L4.temp				-0.00142
L4.temp^2				8.97E-05
L5.temp				0.0221
L5.temp^2				-0.000137
Observations	322	308	280	252
R-squared	0.907	0.945	0.981	0.987
$\sum \mathit{temp}$ effects		-0.0966***	-0.151**	-0.13
$\sum \mathit{temp}^2$ effects		0.00104**	0.00152***	0.00154

The final application of our model is to data for Germany from 1192 to 2013: this data has annual GDP per capita by Land, and weather data from the ECMWF, spatially averaged population weighted daily maximum temperatures averaged over the year. We repeat the analysis of earlier cases, using year and Land-specific fixed effects, and find a clear U-shaped relationship between temperature and GDP per capita, as shown in Table 10. The connection here is weaker than in earlier cases, with no significant effect when no lags are included. However the overall effect is clearly consistent with the upper panel of figure 2.3.

6 Levels and rates of growth

We noted in section 2.2 that the relationship between temperature and productivity has implications for that between temperature and the rate of growth of productivity. As long as $K < K_1$, the relationship between temperature and the rate of growth of productivity will be qualitatively similar to that between temperature and the level of productivity. It certainly seems reasonable to assume $K < K_1$ for a developing country such as China, and for other developing countries in our cross-country data panel. With this in mind, we have also run regressions between the rate of growth of output per capita and temperature.

Table 11: China Productivity Growth

	growth rate of p.c. GDP			
	no lag	1-lag	3-lags	5-lags
VARIABLES	grGDP_prefecture_05rmb			
temp	-0.0192**	-0.0254**	-0.0271**	-0.0207
temp^2	0.000199***	0.000246***	0.000275***	0.000226*
L.temp		0.0227*	0.0242**	0.0207**
L.temp^2		-0.000156	-0.000157*	-0.00011
$\sum temp$ effects		-0.00271	0.00768	0.0637***
$\sum temp^2$ effects		0.0000898	0.000113	-0.000327***

Table 11 shows the results of a regression of the rate of growth of output per capita for Chinese prefecture cities against temperature: it is clear here that we have a U-shaped relationship, as in the cases of temperature and the level of productivity shown in Tables China Quadratic 1, China Quadratic 2 and China 3 bin. One difference between the two cases is that with rates of change, when we disaggregate to the primary, secondary and tertiary sectors we find no significant relationships, whereas with levels of productivity we find the same type of relationship in aggregate and by sector.

Table 12: Germany Growth of Productivity vs Temperature

	growth rate			
	no lag	1-lag	3-lags	5-lags
VARIABLES	g	g	g	g
temp	-0.000279	-0.0102	-0.00364	-0.0133
temp^2	-0.000101	-7.25E-05	-6.70E-05	1.03E-05
L.temp		0.0523	0.0246**	0.0325***
L.temp^2		-0.000431	-0.000180*	-0.000197***
L2.temp			-0.0071	0.000676
L2.temp^2			7.15E-05	2.65E-05
L3.temp			0.00508	-0.0266**
L3.temp^2			-8.81E-05	0.000132
L4.temp				0.0592***
L4.temp^2				-0.000466***
L5.trmp				0.00676
L5.temp^2				-0.000122
$\sum temp$ effects		0.0421	0.0189	0.0592***
$\sum temp^2$ effects		-0.000503	-0.000264	-0.000617***

When we look at the temperature-growth rate of productivity connection for Germany, the impacts of contemporaneous variables are not significant, as shown in Table 12, though some lagged variables do have an impact. In the cases of three and five lags there are significant coefficients which suggest an inverted U rather than a regular U as was shown in the regressions of the level of productivity on temperature. It is possible that in Germany, an advanced industrial country, the capital stock exceeds the key value K_1 (see figure 2.5), in which case as noted in section 2.2 there is no clear presumption that the relationship between temperature and growth of productivity is similar to that between temperature and the level of productivity.

We finally turn to the cross-country data set, and examine the relationship between temperature and the rate of change of productivity in that context. In the study of the effect of temperature on levels of productivity we found clear evidence of an inverted U relationship: productivity rises with temperature at low temperatures and then falls at high temperatures, as shown in figure 1.1. Given the wide range of countries in our data set, it is not clear what we should expect from the temperature-rate of change study. In fact we again find evidence for an inverted U relationship, as in the case of table 1. However the relationship between temperatures and rates of change is less robust statistically than that between temperatures and levels. Table 13 shows one such case: there is weak evidence for an inverted U relating temperature to productivity growth in the coefficients of the unlagged variables when lags are present (but not when there are no lags), but the coefficients on the sums of all temperature effects across all lags are highly significant and suggest a regular U.

Table 13: Cross-Country Growth of Productivity vs Temperature

	no lag	1-lag	5-lags	10-lags
VARIABLES	g	g	g	g
temp	-0.148	0.452	0.525	0.339
temp^2	-0.001	-0.024**	-0.023*	-0.011
L1temp		-1.291***	-0.811**	-1.075***
L1tem^2		0.050***	0.033**	0.044***
$\sum temp$ effects		-0.838**	-1.635***	-1.584**
$\sum temp^2$ effects		0.026**	0.049***	0.053**

Table 14 shows another aspect of the instability problem: it is the same regression as table 13, using the same data, but now we have added country-specific quadratic time trends. The results are very different: the unlagged temperature variables show clear support for an inverted U relationship between temperature and growth of productivity. This is not true for the sum of all lagged effects, though there is some support for an inverted U here too. This result is highly dependent on the inclusion of country-specific quadratic time trends and as table Cross Country Growth shows is not observable without them. We include these trends because other papers in the literature have done so (see Burke et al), but we cannot see a compelling reason for this. The trends are presumably proxies for omitted variables, in which case the best step is to try and identify and measure these variables. It is not immediately clear why there would be omitted variables in the temperature-growth rate regression which are quadratic in time and country-specific.

Table 14: Cross Country Growth with Country-Specific Quadratic Time Trend

	no lag	1-lags	5-lags	10-lags
temp	0.649*	0.809**	0.689*	1.221***
tem^2	-0.036***	-0.041***	-0.037***	-0.058***
L1temp		-0.894**	-0.761**	-0.345
L1wtemp^2		0.032**	0.026*	0.005
$\sum temp$ effects		-0.085	-0.904	4.663**
$\sum temp^2$ effects		-0.009	-0.005	-0.278***

7 Conclusions

We developed a theoretical framework for thinking about the relationship between temperature and productivity, basing this around the widely-recognized impact of temperature on task performance. We used this to show that a variety of different qualitative relationships are possible, depending on the particulars of the region which we are studying. The exact nature of the relationship depends on the importance of income effects (which we show will be larger in countries where people work longer hours) and on the way in which income effects and temperatures are distributed across regions. The temperature-productivity relationship may in theory be U-shaped, ∩-shaped, or monotonically increasing or decreasing, although only find evidence for the first two cases.

The connection between temperature and productivity, when embedded in a Solow growth model, implies a connection between temperature and the rate of growth of productivity. If the capital stock is lower than a certain value denoted K_1 (see figure 2.5) then the relationship between temperature and the rate of growth of productivity is qualitatively the same as that between temperature and the level of productivity: otherwise it may be different.

We test these ideas on different data sets. First we use a cross-country data set similar to that used by Dell et al. [2008], and find clear evidence for a ∩-shaped relationship between temperature and the level of productivity: an increase in temperature raises productivity in cold countries and lowers it in hot ones. However in hot countries with high levels of air-conditioning, an increase in temperature has no impact on productivity. A change in temperature also has no impact on productivity in "straddling" countries, countries which have hot and cold regions which are sufficiently different in temperature that we would expect the impacts of a temperature shock to have different signs. With this cross-country data set we find that there is also a $\cap-$shaped relationship between temperature and the rate of growth of productivity, though it is statistically less robust than that between temperature and the level of productivity.

Our second data set is a novel one that contains GDP and weather data for

286 Chinese prefecture cities for thirteen years. In this case we find a robust U-shaped relationship between temperature and productivity: this is predicted by the theoretical model if income effects are strong, which in turn is predicted if people work long hours, which appears to be the case in China. We also find a U-shaped relationship between temperature and the rate of growth of productivity, which we would also expect from the theory as the capital stock in China can reasonably be assumed to be low.

Finally we work with German data, using Land-level GDP and weather data over the period 1992 to 2015: in this case we find U-shaped connection between GDP per capita and temperature, and no clear relationship between temperature and the rate of growth of productivity. German capital stocks might well exceed the critical level K_1, in which case we do not expect a temperature-growth rate connection similar to the temperature-level connection.

References

Maximillian Auffhammer, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. Using weather data and climate model output in economic analyses of climate change. *Rev Environ Econ Policy*, Summer:181–198, 2013.

M. Dell, B.F. Jones, and B.A. Olken. Climate shocks and economic growth: Evidence from the last half century. Nber working paper 14132, 2008.

M. Dell, B.F. Jones, and B.A. Olken. Temperature and income: Reconciling new cross-sectional and panel estimates. Nber working paper 14680., 2009.

O Deschenes and M Greenstone. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. Nber working paper 13178, 2007.

Geoffrey Heal and Jisung Park. Feeling the heat: Temperature, physiology and the wealth of nations. Working Paper 19725, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138, December 2013.

Geoffrey Heal and Jisung Park. Temperature stress and the direct impact of climate change: A review of an emerging literature. *Review of Environmental Economics and Policy*, 10(2):347–362, August 2016.

Alan Heston, Robert Summers, and Bettina Aten. Penn world table version 7.1,. Technical report, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, November 2012.

J.K. Horowitz. The income-temperature relationship in a cross-section of countries and its implications for global warming. Technical report, 2001.

Solomon M. Hsiang. Temperatures and cyclones strongly associated with economic production in the caribbean and central america. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372, 2010.

Solomon M. Hsiang, Marshall Burke, and Edward Miguel. Quantifying the influence of climate on human conflict. Published online 1 August 2013 [DOI:10.1126/science.1235367], 2013.

K. Matsuura and C. Willmott. Terrestrial air temperature and precipitation: 1900-2006 gridded monthly time series, 2007. URL `http://climate.goeg.udel.edu`.

William D. Nordhaus. Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, 103(10):3510–3517, 2006.

Wolfram Schlenker and M.J. Roberts. Estimating the impact of climate change on crop yields: The importance of non-linear temperature effects. 2006.

O. Seppanen, W. Fisk, and Q. Lei. The effect of temperature on task performance in office environment. Lawrence berkeley laboratory working papers., 2006.

Robert Solow. A contribution to the theory of economic growth. *Quarter*, 70 (1):65–94, February 1956.