NBER WORKING PAPER SERIES

FEELING THE HEAT:
TEMPERATURE, PHYSIOLOGY & THE WEALTH OF NATIONS

Geoffrey Heal
Jisung Park

Working Paper 19725
http://www.nber.org/papers/w19725

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2013

We are very grateful to Raj Chetty, Lawrence Katz, Andrei Shleifer, Edward Glaeser, Sendhil Mullainathan, Joseph Aldy, Roland Fryer, David Cutler, Jeff Miron, Martin Weitzman, Michael Hanemann, Kerry Smith, Sol Hsiang, Kyle Meng, Wolfram Schlenker, Josh Graff-Zivin, Francis Teal, Abdulrahman El-Sayed, Lucas Brown, Jong Ho Hong, Emily Sands, Duncan Gilchrist, Richard Sweeney, Richard Carson, Angus Deaton and Harrison Hong for helpful comments and feedback. Thanks to Nan Zhong and Ratna Gill for excellent research assistance. Part of this research was funded by the National Science Foundation (NSF) Graduate Research Fellowship Program (GRFP). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Feeling the Heat: Temperature, Physiology & the Wealth of Nations
Geoffrey Heal and Jisung Park
NBER Working Paper No. 19725
December 2013, Revised June 2014
JEL No. J01,Q54,Q56

ABSTRACT

Does temperature affect economic performance? Has temperature always affected social welfare through its impact on physical and cognitive function? While many economic studies have explored the indirect links between climate and welfare (e.g. agriculture, conflict, sea-level rise), few address the possibility of direct impacts operating through physiology, despite a deep medical literature documenting the temperature sensitivity of human task performance. This paper attempts a synthesis of these literatures by (1) presenting a microeconomic model of labor supply under thermal stress, and (2) using country-level panel data on temperature and income (1950-2005) to illustrate the potential magnitude of temperature-driven productivity impacts. Using a fixed effects estimation strategy, we find significant temperature sensitivity of per capita income that varies, crucially, with a country's position relative to an optimal temperature zone. Hotter-than- average years are associated with lower output per capita for countries in hot climates and higher output per capita for countries in cold ones: approximately 3%-4% per degree C in both directions. Air-conditioning mediates the adverse impact of hotter years, consistent with the physiological explanation. This more direct causal link between climate and social welfare has important implications for both the economics of climate change and comparative development.

Geoffrey Heal
Graduate School of Business
616 Uris Hall
Columbia University
New York, NY 10027-6902
and NBER
gmh1@columbia.edu

Jisung Park
Columbia University
jisungpark@gmail.com

1. Introduction

What can temperature fluctuations tell us about the relative wealth of nations? How does the climate in which we live and work affect our economic well-being? Specifically, does temperature stress from heat or cold influence our ability to focus or to engage in productive activities? If a temperature-performance relationship does in fact exist, what could this tell us about past and present differences in income and productivity levels across countries and regions, or the potential future impacts of climate change? Exploring more deeply the potential causal relationship between temperature and economic welfare is the primary objective of this study.

We bring together two stylized facts from rather different fields. Each is conventional wisdom in its own field, yet we believe their juxtaposition can add value. They come from economics and physiology. The economic fact is that hotter countries tend to be poor. The physiological fact is that human performance over a range of tasks degrades sharply as temperature rises above or falls below an optimal threshold.

Each of these ideas is at the center of a substantial literature. Scholars have noted for centuries that hotter countries tend to be poor (Montesqieu [1750]; Huntington [1915]). Taking a cross-section of countries in 2000, for example, average per capita income decreases by roughly 8.5% per °C as one moves closer to the tropics (Horowitz [2001]). Sala-i Martin [1997] shows that growth rates decrease sharply with absolute latitude, which is a good proxy for temperature. More recently, Dell et al. [2008] find that hotter than average years are associated with lower than average GDP growth by roughly -1% per degree Celsius for a subset of poor countries, mostly in Sub-Saharan Africa.

That human performance on both physical and intellectual tasks degrades with temperature is also well-established. While economists have noted this only recently[1], similar observations have a much longer history in the medical literature, which suggests that heat can have measurable negative effects on physical and cognitive performance across various metrics. Thermal stress has well-documented effects on athletic performance (Wendt et al. [2007]), and can also adversely impact simple tasks such as manual tracking (e.g. guiding a steering wheel) and cognitive tasks such as sentence completion or basic arithmetic (Grether [1973], Wyon [1974]). Equally well-established is the ability of heating and air-conditioning to offset some of these adverse impacts[2], though the link between air-conditioning and macroeconomic growth has not yet been documented.

Our observation in this paper is that the phsyiological fact can help explain the economic one: that the temperature-performance gradient at the individual level can contribute to explaining the relationship between temperature and economic performance, and ultimately inform our understanding of the impact of future climate change.

Our most policy-relevant result is that annual climate shocks have non-trivial impacts on GDP per capita, but that the direction and magnitude of these impacts are grossly unequal. The economic impact of a warmer world may depend crucially on the initial temperature zone in which one is situated. Warmer-than-average years

[1]For example, Hsiang et al. [2012] show that student performance in standardized math tests falls as the temperature rises above the low 70s Fahrenheit.

[2]For instance, Deschenes and Greenstone [2007] find that local air-conditioning penetration reduces the mortality response to heat shocks in US states.

lead to negative per capita output shocks in hot countries (e.g. Bangladesh); in cold countries (e.g. Sweden), the reverse seems to be true. And while, given the spatial resolution of our data, we cannot rule out the role of other confounders such as agricultural yield or storm intensity, we suggest that this systematic heterogeneity in the treatment effect of temperature on GDP is consistent with the productivity relationships documented in the "sub-micro" literature and formalized in our model. The fact that countries with higher air conditioning per capita are less vulnerable to temperature shocks provides further evidence of a physiologically-mediated causal mechanism.

All of these results are preliminary. They are meant to illustrate the need for further research into the exact nature and scope of a possible pervasive connection between temperature, human physiology, and economic welfare, especially in countries without access to air conditioning and in activities necessarily exposed to external temperatures.

This paper does three things. First, it synthesizes emerging empirical research on the relationship between climate variables and macroeconomic variables such as income per capita (Horowitz [2001]; Dell et al. [2008]; Nordhaus [2006]), in conjunction with a longstanding medical literature on temperature and human task performance at what we call the "sub-micro" level. Second, it presents a model of labor supply decisions under temperature stress that is consistent with these stylized facts and which develops a sufficient statistic for future empirical welfare analysis. The key prediction of the model is that temperature deviations from a biological optimum (be that in the form of heat or cold) will reduce "effective labor supply," defined as the composite of raw labor hours, physiological task productivity, and labor effort, irrespective of the types of contract structures or labor market institutions present. For quasi-linear preferences the willingness to pay for mitigating these effects can be well-approximated by household expenditures on heating and cooling. Third, it provides a preliminary attempt at testing this model empirically, using country-level panel data relating per capita income to average annual temperature fluctuations and air conditioning imports per capita. The key findings are (1) a universally concave relationship between temperature and income levels that is dependent on the level of exposure to thermal stress, and (2) a mediating role played by AC penetration per capita.

The rest of the paper is organized as follows. Section 2 presents a synthesis of work on climate-economy interactions, through historical and prospective lenses. Section 3 presents some old and new facts about temperature and human activity at the level of the individual, which draws heavily from the medical and epidemiological literature. Section 4 presents the model and some empirical predictions that arise from it. Section 5 presents a simple empirical framework for identifying causal impacts of temperature on income at the country level, and presents the results from international panel data. Section 6 concludes.

2. The Evolving Economics of Geography, Temperature, and Climate Change

A casual scatterplot of log GDP and (population-weighted) average annual temperatures reveals a striking temperature-income gradient (Figure 2.1). While there is still considerable disagreement over how much of this cross-sectional relationship is driven by institutions (Acemoglu et al. [2001] among others) or other geographical correlates such as disease burden (Sachs et al. [2001]), more recent empirical evidence suggests that a large proportion of the causal effect is driven by climate variables (Dell et al, 2013).

Studies using national and sub-national cross-section data (Dell et al. [2009]; Horowitz [2001]), suggest that the income-temperature relationship exists not only across OECD and non-OECD countries, but also across provinces and counties within countries. If this is true, and, more importantly, if we can say something about why it is the case, the potential implications for both development theory and climate policy would be substantial.

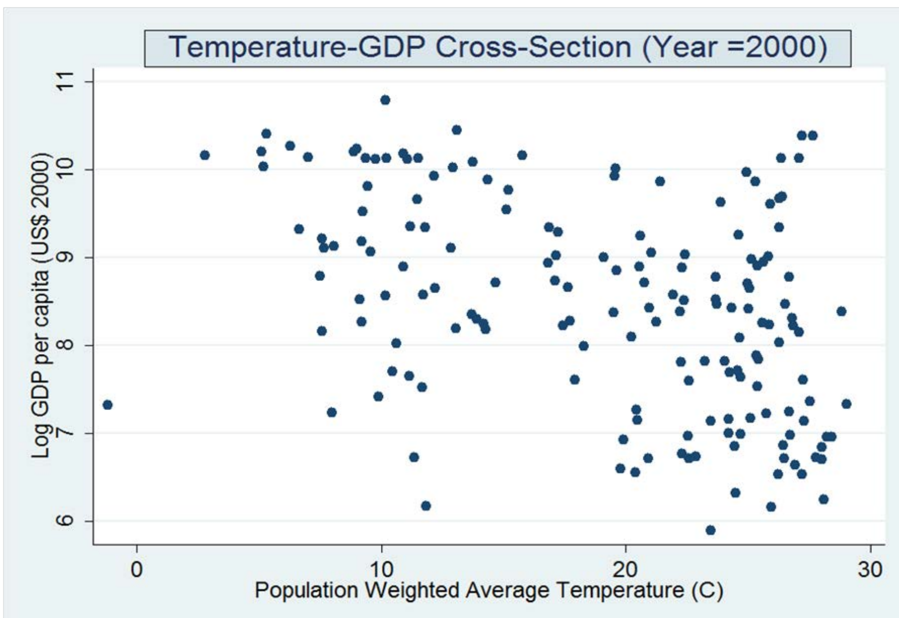

FIGURE 2.1. Countries by log income per capita and population-weighted average temperature

Dell et al. [2009] also show that hotter counties and municipalities are, on average, 1.2%-1.9% poorer per degree C average annual temperature (across 7,793 municipalities in 12 countries in the Americas), confirming that omitted country characteristics are not wholly driving the cross-sectional relationship (Dell et al. [2009]). Even among only OECD countries, +2°F is associated with –3.7% to -4.0% GDP (Horowitz [2001]). Simply extrapolating the existing cross-sectional relationship without accounting for adaptation or institutions might suggest that an average warming of +6-7°F in the future could lead to an average decrease of approximately -13%-14% of GDP worldwide, a much higher figure than most bottom-up climate damage estimates suggest (Horowitz [2001]).

Most of these studies emphasize the impact of heat (as opposed to temperature per se), motivated perhaps by the projected rise in global mean temperatures due to global warming. However, it is plausible that it is the extremity of climate, rather than the hotness per se, that adversely impacts human activity. There is much evidence to suggest that extreme weather events (e.g. hurricanes) have large and persistent negative effects on GDP. By this token, then, there is no a priori reason to expect a monotonic relationship between temperature and economic productivity; that is, the fundamental relationship between temperature and productivity may, in fact, be single-peaked, implying something akin to an optimal temperature zone for human activity.

While there is strong evidence for an optimal temperature zone at the micro level, causal evidence at the macro level has so far been thin. Nordhaus [2006] uses geospatially indexed economic and climate data at the grid cell level ("gross cell product") and finds a relationship between average annual temperature and output per grid cell that is robust and single-peaked. The fall-off in productivity toward hotter and colder extremes suggests an optimal temperature zone for human economic activity.

But what is the causal pathway underlying these relationships? Are these correlations due to the effect of temperature on institutions, or the incidence of disease and violent conflict? Or are other omitted variables driving the relationship? The human being, as with the rest of life on earth, is a biological organism evolved to function more effectively in some environments than others. And yet the question of whether and to what extent temperature affects economic wellbeing causally remains unresolved in the literature. While most of these studies have steered clear of emphasizing one causal pathway over another, we believe that insofar as most plausible pathways operate through human performance and human interaction, there may be a pervasive and perhaps universal role played by the effect of thermal stress on the human body.

Viewing the problem through this lens also leads to an important methdological shift. Whereas many studies have treated a $+1^{\circ}$C weather shock as the same "treatment" across all countries and regions, our approach suggests significant heterogeneity in treatment effect a priori. Whether a hotter year leads to adverse (or beneficial) outcomes depends crucially on whether this shock pushes one away from or toward the thermoregulatory optimum. [3]

3. Some Old and New Facts About Heat and Human performance

That extreme temperatures can hinder human activity at the individual level is almost tautologically true. Heat or cold can influence human behavior by making one less effective at any activity (e.g. working or exercising), and also by nudging one to choose certain activities over others (e.g. staying in the shade versus working out in the field). For example, the effect of heat waves on mortality – particularly among the elderly – is well documented in the epidemiological literature (Curriero et al. [2002]; Kilbourne [1997]; Kovats and Hajat [2008]; McMichael and et al [2008], etc). A growing number of studies have shown that, even in rich countries, extreme heat waves cause a large number of deaths. In 2003 for example, France suffered

[3] Of course, there are a number of documented links between climate and economic output that may be somewhat orthogonal to human physiology. Crop yields are adversely impacted by heat after a certain point (Schlenker and Roberts [2006]). Sea-level rise will no doubt damage many low-lying coastal assets (Yohe et al. [1996]). Changing rainfall patterns and storm intensity may affect the availability of water resources in different parts of the world, likely making dry areas drier, and wet areas wetter (Pachauri and Reisinger [2007]).

approximately 14,000 heat-related deaths mostly among the elderly, and Europe as a whole roughly 40,000.

The slope of the temperature-mortality response is heterogeneous, and in general not predicted by latitude, as shown by comparisons of cities in the US, Europe, or around the world (Curriero et al. [2002]; McMichael and et al [2008]). While some of this has to do with demographics (the relative densities of old and infirm), it has been suggested that a significant proportion of this variability is related to the prevalence of air conditioning (Kovats and Hajat [2008]), a key variable in the model presented in this paper. Deschenes and Greenstone [2007] show that hot days have historically led to very high mortality rates, and that the spread of air conditioning (AC) in the United States can account for up to 80% of the decline in heat-related mortality. They suggest that many developing countries which have much lower levels of residential AC penetration than the US may suffer increasingly severe mortality shocks from future climate change.[4]

But heat can also affect human welfare at less extreme temperatures, and in less extreme ways than outright mortality or morbidity. Task productivity has been shown to decline systematically with thermal stress (Wendt et al. [2007]). Even test scores, controlling for individual ability, appear to be sensitive to ambient temperatures, though the effect is, interestingly, significant for math but not for reading scores (Hsiang et al. [2012]).

There also seems to be evidence for behavioral responses by individuals in labor and leisure settings. Anticipating lower productivity and/or direct disutility from higher core body temperatures, individuals choose to exert less effort or devote less time to effort-involving tasks. A recent report by the Center for Disease Control and Prevention shows that residents of hotter regions in the US are generally less physically active (Centers for Disease Control and Prevention, 2011; Figure (6.1)). There is also evidence emerging from the behavioral psychology literature suggesting that individuals' anxiety levels, depression incidence, and propensity toward aggression are significantly correlated with temperature, sunlight, and cloud cover (Keller et al. [2005]). Insofar as GDP is a cumulative measure of productive activity over a year, even such subtle environmental factors could in principle create accumulated advantages or disadvantages over time. Indeed, there is a documented relationship between wages and climate amenities at the local level (Blomquist, ; Sinha and Cropper,), which may be related to productivity differences.

Using data from the American Time-Use Survey, Graff Zivin and Neidell [2010] find evidence for changes in time-use decisions resulting from temperature shocks. In industries with high exposure to climate, workers report lower time spent at work on hot and cold days, as well as in time spent on outdoor leisure activities. While Graff Zivin and Neidell do not show this, intuitively one might think that extreme temperature and weather events lead to a reduced average flow intensity of economic activity if measured at a high enough level of aggregation. Similarly, Adhvaryu et al (2013) show that manufactering worker efficiency at the plant level declines substantially on hotter days, an effect that is driven primarily by on-the-job task productivity declines as opposed to increased absenteeism. [5]

[4]Lee Kwan Yu once declared that air conditioning was the single most important inventions in history, and that, without it, Singapore could never have grown to the thriving tropical megapolis that it is today.

[5]This is a key intuition that justifies our use of country-level data in the empirical analysis. For example, if a hotter-than-average year leads to five more days of above-100 degree temperatures,

Meta-analyses of this vast and growing literature confirm the presence of a non-linear relationship between thermal stress and productivity (Seppanen et al. [2006]; Hancock et al. [2007]).[6] The stylized empirical trend seems to be a single-peaked relationship between temperature and productivity, where negative productivity impacts increase non-linearly the further one deviates from the biological comfort zone (approximately 18°C to 22°C), a trend consistent with existing models of human physiology (Figure 3.1).[7]

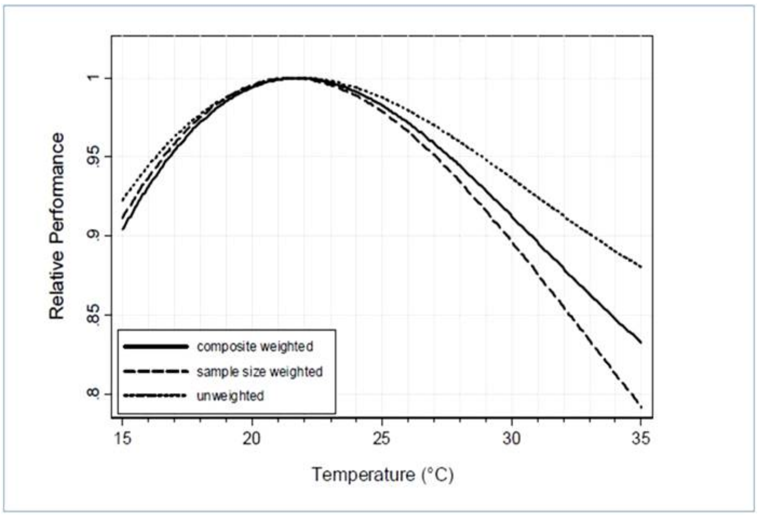

FIGURE 3.1. Task performance vs temperature. Maximum performance is normalized to 1 at 22 C. Source: Seppanen et al. [2006]

In summary, a large number of studies from various disciplines show physical and cognitive performance to deteriorate with temperature deviations beyond a biologically optimal zone. In other words, there is a single-peaked and non-linear relationship between temperature and task effectiveness at the micro or sub-micro level.[8] The biological mechanism through which this effect works is that of thermoregulation. We believe this biological mechanism is fundamentally related to many of the documented climate-economy links in the literature (Table 2).

---

which leads to the cancellation of several workdays or meetings that were meant to be held during those days, one would expect a noticeable impact on annual output, unless these shocks were made up for by cannibalizing leisure time. From a social welfare perspective, however, even if individuals engage in forced "make-up" work by taking away from leisure time, in the absence of parallel preference shifts, this is a clear welfare loss, even if nominal output may remain the same.

[6]Seppanen et al. [2006] and Hancock et al. [2007] conduct meta-analyses of 24 and 49 lab and field studies respectively and find robust single-peaked relationships between ambient temperature and objective metrics of worker productivity in indoor, office environments. Both groups of authors are cautious to select only those studies that use "objective" measures of productivity, as opposed to subjective measures such as self-reported productivity or peer-evaluations. They also weight the studies by sample size, which vary from 9 to 500 individuals per study. The tasks measured include

Potential Impacts of Thermal Stress on Human Welfare that Operate through Physiology and Thermoregulation			
Utility		**Health and Human Capital**	
Direct disutility	Connolly [2013]	Mortality	Kovats and Hajat [2008]
Travel amenity	Unexamined	Morbidity	Deschenes and Greenstone [2007]
		Cognitive Function	Hsiang et al. [2012]
Effective Labor Supply		**Interactive/Political**	
Labor Hours	Graff Zivin and Neidell [2010]	Innovation	Dell et al. [2008]
Labor Effort	Unexamined	Crime and Violence	Hsiang et al. [2013]
Labor Productivity	Seppanen et al. [2006]	Political Instability	Dell et al. [2008]

Table 2. Categorization of Potential Causal Impacts of Thermal Stress on Human Welfare

4. A Model of consumer behavior under thermal stress

We next develop a simple formal model that reflects the issues reviewed above. It combines elements of the standard labor-leisure choice model from labor economics with the thermoregulatory factors that emerge as important influences on labor productivity as temperatures vary. The result is an optimizing model of the choice of labor hours and effort, leading to a physiological-economic model of labor supply.

All human beings regulate core body temperature to keep it as close as possible to a biological optimum (98.6°F, 37°C) (Kovats and Hajat [2008]). Scientific evidence suggests that we do this both subconsciously – through sweating or involuntary physical activity modulation (for example, shivering) – and consciously – by putting on or taking off clothing, or turning on the air-conditioning or heating if it is available. Core body temperature is affected by a host of factors which can be grouped into the following three categories: 1) physiological factors, including level of physical activity, and involuntary acclimatizing activities such as sweating, shivering, or long-term physical acclimatization (biologists refer to this as the metabolic rate), 2) ambient temperature and humidity, and 3) the built environment (e.g. the availability of heating and air conditioning). As the core body temperature moves further away from the biological optimum, we devote more and more energy to trying to bring it back: more energy to shivering if it is too low and to sweating

office type work, text processing, length of customer service time, simple numerical calculations, and total handling time per customer for call-center workers.

[7]The authors suggest that these results likely underestimate the true magnitude of the effect on productivity, due to the short term nature of many of the lab experiments reviewed (Seppanen et al, 2005).

[8]We call these "sub-micro" studies in that the effect often occurs without conscious decisions or awareness on the part of the agents themselves. Micro-economics typically applies to models of individual utility maximization and the choices that individuals make, not subconscious processes.

if too high (Parsons [2003], Kilbourne [1997]). And when the temperature is too high, the body automatically circulates more blood near the skin in order to take advantage of cooling opportunities, and limiting the supply to key organs. These cooling opportunities are more limited if the external environment is hot or humid. It takes only a small deviation from the optimal core body temperature for a person to be very sick – consider a temperature of 101°F, only three degrees above the optimum, yet high enough to make it difficult to function. A temperature of 104°F maintained for several days can prove fatal.

One of the principal mechanisms through which temperature affects performance appears to be the ability of the brain to dispose of waste heat: on average the brain generates 20% of all the heat generated by the human body, and its performance is temperature-sensitive, so that it needs to dispose of waste heat (Schiff and Somjen [1985], Yablonskiy et al. [2008]). This becomes harder as the ambient temperature rises.

In economic terms, the consequences of thermal stress (a shock to body temperature pushing it away from the optimum) are threefold: 1) feeling excessively hot or cold, which we model as a direct loss of utility or welfare, 2) a drop in task performance which lead to a reduction in earning power, and 3) behavioral adjustments by the agent to reoptimize subject to the new temperature. Body temperature is determined by the external temperature, by the level of physical activity, and by expenditures on cooling, such as air conditioning.

Utility is assumed to depend on income, leisure, effort supplied and core body temperature. So we have that $U = U(Y, L, A, T)$ where Y is income, L is leisure, A is effort supplied to the work (related to the physiological concept of metabolic rate) and T is core body temperature. U is increasing in Y, L and decreasing in A. Utility is a concave function of core body temperature, increasing at low values of T and decreasing at high values. Hence the derivative of U with respect to T, U_T, changes sign as T increases, and the second derivative $U_{T,T}$ is negative.

These variables are interrelated:

$$Y = (1 - L)\, AP(T)\,,\ T = T(E, A)$$

where E is the environmental (external) temperature and $P(T)$ is labor performance.[9], a function of core body temperature. We normalize the wage rate to unity. Hsiang et al. [2012] note for example that math test scores decline with temperature: this is an aspect of performance, even if it might not be classified as a change in productivity. Performance increases with temperature at low temperatures and decreases at high temperatures, so that P_T, the derivative of P, changes sign from positive to negative and $P_{TT} < 0$. Income is hours worked multiplied by both effort and performance. More effort means working harder, and greater performance means that a given level of effort leads to more output. The core body temperature T is influenced by external temperature E and effort or metabolic rate A.

The total supply of labor is taken to be 1. Hence

$$U = U((1 - L)\, AP(T)\,, L, A, T(E, A))$$

gives the full specification of utility. In this relationship, E is a parameter given by the external environment, T and P are functional forms given by physiological

[9]By using the word performance we intend to include a broader range of effects than would be indicated by productivity.

considerations, and A and L are choice variables selected to optimize U subject to the relationships between the variables. In particular for given functions U, P and T the choices of A and L depend on the external temperature E: denote the maximizing values by $A^*(E)$ and $L^*(E)$. We can then write the indirect utility function

$$V\left(L^*(E), A^*(E)\right) = max_{A,L} U\left(\left(\left(1-L\right) AP\left(T\right), L, A, T\left(E, A\right)\right)\right)$$

More generally we will write

$$W\left(L, A : E\right) = U\left(\left(1-L\right) AP, L, A, T\left(E, A\right)\right)$$

as a simplified representation of utility, showing its dependence on the choice variables L, A and the external parameter E.

From this general framework, we will specialize to a particular functional form and assume that utility is quasi-linear in income:

$$U\left(Y, L, A, T\right) = Y + f\left(L, A, T\right) \tag{4.1}$$

as this makes possible a more precise understanding of the mechanisms at work. In this specification we are assuming that the interactions between leisure, effort and temperature are independent of income. We will also adopt a more specific functional form for the relationship between body temperature T, external temperature E and effort or metabolic rate A. We will assume

$$T\left(E, A\right) = \alpha + \beta E + g\left(A\right) \tag{4.2}$$

where α, β are constants and $g\left(.\right)$ is a concave increasing function. This is consistent with the physiological literature, which again suggests that core body temperature is non-decreasing with effort.

Optimizing behavior is characterized by the two obvious first order conditions:

$$\frac{\partial W}{\partial A} = 0, \ \frac{\partial W}{\partial L} = 0 \tag{4.3}$$

and we can treat these as implicit functions relating L, A and E and differentiate these by the implicit function theorem to obtain comparative static results on how the optimal choices of A and L respond to an increase in temperature E. The results are

$$\frac{dA}{dE} = -\left\{\frac{W_{A,E}}{W_{A,A}}\right\}, \ \frac{dL}{dE} = -\left\{\frac{W_{L,E}}{W_{L,L}}\right\} \tag{4.4}$$

where $W_{A,E} = \frac{\partial^2 W}{\partial A \partial E}$ etc.

We need to sign the expressions in (4.4). Consider the denominators $W_{A,A}$ and $W_{L,L}$: we assume the problem to be such that the optimal choices of both A and L are interior maxima. (Below we verify that this condition is in fact satisfied.) In this case the second derivative of W with respect to each is at least locally negative, implying that at an optimum

$$W_{A,A} < 0, \ W_{L,L} < 0$$

Hence the signs of the derivatives in (4.4) are those of the numerators in the parentheses, which we investigate next. It is easy to verify that the sign of $\partial A / \partial E$, the derivative of effort with respect to external temperature, is equal to that of

$$\left(1-L\right) P_T \beta + \left(1-L\right) AP_{TT} \beta g_A + f_{A,T} \beta + f_{T,T} \beta g_A \tag{4.5}$$

In this expression, we know that $(1-L), \beta, g_A > 0$. We also know that $P_{TT}, f_{T,T} < 0$. P_T changes sign from positive at low body temperatures to negative at high temperatures. We have not yet assigned a sign to $f_{A,T}$.

The issue in this case is: does the marginal disutility of effort rise or fall with body temperature? We assume $f_{A,T} < 0$, so that the marginal disutility of effort becomes more negative at higher temperatures.

The combined effect of these conditions is that the sign of (4.5), which is the sign of the derivative of effort with respect to external temperature, is negative at high temperatures (those at which productivity falls with temperature) and could be positive at low temperatures if P' is sufficiently large.

Next we check the sign of $\partial L/\partial E$, the effect of the external temperature on the amount of leisure chosen. This is equal to the sign of

$$-AP_T\beta + f_{L,T}\beta \tag{4.6}$$

Here $A, \beta > 0$, and as we have already noted P_T changes sign from positive to negative. $f_{L,T}$ shows the impact of body temperature on the marginal utility of leisure. Under the assumption that working in extreme conditions, be they heat or cold, is difficult and unpleasant, it seems reasonable that the marginal utility of leisure will be greater at high and low temperatures and lower at intermediate temperatures: this implies that f_L as a function of T is U-shaped and $f_{L,T}$ is negative and then positive. Hence $\frac{dL}{dE}$ is first negative and then positive: leisure (work) is decreasing (increasing) then increasing (decreasing) in external temperature. Hence we have established

Proposition 1. *With quasi-linear preferences and under the specified assumptions about the signs of $f_{A,T}$ and $f_{L,T}$, an increase in environmental temperature will lower the amount of effort A supplied at high temperatures, may raise the effort supplied at low temperatures, and will raise the hours worked at low temperatures but lower the hours worked at high temperatures.*

This clearly implies that productivity in terms of output per person will fall with an increase in temperature at high temperatures: people work less hard for fewer hours. Output per person may rise as temperature rises at low temperatures, as hours worked rise and effort may also rise, but only if the direct impact on performance is large enough.

4.1. **Spending on Thermoregulation.** Next we develop a simplified model that allows us to analyze spending on thermoregulation, and establish a relationship between this spending and the welfare losses from temperature changes. The model specifies only the bare essentials:

$$U = U(Y - S, T - rS)$$

where Y is income, T temperature before cooling as before, and S is the amount the agent spends on cooling. Each dollar spent on cooling reduces temperature by r degrees, and of course net income is reduced by S. Clearly the first order condition for the optimal choice of cooling C is

$$r = -\frac{U_Y}{U_T}$$

which just tells us that the marginal rate of substitution between income and temperature should equal the cost of reducing the temperature.

Now the loss of welfare from a temperature shock ΔT is

$$\Delta U = U_T \Delta T$$

Next we find the change in spending on cooling as a result of this temperature shock. For this we need the derivative of S with respect to T when the first order condition is satisfied. This is

$$\frac{\partial S}{\partial T} = -\left\{ \frac{-U_{YT} - rU_{TT}}{U_{YY} + 2U_{YT} + r^2 U_{TT}} \right\}$$

which in the quasi-linear case reduces to $1/r$. The welfare loss is $\Delta U = \Delta T U_T = \Delta T U_Y / r = \Delta T / r$ as $U_Y = 1$ in the quasi-linear case. But the increment to spending is $\frac{\partial S}{\partial T} \Delta T = \frac{\Delta T}{r}$. Hence in this case **the increment to spending as a result of the temperature shock exactly equals the associated welfare loss**[10].

Proposition 2. *With quasi-linear preferences the welfare cost of a temperature shock is exactly equal to the extra spending that results from the shock.*

4.2. **Implications for empirical work.** Several points from this theoretical analysis have implications for our empirical work.

(1) For quasi-linear preferences, the increase in spending on cooling (or heating) as a result of an increase (decrease) in temperature is exactly equal to the welfare loss from this increase. For more general preferences, the increase in spending is a lower bound on the welfare loss.
(2) Holding external temperature constant, changes in effort (or other factors that influence metabolic rates such as whether or not someone is working) will affect the expenditure on cooling or heating.
(3) With a group of people who have identical (or, strictly speaking, very similar) quasi-linear preferences, then in aggregate they behave as one person with quasi-linear preferences.[11] This means that with the model developed here, we can move freely between different levels of aggregation – from individuals to households to larger groups and even nations.
(4) At high temperatures, an increase in temperature will lead to a drop in performance, via decreases in both effort and hours worked – what we call "effective labor supply."
(5) At low temperatures, an increase in temperature will lead to an increase in hours worked and possibly in effort, and may lead to an increase in output per person.
(6) For a given adverse shock to the external environment (ambient temperature), an individual with more installed thermoregulatory capital (higher expenditures on cooling and heating) will suffer a smaller shock to productivity.

According to points (4) and (5) above, we expect that in a study of the impacts of temperature changes, we will see different responses in hot and cold environments, with output responding negatively to a temperature increase in hot environments

[10]Note that a change in core body temperature $\triangle T$ can be caused by a change in the external temperature or by a change in the level of physical or mental activity, which will change the the metabolic rate. If we compare the responses of people with different metabolic rates, those with higher rates will have a greater change in core body temperature in response to a given temperature shock.

[11]See Mas-Collel et al. [1995]

and possibly positively in cold ones. Point (6) suggest that, at the macro level, countries with varying levels of thermoregulatory capital may react differently to a given temperature shock. We do in fact find evidence of all three effects in the analyses that follow.

5. Empirical Results

In the following section, we take our model to cross-country data on climatic shocks and income per capita in an attempt to revisit the age old question: what is the role of climate in explaining the relative wealth of nations? We find suggestive evidence of a physiological effect of climate on economic activity at the macro level, one which may profoundly influence the way policymakers think about the welfare consequences of future climate change. Inasmuch as climate change will push already heat-stressed countries - which tend to be much poorer on average - toward more heat-stressed extremes, it may exacerbate existing income inequalities at the global level.

Whereas previous studies have focused on the role of heat (or low latitude) in predicting GDP, we predict that deviations from the thermoregulatory optimum, as opposed to hotter temperatures per se, are what dictate the magnitude of climate-GDP impacts. In some sense, we use the medical literature on thermal stress and human performance to inform our prior about treatment-effect heterogeneity; a warmer-than-average year (the "treatment) would have a very different impact on productivity and GDP for warmer, tropical countries than it would on cooler, temperate ones. It turns out that allowing for this particular form of effect heterogeneity makes a big difference in interpreting even the most well-studied macroeconomic datasets.

Our analysis suggests that the relationship between temperature and income is nearly universal (i.e. not necessarily limited to poor countries) and single-peaked, in line with what the physiological literature and our model imply. The causal effect of thermal stress is highly negative in already hot environments such as Thailand and India (as much as -3.9% annual output per capita per degree Celsius) and highly positive (up to +4.1%) in cool environments such as Canada and Sweden, with an indeterminate effect in temperate zones. In the time period surveyed (1950-2006) a one degree C hotter-than-average year occurs roughly once every 17 years. While we hesitate to extrapolate directly to future climate change scenarios, it is worth noting that such a two-sided "dose-response" to global warming could have serious political, economic, and philosophical consequences. [12]

Figure 6.2

As we note, there are many potential confounders that limit one's ability to interpret these estimates literally. While the single-peaked relationship between temperature and output per capita is certainly consistent with a model of thermoregulatory stress, it may also be driven by other, correlated causal factors – for example changes in agricultural yield. In principle it may also arise from spurious correlation resulting from secular time trends in temperature and total factor productivity (TFP). We attempt to control for these confounders by using air conditioning data, as well as allowing for flexible, country-specific time trends, discussed

[12]The point estimates reported here refer to the contemporaneous impact of temperature on log per capita income allowing for up to 10 lags in temperature, controlling for precipitation, country and year fixed effects, in addition to capital stock variables. See Table 4

in more detail below. The core result – a single-peaked relationship between temperature and output – is robust to a wide range of specifications.

5.1. **Empirical Framework.**

Before setting out our estimation strategy, we note that there are two important dimensions to consider when exploring the effect of temperature fluctuations on macroeconomic aggregates.

First, the initial climate in which an economy is situated matters. Our model suggests that the impact of a hotter-than-average year will not be the same across different "original climates." A one-degree C hotter-than-average year may lead to diminished overall labor performance in an already warm environment (Namibia), but it may actually lead to increased overall labor yield in a cold country (Norway). Second, in moving from a microeconomic model of thermoregulation to an analysis of macroeconomic variables, we must take into account industry composition: that is, the relative compositional sensitivity of the economic activity in a country or region to the effects of thermal stress on productivity. Occupations more intensive in outdoor labor are likely to be more sensitive to thermal stress, and countries with a higher share of economic activity concentrated in these industries to be more sensitive to temperature shocks[13]. Crucially for this analysis, the sensitivity of GDP to temperature stress may also be related to the degree of thermoregulatory capital available: that is, electrification, air conditioning, and access to heating systems and heat fuel. Given the asymmetric impact of physical activity in cold and hot environments mentioned above, as well as the relatively advanced technological requirements of AC (which results in more cross-country variation) we focus on thermoregulatory capital at the top end: capital that "defends" against heat stress in particular. Using a novel data set on air conditioning penetration by country that we construct from international trade data, we test whether the sensitivity of GDP to temperature is mediated by air conditioning, and find that it appears to be highly dependent on the amount of AC expenditure per capita.

Following DJO, we use historical fluctuations in temperature within countries to identify its effect on aggregate economic outcomes. Unlike DJO, we focus on the effect of temperature on the *level* of income per capita (as opposed to the growth rate), noting that the impact of thermal stress on labor productivity is mostly contemporaneous.[14]

Suppose each country's annual per capita GDP, Y_{it}, is produced using a combination of capital and effective labor input:

$$Y_{it} = Y(\theta_i, N_{it}, K_{it})$$

where once again the inputs are expressed in per capita terms. K_{it} denotes a holistic measure of capital (human and physical), N_{it} is a measure of effective labor supply, and θ_i is some country-specific measure of factor productivity that might be thought of as the institutional environment in country i.[15] Per capita output is increasing in effective labor supply.

[13]Cachon et al (2012) find that, even in automobile manufacturing plants in the United States, temperature shocks have a significant adverse impact on productivity, suggesting that even indoor manufacturing occupations may not be immune to the effects of thermal stress.

[14]As some recent studies (for example, Hsiang [2010]) have shown, there may be lagged impacts insofar as temperature effects investment that would have paid out in future years. It is unclear how large these effects might be.

[15]We abstract away from population growth for simplicity.

Define effective labor input, N_{it}, as a composite of labor hours $(1-L)$, labor effort (A), and labor performance (P), a function of the ambient temperature, T:

$$N_{it} = N_{it}((1-L), A_{it}, P(T_{it}), T_{it})$$

Insofar as the level of effective labor supply depends on the ambient temperature experienced by workers in the country (T_{it}), we would expect per capita output to be a function of experienced temperature:[16]

$$Y_{it} = Y_{it}(N_{it}(T_{it}), A_i, K_{it}, T_{it}).$$

Abstracting from capital inputs, we focus on the role of effective labor inputs:

$$Y_{it}(N_{it}, A_{it}, T_{it})$$

According to the model presented in section 4, and the mapping from changes in T_{it} to changes in N_{it} described therein, we expect the relationship between per capita output and temperature to be single-peaked: with Y_{it} decreasing in both directions away from the optimal zone. We attempt to estimate this relationship by utilizing within-country variation in historical annual temperature realizations, using panel data analagous to that used by DJO (Dell et al. [2008]).

5.2. **Data.**

5.2.1. *Climate Data.*

Annual average temperature and precipitation data at the country level are taken from DJO (Dell et al. [2009]). Temperature is measured in degrees Celsius, precipitation in mm per year. Their data is derived from Terrestrial Air Temperature and Precipitation: 1900-2006 Gridded Monthly Time Series, Version 1.01 (Matsuura and Willmott [2007]), and is weighted by population. 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). [17]

5.2.2. *International Economic Data.*

We use income data from the UN National Accounts. Real GDP per capita is measured in terms of USD$ (2000) using Laspreyes constant prices. Like DJO (Dell et al. [2008]), 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 1950-2006, and a total of 6,101 observations.

5.3. **Statistical Model.**

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

$$y_{it} = f(T_{it}) + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{5.1}$$

[16]This is one reason why population-weighted average temperature is a more relevant metrix than a raw geographic average.

[17]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.

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 (e.g. quadratic) relationship between income and temperature, as the medical and experimental literature suggests and summarized in the model of section 4.:

$$y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{5.2}$$

In this case, our main hypothesis is that the coefficients on T and T^2 are positive and negative respectively. That is, the relationship between temperature and income is globally single-peaked around some optimal zone. More specifically, we hypothesize that the GDP-residual, controlling for institutions, capital stock, and education, is dependent on temperature.

In an ideal experiment, we would expose otherwise identical economies to a series of random temperature shocks, and would do so for the whole range of base climates. This is for obvious reasons impossible at the macro level. Our econometric challenge is to come as close to such an experiment as possible with the data that we have.

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[18].

We confirm that there exists a strongly negative cross-sectional relationship between temperature and income in countries where population-weighted average temperatures are above 20°C. 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[19].

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

18

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. Dell et al (2009) find a strong, statistically significant negative relationships between temperature and income in a cross-section, of slighly smaller magnitude. In Nordhaus' case, the finding is of a strongly single-peaked relationship.

19 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 idiosyncracies. 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.

though a much larger number of countries have regions with very cold climates (e.g. Alaska and the Upper Midwest in the US; the Tibetan Plateau in China). 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: be that positive or negative.[20]

The panel nature of the dataset allows for one 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 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.[21] 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, though they may be correlated over time in the short run.

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

$$y_{it} = f(T_{it}) + \beta_3 K_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{5.3}$$

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 the section below, as well as in the Appendix.

5.4. Main Results.

We begin by estimating a single-peaked (quadratic) relationship between temperature and income per capita. Table 3 presents the coefficients from estimating equation (5.2) above. We allow for the possibility that temperature may affect GDP 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.[22]

---

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

[21] Both variables are taken from the Penn World Tables, version 8.0 (Heston et al, 2013).

[22] 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.

Our coefficient of interest, therefore, is the contemporaneous impact of temperature in year t on income in year t. Columns (9) through (12) suggest 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.[23]

Table 3

Next, we consider a more flexible functional relationship between temperature and GDP (5.1), 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" (15-20°C), "cold" (10-15°C), and "very cold" (10°C and below)[24]. The results suggest a single-peaked relationship, with the implied peak again occuring somewhere between 15° and 20° Celsius (Figure 6.2). 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).

Table 4

Table 5

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, notably, 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.

5.4.1. *Robustness Checks for Omitted Variables, Adaptation, and Spurious Correlation.*

We have established a single-peaked relationship between temperature and output per capita, and posited that this arises in part from the physiological factors discussed in earlier sections. There are of course alternative mechanisms which could lead to this relationship, as well as possible time-series properties of the data

[23]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.

[24]The number of observations in each bin are 1384, 1151, 470, 544, and 442 respectively.

that might bias our results. These include standard omitted variable bias, the distinction between weather and climate (and potential for adaptation over time), and spurious correlation between temperature trends and TFP trends. In this section, we tackle each issue in turn.

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 precipitously (Schlenker and Roberts [2006]). This suggests that, looking only at 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, in principle, explain a portion of the observed temperature-dose-GDP-response relationship. 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 (Table 6).

Table 6

There is also evidence to believe that negative public health impacts from extreme temperatures work through a diverse range of mechanisms, including influenza outbreaks, the spread of tropical disease vectors, and the effects of heat stress on outright mortaility. While the focus to date 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.

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 attempt to control for this difference by allowing for the treatment to be defined over 3-year intervals (3-year block averages, 3-year moving averages) as well, but find relatively minor differences in effects, suggesting that, at least over the short-to-medium run, economies seem to adapt only marginally[25].

[25] Using longer time intervals might be preferred in testing the weather-climate distinction, but doing to reduces our sample size significantly.

Table 7

Note that the reported coefficients for the 3-year block averages are much larger than the original specification. This is due to the fact that the "shock" in question is now a much rarer and more severe event. Whereas the original regressions use a 1° C hotter than average year as the treatment unit, this specification uses a period of 3 years with 1°C hotter than average temperatures. Seen in this light, the magnitudes seem consistent with a view that partial adaptation is possible. The 3-year impact is greater than the 1-year impact, but less than three times larger, suggesting some adaptation to prolonged extreme temperatures.

Table 8

Looking at the 3-year moving average estimates, it is easy to see that the shape of the relationship remains the same, and that, controlling for lagged impacts, one sees a similar response to annual fluctuations as one does to longer-term shifts in climate.

5.4.2. *The Role of Air Conditioning.*

Additional evidence strengthens the case for physiological impacts as a key causal mechanism. We test for the impact of thermoregulatory capital on the temperature-output gradient, by utilizing data on country-specific air conditioning penetration. Insofar as thermoregulatory capital 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 thermoregulatory capital.

Using the quadratic model, we attempt to examine whether access to air conditioning attenuates the effect of thermal stress at high temperatures. Because country-specific data on air conditioning penetration per capita is not readily available, we construct a proxy for air conditioning penetration per capita by imputing the value of air conditioning 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. Because we expect AC to protect against heat shocks but not extreme cold (though AC penetration is likely correlated with heating equipment as well, seeing as though the adoption of electrification and AC comes later than basic space heating in many development contexts), we restrict our sample to the subset of temperate, hot, and very hot countries.

Using this data for this subset of countries, we stratify the sample based on whether a country resided in the top third or bottom third of air conditioning penetration per capita averaged over the sample period 1960-2010. Table 11 presents the results for the two subsets of countries, allowing for lagged impacts once again. Consistent with the notion that higher levels of thermoregulatory capital dampen the impact of thermal stress on productivity, the subset of countries with air conditioning penetration in the top-third of the sample feature a less concave relationship between temperature and income per capita. The temperature-income gradient implied by the coefficients on temperature and temperature squared in columns (2), (4), (6), and (8) – the subset countries with top-third air conditioning – is shallower than that implied by the coefficients in columns (1), (3), (5), and (7) – which represent the subset of countries in the bottom third. Countries with lower levels of AC suffer large negative impacts from temperature shocks - between -3.1 and -9.2 log points of income per capita per year[26].

Table 11

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 thermoregulatory capital tend to be richer on average, there are also relatively hot and poor countries with high air conditioning penetration (for instance, Libya; see Table 6.3). 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[27].

6. Conclusion

Four main implications emerge from our analysis.

First, it seems that labor productivity may be a key link between climate shocks and economic outcomes at the macro level. While many have documented this causal mechanism at the micro and "sub-micro" level, this study is the first to demonstrate a connection at the macro level. We take the globally single-peaked relationship between temperature and output per capita, combined with the mediating impact of air-conditioning per capita among temperate and hot countries, to be highly suggestive of a causal link between climate shocks and economic output that operates through the impact of temperature on human physiology. The magnitude of these impacts may be as large as 3-4% per degree Celsius – in both positive and negative directions. To our knowledge, this paper also presents the first micro-founded model of economic behavior under temperature stress.

Second, in the context of economic development, this study sheds new insight into the old question of what determines the relative wealth of nations; it seems that temperature per se has always mattered for economic productivity, despite its well-documented correlation with institutional variables. The fact that temperature shocks away from an optimal zone can have significantly negative impacts for countries with low levels of air-conditioning suggests that Sub-Saharan Africa remains poor in part because of its inability to provide adequate thermoregulatory infrastructure to an already heat-stressed workforce. While it is likely that, in the real world, better institutions, better infrastructure, better education – and many other documented "deep determinants" of growth – go hand in hand, our study suggests that the growth handicap associated with hotter climates per se may be worth noting in earnest.

Third, these results have important implications for social cost of carbon estimates, which are crucial in determining whether and how much countries should

[26]The sum of all lagged temperature impacts in the 5- and 10-lag cases, though not reported in the table above, are significantly negative: -0.21*** and -0.32*** respectively.

[27]Note that one cannot estimate the impact of AC directly by including AC per capita on the RHS of the regression. This is because AC per capita is endogenous in a fixed effects model.

act to mitigate climate change. Our estimates suggest that current integrated assessment models may be missing an economically significant causal link in the form of labor productivity and labor supply impacts. While many bottom-up analyses of climate damages include the effects of sea-level rise and agricultural losses among a host of other causal channels, most do not include a direct effect of heat on human productivity (Heal, 2008). Given the medical literature on the subject, in conjunction with the country-level relationships that we document, this seems a non-trivial omission.

Fourth, the damages from future climate change may in fact be much more heterogeneous than previous studies have assumed, and, more specifically, that climate change may widen existing wealth and income disparities. Global warming may exacerbate income inequalities insofar as it pushes hotter countries, which on average tend to be much poorer, further away from the thermoregulatory optimum; and vice versa for relatively temperate and rich countries. The fact that cooler countries such as Canada or Russia may benefit from warmer weather – due, in part, to the advantageous impact on days and hours worked, and the relative productivity of those labor hours – is a difficult one, given the already complex international politics of climate change.

It seems that temperature does, in fact, play a role in determining the relative wealth of nations, through its impact on something very basic: human physiology. Inasmuch as anthropogenic climate change may make some parts of the world less friendly to human physiology (and some places more), it may have important consequences not only for local economic productivity, but also for global income inequality.

References

D. Acemoglu, S. Johnson, and J.A. Robinson. The colonial origins of comparative development: an empirical estimation. Nber working paper 7771, 2001.

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.

Marie Connolly. Some like it mild and not too wet: The influence of weather on subjective well-being. *Journal of Happiness Studies*, pages 1–17, 2013.

Frank C. Curriero, Karlyn S. Heiner, Jonathan M. Samet, Scott L. Zeger, Lisa Strug, and Jonathan A. Patz. Temperature and mortality in 11 cities of the eastern united states. *American Journal of Epidemiology*, 155(1):80–87, 2002.

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.

Joshua Graff Zivin and Matthew Neidell. Temperature and the allocation of time: Implications for climate change. Nber working paper 15717, 2010.

W.F Grether. Human performance at elevated environmental temperatures. *Aerospace Medicine*, 1973.

P.A. Hancock, J.M. Ross, and J.L. Szalma. A meta-analysis of performance response under thermal stressors. *Human Factors*, 49(5):851–877, 2007.

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, Joshua Graff Zivin, and Matthew Neidell. Mathematical reasoning deteriorates in warm environmental temperatures. June 2012.

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.

E. Huntington. *Climate and Civilization*. Yale University Press, 1915.

BF. Jones and BA. Olken. Climate shocks and exports. Technical Report 15711, NBER, 2010.

Matthew C. Keller, Barbara L. Fredrickson, Oscar Ybarra, Stéphane Côté, Kareem Johnson, Joe Mikels, Anne Conway, and Tor Wager. A warm heart and a clear head a warm heart and a clear head: The contingent effects of weather on mood and cognition. *Psychological Science*, 16(9):724–731, 2005.

E.M. Kilbourne. *Heat waves and hot environments*, pages 245–269. Oxford University PressE. Noji, 1997.

R.S. Kovats and S. Hajat. Heat stress and public health: A critical review. *Ann Rev Public Health*, 21:41–55, 2008.

Andreu Mas-Collel, Michael Whinston, and Jerry Green. *Microeconomic Theory*. Oxford University Press, 1995.

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

Antony J. McMichael and et al. International study of temperature, heat and urban mortality. *The International Journal of Epidemiology*, 37(5):1121–1131, 2008.

Baron de Montesqieu. *The Spirit of Laws*. 1750.

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

P.K. Pachauri and A. Reisinger, editors. *Climate Change 2007: Synthesis Report*. IPCC, Geneva, Switzerland, 2007.

K.C. Parsons. *Human Thermal Environments: The Effects of Hot, Moderate, and Cold Environments on Human Health, Comfort, and Performance*. Taylor and Francis, 2003.

J.D. Sachs, A.D. Mellinger, and J.L. Gallup. The geography of poverty and wealth. 2001.

Xavier Sala-i Martin. I just ran two million regressions. *American Economic Review*, 87(2), 1997.

S.J. Schiff and G.G. Somjen. The effect of temperature on synaptic transmission in hippcompal tissue slices. *Brain Research*, 345:279–284, 1985.

Wolfram Schlenker and Craig McIntosh. Identifying non-linearities in fixed effects models. Technical report, IRPS, UC San Diego, 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.

D. Wendt, LJC. van Loon, and van Marken-Lichtenbe. Thermoregulation during exercise in the heat. *Sports Medicine*, 37(8):669–680, 2007.

D.P. Wyon. The effects of moderate heat stress on typewriting performance. *Ergonomics*, 1974.

D.A. Yablonskiy, J.J.H. Ackerman, and M.E. Raichle. Coupling between changes in human brain temperature and oxidative metabolism during prolonged visual stimulation. *Proceedings of the National Academy of Sciences*, 97:7603–7606, 2008.

Garry Yohe, James Neumann, Patrick Marshall, and Holly Ameden. The economic cost of greenhouse-induced sea-level rise for developed property in the united states. *Climatic Change*, 32:387–410, 1996.

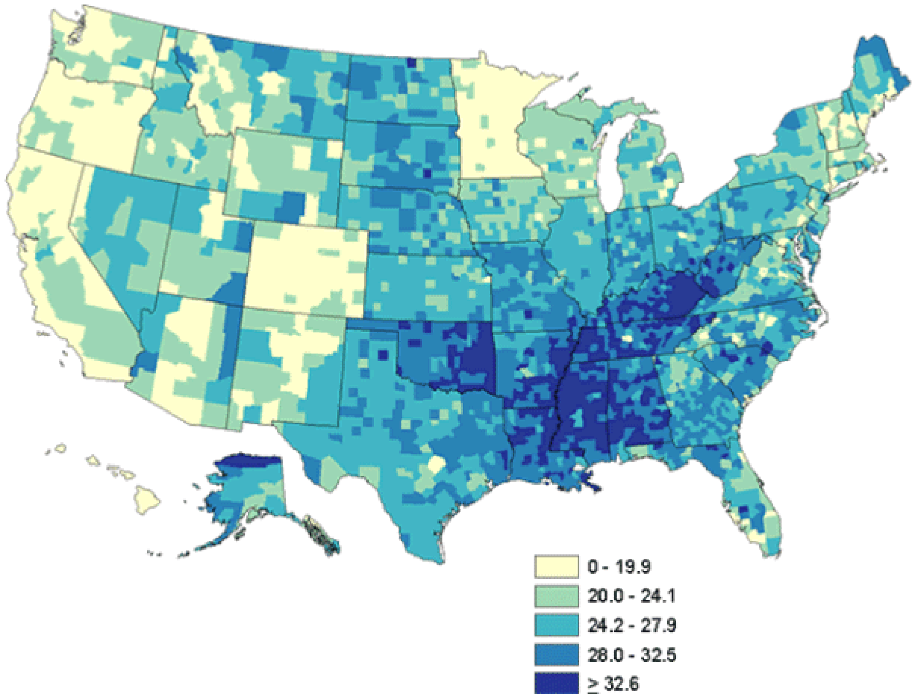

FIGURE 6.1. Percentage of adults who are physically inactive (2011)

PREFERRED MODEL WITH QUADRATIC IN TEMPERATURE				
	no lag	1-lag	5-lags	10-lags
	(9)	(10)	(11)	(12)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita
TEMPERATURE (C°)	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 (MM)	-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)
LAGGED TEMP		0.065***	0.041***	0.034***
		(0.013)	(0.014)	(0.013)
LAGGED TEMP^2		-0.002***	-0.001***	-0.001**
		(0.000)	(0.000)	(0.000)
LAGGED PRECIP		-0.004**	-0.002	-0.003*
		(0.001)	(0.002)	(0.001)
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				

TABLE 3. The contemporaneous impact of temperature and temperature squared on log income per capita, allowing for up to 10 lag terms in temperature.

PREFERRED MODEL STRATIFIED BY TEMPERATURE BIN:
VH = VERY HOT (>25C), H = HOT (20C-25C), M = MILD (15C-20C), C = COLD (10C-15C), VC = VERY COLD (10C>)

	no lag	1-lag	5-lags	10-lags
	(9)	(10)	(11)	(12)
VARIABLES	Log income per capita	Log income per capita	Log income per capita	Log income per capita
HUMAN CAPITAL INDEX	0.058	0.034	0.017	0.031
	(0.040)	(0.041)	(0.044)	(0.047)
LOG CAPITAL STOCK	0.334***	0.323***	0.287***	0.228***
	(0.022)	(0.022)	(0.024)	(0.028)
VH	-0.121***	-0.098***	-0.055**	**-0.042****
	(0.021)	(0.024)	(0.022)	(0.021)
H	-0.074***	-0.051***	-0.033**	**-0.030****
	(0.012)	(0.014)	(0.015)	(0.015)
M	0.002	0.002	0.003	**0.014**
	(0.018)	(0.019)	(0.017)	(0.014)
C	0.046***	0.030*	0.023	**0.030***
	(0.015)	(0.016)	(0.016)	(0.016)
VC	0.045***	0.030***	0.032***	**0.024*****
	(0.008)	(0.008)	(0.008)	(0.007)
PRECIPITATION	-0.005***	-0.003*	-0.003*	-0.004***
	(0.001)	(0.002)	(0.002)	(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				

TABLE 4. The impact of a +1°C hotter-than-average year temperature shock on log income per capita that year, stratified by temperature zone, allowing for up to 10 lags in temperature.

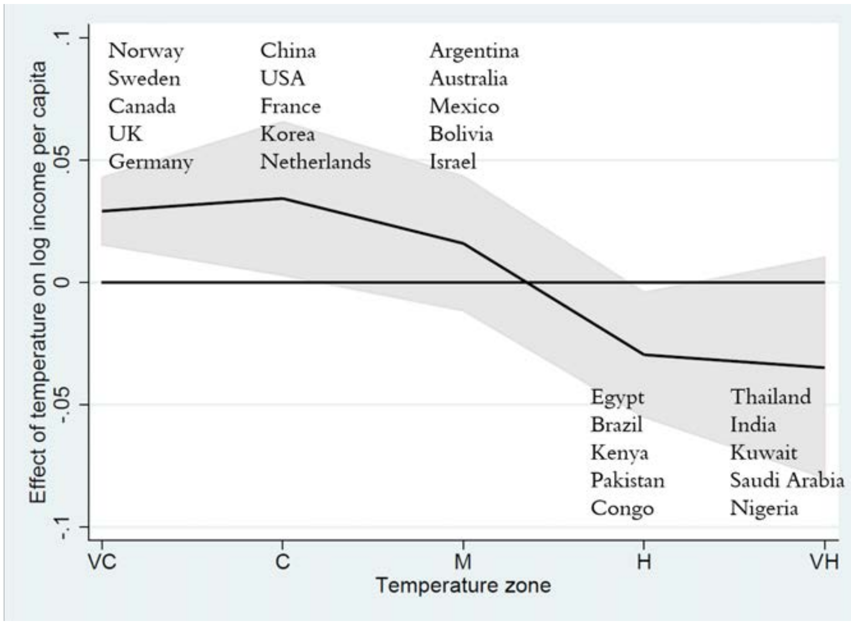

FIGURE 6.2. A 1°C warmer year results in negative output shocks in warmer countries, positive output shocks in cooler countries. Shaded bands denote 95% confidence intervals. Representative countries for each zone.

PREFERRED MODEL STRATIFIED BY TEMPERATURE BIN:
H = HOT (20C<), M = MILD (10C-20C), C = COLD (10C>)

	no lag	1-lag	5-lags	10-lags
	(9)	(10)	(11)	(12)
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)
L1WTEM_H		-0.048***	-0.019	-0.009
		(0.014)	(0.014)	(0.014)
L1WTEM_M		0.002	-0.004	-0.000
		(0.020)	(0.017)	(0.014)
L1WTEM_C		0.032***	0.021**	0.021**
		(0.009)	(0.009)	(0.009)
L1WPRE		-0.004***	-0.002*	-0.003*
		(0.001)	(0.002)	(0.001)
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				

TABLE 5. The pattern of positive impacts in colder countries, indeterminant impacts in temperature countries, and negative impacts in hot countries persists across multiple climate classifications

Agricultural Value Added as Percentage of GDP

Country Name	Average (1960-2008)
United Kingdom	1.714987
Germany	1.868958
United States	2.228492
Japan	2.466964
High income: OECD	3.140768
OECD members	3.481743
European Union	3.708261
High income: nonOECD	4.200496
Europe & Central Asia (all income levels)	4.449546
Middle East & North Africa (all income levels)	8.446136
Latin America & Caribbean (all income levels)	9.200476
East Asia & Pacific (all income levels)	9.533897
Sub-Saharan Africa (all income levels)	18.24035

TABLE 6. Agricultural value-added as proportion of GDP (Select countries and regions; 1960-2006)

PREFERRED MODEL STRATIFIED BY TEMPERATURE BIN: 3-YEAR BLOCK AVERAGES

	no lag
	(1)
VARIABLES	Log income per capita
HC	0.016
	(0.066)
LCK	0.336***
	(0.035)
WTEM5_VH	-0.189***
	(0.042)
WTEM5_H	-0.121***
	(0.026)
WTEM5_M	0.008
	(0.038)
WTEM5_C	0.078***
	(0.029)
WTEM5_VC	0.062***
	(0.016)
WPRE	-0.008**
	(0.004)
OBSERVATIONS	1,193
R-SQUARED	0.988
ROBUST STANDARD ERRORS IN PARENTHESES	
***** P<0.01, ** P<0.05, * P<0.1**	

TABLE 7. 3-year block averages for climate and economic variables

PREFERRED MODEL STRATIFIED BY TEMPERATURE BIN:
3-YEAR MOVING AVERAGES

	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
HUMAN CAPITAL INDEX	0.069*	0.052	0.025	0.042
	(0.040)	(0.041)	(0.044)	(0.046)
LOG CAPITAL STOCK	0.337***	0.329***	0.293***	0.240***
	(0.023)	(0.023)	(0.025)	(0.028)
VH	-0.174***	-0.215***	-0.160***	-0.141**
	(0.026)	(0.055)	(0.059)	(0.060)
H	-0.106***	-0.091***	-0.097***	-0.086**
	(0.014)	(0.032)	(0.034)	(0.037)
M	0.006	0.015	0.012	0.067*
	(0.023)	(0.047)	(0.045)	(0.039)
C	0.076***	0.037	0.061	0.076*
	(0.020)	(0.034)	(0.042)	(0.044)
VC	0.071***	0.043***	0.078***	0.070***
	(0.011)	(0.016)	(0.022)	(0.021)
PRECIPITATION	-0.009***	-0.005	-0.008**	-0.012***
	(0.002)	(0.003)	(0.004)	(0.004)
OBSERVATIONS	3,363	3,256	2,840	2,333
R-SQUARED	0.988	0.988	0.990	0.992
ROBUST STANDARD ERRORS IN PARENTHESES				
***** P<0.01, ** P<0.05, * P<0.1**				

TABLE 8. 3-year moving averages for climate and economic variables

	5-bin stratification				3-bin stratification			
	no lag	1-lag	5-lags	10-lags	no lag	1-lag	5-lags	10-lags
	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
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	Log income per capita	Log income per capita
VH	**-0.1227*** **	**-0.1007*** **	**-0.0554** **	**-0.0399***				
	(0.023)	(0.026)	(0.024)	(0.023)				
H	**-0.0806*** **	**-0.0638*** **	**-0.0433*** **	**-0.0328** **				
	(0.011)	(0.013)	(0.013)	(0.014)				
M	**0.0047**	**0.0042**	**0.0041**	**0.0131**				
	(0.018)	(0.019)	(0.017)	(0.014)				
C	**0.0453*** **	**0.0304***	**0.0252**	**0.0410** **				
	(0.015)	(0.016)	(0.016)	(0.017)				
VC	**0.0560*** **	**0.0383*** **	**0.0415*** **	**0.0368*** **				
	(0.008)	(0.008)	(0.008)	(0.007)				
H+VH * Temp (20C and above)					**-0.0976*** **	**-0.0791*** **	**-0.0467*** **	**-0.0309** **
					(0.012)	(0.014)	(0.013)	(0.013)
M * Temp (15C to 20C)					**0.0055**	**0.0049**	**0.0045**	**0.0157**
					(0.018)	(0.019)	(0.017)	(0.014)
C+VC * Temp (15C and below)					**0.0516*** **	**0.0352*** **	**0.0354*** **	**0.0313*** **
					(0.008)	(0.009)	(0.009)	(0.009)
Observations	3,118	3,022	2,650	2,190	3,118	3,022	2,650	2,190
R-squared	0.986	0.986	0.988	0.991	0.986	0.986	0.989	0.991

TABLE 9. Allowing for lagged impacts (up to 10 years).

	5-BIN CLASSIFICATION				3-BIN CLASSIFICATION			
	no lag	1-lag	5-lags	10-lags	no lag	1-lag	5-lags	10-lags
	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
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	Log income per capita	Log income per capita
HUMAN CAPITAL	0.077*	0.064	0.069	0.104**	0.077*	0.063	0.066	0.095**
	(0.041)	(0.041)	(0.044)	(0.047)	(0.041)	(0.041)	(0.044)	(0.046)
LOG CAPITAL STOCK	0.342***	0.333***	0.295***	0.226***	0.342***	0.333***	0.297***	0.230***
	(0.022)	(0.023)	(0.024)	(0.028)	(0.022)	(0.023)	(0.024)	(0.027)
VH	0.049*	0.015	0.016	**-0.000**				
	(0.026)	(0.026)	(0.024)	(0.023)				
H	0.004	-0.004	-0.014	**-0.017**				
	(0.016)	(0.016)	(0.016)	(0.017)				
M	-0.020	-0.013	-0.007	**0.015**				
	(0.022)	(0.021)	(0.018)	(0.015)				
C	-0.004	-0.005	0.000	**0.018**				
	(0.017)	(0.017)	(0.018)	(0.017)				
VC	0.003	0.002	0.013	**0.029*****				
	(0.008)	(0.008)	(0.009)	(0.008)				
HOT					0.024	0.005	-0.001	**-0.008**
					(0.015)	(0.015)	(0.014)	(0.014)
MILD					-0.020	-0.013	-0.007	**0.015**
					(0.022)	(0.021)	(0.018)	(0.015)
COLD					0.000	-0.000	0.008	**0.025*****
					(0.008)	(0.009)	(0.009)	(0.009)
OBSERVATIONS	3,363	3,256	2,840	2,333	3,363	3,256	2,840	2,333
R-SQUARED	0.987	0.987	0.989	0.992	0.987	0.987	0.989	0.992

Table 10. The impact of temperature shocks on log output per capita, controlling for country-specific temperature trends, stratified by 5 and 3 different temperature zone classifications.

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		10-lags	
	bottom third	top third	bottom third	top third	bottom third	top third	bottom third	top third
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
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	Log income per capita	Log income per capita
TEMPERATURE	-0.094***	0.001	-0.068***	-0.003	-0.030	0.005	-0.031	0.005
	(0.021)	(0.021)	(0.024)	(0.023)	(0.022)	(0.021)	(0.022)	(0.020)
PRECIPITATION	-0.001	-0.006**	0.001	-0.005**	0.006	-0.006**	0.003	-0.006**
	(0.004)	(0.002)	(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*	-0.083	0.152
	(0.084)	(0.101)	(0.089)	(0.102)	(0.116)	(0.110)	(0.144)	(0.122)
LOG CAPITAL STOCK	0.139***	0.664***	0.132***	0.660***	0.091*	0.631***	-0.003	0.571***
	(0.042)	(0.033)	(0.043)	(0.034)	(0.046)	(0.039)	(0.049)	(0.048)
OBSERVATIONS	795	746	765	723	657	631	535	516
R-SQUARED	0.952	0.969	0.954	0.969	0.962	0.972	0.971	0.978
ROBUST STANDARD ERRORS IN PARENTHESES								
***** P<0.01, ** P<0.05, * P<0.1**								

Table 11. The effect of air conditioning expenditure per capita (proxied by import value) on the relationship between population weighted average annual temperature and income per capita. Top third/bottom third denotes whether or not countries were in the top or bottom third of the sample in terms of average per capita AC expenditure for the sample period.

	country	acRank	pcAC
1.	Kuwait	1	.0016197
2.	United Arab Emirates	2	.0009899
3.	Brunei	3	.0007541
4.	Canada	4	.0005265
5.	Oman	5	.0005257
6.	Cyprus	6	.0004963
7.	Saudi Arabia	7	.0004307
8.	Bahamas, The	8	.0003085
9.	Libya	9	.0002516
10.	Netherlands	10	.0001903
11.	Greece	11	.0001901
12.	Norway	12	.0001882
13.	Sweden	13	.0001792
14.	Switzerland	14	.0001746
15.	Australia	15	.0001642
16.	Ireland	16	.0001598
17.	Finland	17	.0001391
18.	Portugal	18	.0001365
19.	Austria	19	.0001355
20.	Panama	20	.0001338
21.	Trinidad and Tobago	21	.0001249
22.	Spain	22	.0001137
23.	United Kingdom	23	.0001036
24.	Israel	24	.0001036
25.	Denmark	25	.0000897
26.	Gabon	26	.0000859
27.	France	27	.0000837
28.	New Zealand	28	.0000832
29.	Suriname	29	.0000728
30.	Germany	30	.0000718
31.	Malaysia	31	.000071
32.	Italy	32	.0000686
33.	Belize	33	.0000678
34.	Mauritius	34	.0000671
35.	United States	35	.0000649
36.	Fiji	36	.0000628
37.	Venezuela	37	.0000486
38.	Hungary	38	.0000413
39.	Iceland	39	.0000402
40.	Jordan	40	.0000389

FIGURE 6.3. Per capita AC expenditure by country, in hundreds of thousands of dollars.

7. Data Appendix

7.1. **Air Conditioning Data.** The AC imports data comes from WITS, the World Integrated Trade Solution data set (http://wits.worldbank.org/wits/), specifically, the United Nations COMTRADE database, a subset of WITS, which offers large country and period coverage of trading data (from 1962 and virtually all countries). According to the WITS User Manual, data from the data base is reported by statistical offices of each country to relevant international organizations.

In the COMTRADE database, data are recorded using several nomenclatures. The nomenclature we use for AC imports data is SITC Revision 1 (a trade classification maintained by the UN). The reason we chose this nomenclature is because it includes a product category of "air conditioning machines" and provides the longest time period (from 1962). These data are double-entried – that is, the same good is accounted for as an import by the importing country and an export by the exporting country, by two separate book-keeping entities. Given that imports are considered to be a more accurately recorded than exports, mostly due to the political economy of tariff revenue collection, we use import records to establish AC expenditure. The unit we use to measure trade flow is trade value (in million dollars). We use this measure instead of quantity measure because some countries report trade quantity in weight (kg), others in number of items, which are often inconsistent. Trade value, on the other hand, is consistently recorded for all countries and years.

We construct a variable that represents cumulative AC import value per capita for each country-year recorded in the income data above. One would be skeptical of using this as an explanatory variable if it is perfectly or very highly correlated with income. Regressing income per capita on AC expenditure per capita reveals that this is not the case (r = 0.52). The list of the top countries by per capita AC expenditure shows some very rich countries (e.g. Saudi Arabia) and poorer countries (e.g. Libya) as having high AC expenditures per capita. Obviously this measure understates AC expenditure by countries who are large domestic producers of AC units, notably the US, South Korea, and China. However, the AC expenditure variable as currently constructed allows us to identify countries that have the lowest levels of "thermoregulatory capital", which is the sub-population of interest.

7.2. **Robustness Checks for International Panel Results.** 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 the GDP residual (from TFP growth, for example) to climate fluctuations. There is a subtle but important interpretation issue here. Insofar as we believe that the evolution of capital stock variables – be that physical or human capital – is mediated by the ambient temperature in a country or region, we might still be able to attribute causal significance to temperature even if there is correlation between omitted capital stock variables and the temperature series. The rapid (or slow) accumulation of capital stock of an economy may be the proximal cause of higher (or lower) output or income, but temperature may have some ultimate causal role. For this to be true, however, it must be true that the temperature series and the omitted capital stock variables are not cointegrated (i.e. both cannot contain unit roots).[28]

[28]We discuss this issue in more detail in the Appendix.

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 (Table 10).

Table 10

Income per capita is often considered to be an AR1 process, or to be non-stationary. If the explanatory variable of interest – temperature in this case – is also non-stationary, this might lead to spurious correlation simply by virtue of the time-series properties of the data. Note that there is a distinction between non-stationarity of the series and whether or not there are time trends. It seems that income per capita and global average annual temperatures have clear time trends. Whether each country-specific temperature series in our panel (1) has a time trend, and (2) has a unit root is not immediately clear.

Population-weighted temperature (wtem), despite an apparent time trend for the global average, appears to be stationary across the panel, though some individual country series may have unit roots. The Pesaran (2007) panel unit root test suggests stationary of the wtem variable, even allowing for a series of lags. For the purposes of this analysis we assume average annual temperature to be a trend-stationary process. It is unclear exactly how one should account for time-trends in this context. Should we de-trend the temperature series for each country around its specific time-trend, or should it be with respect to the global average time-trend?

As Jones and Olken [2010] note, if hot climates were to cause low-quality institutions, which in turn cause low income, then controlling for institutions in a cross-sectional levels regression can have the effect of partially eliminating the explanatory power of climate, even if climate is the underlying fundamental cause. By the same token, if TFP growth was caused in part by climate variation especially over the long term, then controlling for TFP trends can have the effect of partially eliminating the explanatory power of climate. The fact that GDP per capita still has a clear time trend, even after controlling for capital accumulation and human capital (as well as institutions via country fixed effects), might be interpreted as a "secular" growth in TFP. We think that regressing this GDP residual on temperature in the presence of clear positive time trends in temperature might lead to spurious correlation insofar as we would be attributing "secular" TFP changes to temperature changes. But inasmuch as we believe that temperature is itself a determinant (if not the sole determinant) of TFP changes, then attempting to correct for this spurious correlation by detrending the temperature series would actually have the effect of partially eliminating the explanatory power of climate, just as in the cross-sections.

Moreover, we would need temperature and the GDP residual to be rising in cool countries and temperature to be falling and GDP residual to be rising in hot countries, or for the relative rates of increase to be significantly different among these groups. Neither seem to be true. The fact that the impact on hot countries becomes insignificant when controlling for country-specific time trends is somewhat puzzling, but we speculate that this might be due to 1) the reduced power due to reduced variation in the x-variable, and 2) the heterogeneity of the countries in the "hot" and "very hot" groups; Saudi Arabia and Mali are probably much more different than Canada and Switzerland.

7.3. **Estimating Non-linear Relationships with Fixed Effects.** Our panel estimation involves using the fixed effects regression to test for a non-linear relationship. An important distinction to bear in mind is whether or not the non-linear (single-peaked) relationship between temperature and productivity is global or "within-group." Is it that a warmer year leads to lower productivity if you're already in a hot climate, but higher productivity if you're in a cold climate, as the literature suggests? Or is it that small deviations around any point have a positive effect, but large deviations around any point have a diminishingly positive effect? We wish to test for whether the former is true – that is, whether there is a globally non-linear relationship between temperature and income. As such we insert quadratic terms directly into the estimation equation. That is, we allow the fixed effects estimator to de-mean the squared values of temperature, rather than taking the square of the de-meaned values (which is what one would do if one expected a "within-group" quadratic relationship). For a detailed description of using fixed effects to test for non-linear relationships, see Schlenker and McIntosh [2006].