

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 and Richard Sweeney 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.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Geoffrey Heal and Jisung Park. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Feeling the Heat: Temperature, Physiology & the Wealth of Nations
Geoffrey Heal and Jisung Park
NBER Working Paper No. 19725
December 2013
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 studies have explored the indirect links between climate and welfare (e.g. agricultural yield, violent conflict, or sea-level rise), few address the possibility of direct impacts operating through human physiology. This paper presents a model of labor supply under thermal stress, building on a longstanding physiological literature linking thermal stress to health and task performance. A key prediction is that effective labor supply – defined as a composite of labor hours, task performance, and effort – is decreasing in temperature deviations from the biological optimum. We use country-level panel data on population-weighted average temperature and income (1950-2005), to illustrate the potential magnitude of the effect. Using a fixed effects estimation strategy, we find that hotter-than-average years are associated with lower output per capita for already hot countries and higher output per capita for cold countries: approximately 3%-4% in both directions. We then use household data on air conditioning and heating expenditures from the US to provide further evidence in support of a physiologically based causal mechanism. 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
Department of Economics
Harvard University
Littauer Center
1805 Cambridge Street
Cambridge, MA 02138
jisungpark@gmail.com

1. INTRODUCTION

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 differences in income 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 ideas. They come from rather different fields, and each is commonplace in its own field, yet we believe their juxtaposition can add value. These ideas come from economics and physiology. The economic idea is that hotter countries tend to be poor, and most rich countries are found in temperate regions. The physiological idea 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. That hotter countries tend to be poor has been recognized for quite some time (Montesquieu [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 of both physical and intellectual tasks degrades with temperature is also well-established. While economists have noted this only recently - for example, Hsiang et al. [2012] show that student performance in standardized math tests falls as the temperature rises above the low 70s Fahrenheit - similar observations have a much longer history in the physiological 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]).

Our observation in this paper is that the physiological fact can help explain the economic one: that the fall-off in human performance with temperature can contribute to explaining the negative relationship between temperature and economic performance.

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 two different data sets: country-level panel data relating per capita income to average annual temperatures, and household expenditure data on heating and cooling in a cross-section of US households. 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 gradient in US households' willingness to pay for heating and cooling which depends on factors that relate to the physiology of thermoregulation.

Our most policy-relevant result is that annual climate shocks have historically had non-trivial impacts on GDP per capita, and that the direction and magnitude of these impacts depend crucially on the initial temperature zone or climatic zone. Hotter-than-average years lead to positive per capita output shocks in cold countries, and negative per capita output shocks in hot ones. 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 have systematically lower adverse climate impacts provides further suggestive evidence of a physiologically-mediated effect.

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.

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 presents further suggestive evidence of the physiological causal mechanism using household heating and cooling data from the United States. Section 7 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.

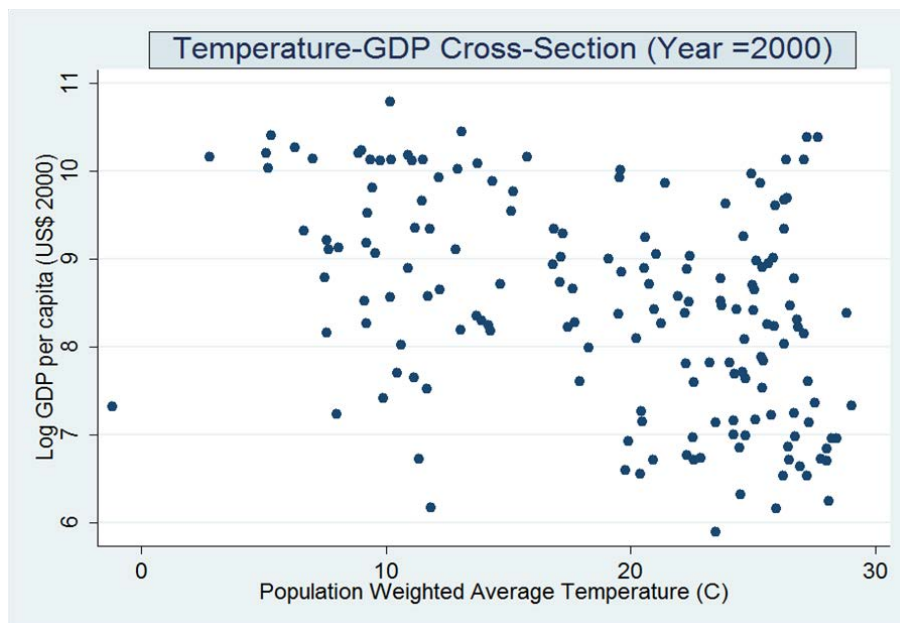


FIGURE 2.1. Countries by log income per capita and population-weighted average temperature

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 OECD countries, $+2^{\circ}\text{F}$ 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 $+6\text{-}7^{\circ}\text{F}$ in the future could lead to an average decrease of approximately $-13\%\text{-}14\%$ of GDP worldwide, a much higher figure than most bottom-up climate damage estimates suggest (Horowitz [2001]).

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.¹

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 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 (i.e. the relative densities of old and infirm), it has been suggested that a significant proportion of this variability is related to the extent to which nursing homes had 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.²

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]).

¹Of 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]).

²Lee 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.

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 (7.2)). 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.

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.³

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]).⁴ 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).⁵

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.⁶ 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).

³This 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, 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.

⁴Seppanen 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

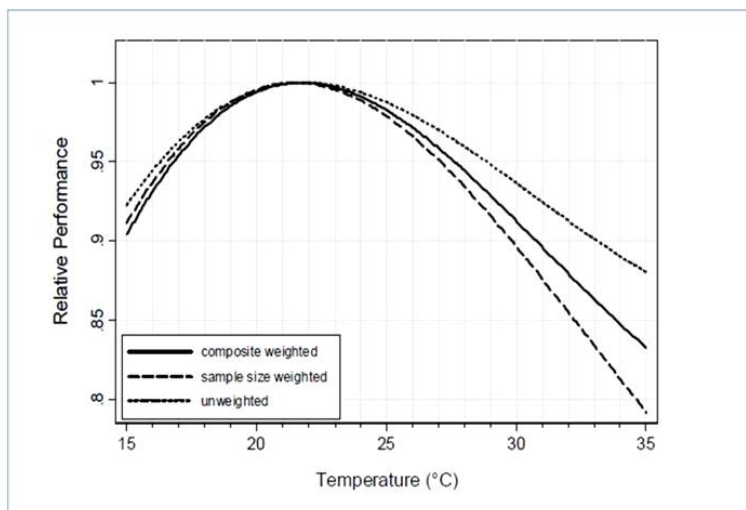


FIGURE 3.1. Task performance vs temperature. Maximum performance is normalized to 1 at 22 C. Source: Seppanen et al. [2006]

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

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

⁵The 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).

⁶We 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.

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 physiological factors that emerge as important influences on labor productivity as temperatures vary. The physiological factors are incorporated into 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 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.

An important concept is that of Basal Metabolic Rate or BMR. This is a measure of the amount of energy a person uses in just staying alive and carrying out the most basic bodily functions. In effect, it is the energy expended at rest, which is given off in the form of waste heat, and is a part of what each person has to lose to maintain core body temperature in the appropriate range. The actual metabolic rate increases with one's activity level and body mass, and decreases with age. A higher BMR increases the rate of energy use and the amount of energy to be dissipated to maintain core temperature. People with a high metabolic rate thus would be expected to have a greater need for air conditioning and a lesser need for heating. We will see below that BMR influences willingness to pay for both heating and cooling, increasing the latter and decreasing the former, as the theory suggests.

In economic terms we can say that the consequences of a body temperature away from the optimum are feeling excessively hot or cold, which we model as a loss of utility, a direct effect on welfare, and a drop in performance, leading to

a reduction in earning power. 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:

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

where E is the environmental (external) temperature and $P(T)$ is labor performance,⁷ 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 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

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

More generally we will write

$$W(L, A : E) = U((1 - L) AP, L, AT(E, A))$$

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:

$$(4.1) \quad U(Y, L, A, T) = Y + f(L, A, T)$$

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

⁷By using the word performance we intend to include a broader range of effects than would be indicated by productivity.

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

$$(4.2) \quad T(E, A) = \alpha + \beta E + g(A)$$

where α, β are constants and $g(\cdot)$ 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:

$$(4.3) \quad \frac{\partial W}{\partial A} = 0, \quad \frac{\partial W}{\partial L} = 0$$

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

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

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

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

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

$$(4.5) \quad (1 - L) P_T \beta + (1 - L) A P_{TT} \beta g_A + f_{A,T} \beta + f_{T,T} \beta g_A$$

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

$$(4.6) \quad -A P_T \beta + f_{L,T} \beta$$

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

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

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.**

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.*

Note that a change in core body temperature ΔT 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 metabolic rate. And 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.

4.2. Implications for empirical work. There are several points that emerge from this theoretical analysis that 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. In the section 7 we examine how spending on air conditioning and heating in the US responds to a change in the number of cooling degree days (heating degree days), and this result allows us to interpret this spending as an estimate of the welfare loss from exposure to more degree days.
- (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.⁸ 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.

According to point (2) above, we expect households’ expenditure on heating and cooling to depend not only on environmental factors such as local temperatures, but also on factors related to physiology and task performance. Because of 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 and possibly positively in cold ones. We do in fact find evidence of all three effects in the analyses that follow.

5. EMPIRICAL RESULTS I: INTERNATIONAL PANEL DATA

In the following two sections, we take our model to two separate data sets, one at the international level, another at the household level, to provide suggestive evidence of a physiological effect of climate on welfare.

As a crude first pass, we take the model’s key predictions to cross-country panel data. In effect, we revisit the age old question: what is the role of climate in explaining the relative wealth of nations?

⁸See Mas-Collel et al. [1995]

Whereas previous studies have focused on the role of heat or low latitude, we predict that deviations from the thermoregulatory optimum, as opposed to hotter temperatures per se, are what dictate the magnitude of climate-GDP impacts. Our analysis suggests that the relationship between temperature and income is nearly universal (i.e. not necessarily limited to poor countries as in DJO) 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.⁹

Figure 7.1

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 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 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). A new look at the cross-country panel data seems to confirm this intuition.

Second, in moving from a microeconomic model of thermoregulation to an analysis of macroeconomic variables, we must take into account 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 GDP coming from these industries to be more sensitive to temperature shocks¹⁰. Crucially for this analysis, the sensitivity of GDP to temperature stress may also be related to

⁹The 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

¹⁰In work currently in progress, we attempt to estimate differential impacts of temperature stress on particular sectors, using a panel of sub-national output data for US states and municipalities.

the degree of thermoregulatory capital available: that is, electrification, air conditioning, and access to heating systems and heat fuel. 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, noting that the impact of thermal stress on labor productivity is mostly contemporaneous.¹¹

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 .¹² Per capita output is increasing in effective labor supply.

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:¹³

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

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. We attempt to estimate this relationship by utilizing within-country variation in historical annual temperature realizations, using panel data analogous 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

¹¹As 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.

¹²We abstract away from population growth for simplicity.

¹³This is one reason why population-weighted average temperature is a more relevant metric than a raw geographic average.

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).

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.¹⁴

5.2.2. *International Economic Data.*

Following DJO (2008), we use the Penn World Tables Version 8.0 (Heston et al. [November 2012]). 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. To ensure robustness of the results, we run the same analyses using an alternative measure of income, taken from the UN national accounts. This data covers the same countries and years, but uses different inflation adjustments and price deflators. The results are broadly consistent across the different measures of income.

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:

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

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.:

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

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 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.

¹⁴We are, however, in the process of constructing a panel using data from the National Oceanic and Atmospheric Annals (NOAA) that imputes CDDs for all countries and regions of the world over the relevant time span.

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.¹⁵ We confirm that there exists a strongly negative cross-sectional relationship between temperature and income, particularly in countries where population-weighted average temperatures are above 20°. Of course, a key limitation of the existing cross-sectional analyses is that they may miss country-specific factors such as natural resource endowments or institutions. Researchers often point to the starkly different fortunes of North and South Korea as indicative of the crucial role of institutional factors.¹⁶

It is worth noting, furthermore, that previous studies which emphasize the monotonic cross-sectional relationship between temperature (latitude) and income (growth) may miss a significant component of the relationship, due to the limited number of cold countries in most samples. For example, in our sample there are only 5 countries which have annual average temperatures below 5° Celsius, even though a much larger number of countries have regions with very cold climates. More research is needed to uncover the temperature-income gradient within countries, especially those that have significant cold regions. At the very least, the temperature-income gradient in the cross-section provides us with an upper bound for any contemporaneous impact of temperature on income: be that positive or negative.¹⁷

The panel nature of our dataset allows us to control for time-invariant, country-specific unobservables that may influence income per capita: for instance, institutions or natural resource endowments (θ_i), and average climate (\bar{T}_i). In addition, we control for country-specific factors that may be changing over time by adding measures of country-specific capital stock directly. Using data from the Penn World Tables, we control for physical capital (log capital stock per capita) and human capital accumulation (in the form of an index).¹⁸ One way to think of this is that we are

15

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

¹⁶Selection 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.

¹⁷Selective migration based on the intensity of preferences for climate amenities (or adaptive capacity) notwithstanding.

¹⁸Both variables are taken from the Penn World Tables, version 8.0 (Heston et al, 2013).

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

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

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

Of course, 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.

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.

5.4. 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.¹⁹

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.²⁰

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). The results suggest a single-peaked relationship, with the implied peak again occurring somewhere between 15° and 20° Celsius (Figure 7.1). 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 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 and Spurious Correlation.*

¹⁹While 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.

²⁰These 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.

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. However there are of course alternative mechanisms which could lead to this relationship. We know for example that the connection between crop yield and temperature is highly non-linear, with yields increasing in temperature up to a point and then falling rapidly (Schlenker and Roberts [2006]). This suggests that looking across agricultural societies, we could find a single-peaked connection between temperature and output. One would not expect this relationship to persist across industrial countries, but it could be an explanation for our findings for a part of our sample. However, average agricultural value-added as a proportion of GDP in OECD countries is roughly 3% (over the period 1960-2006), and even in many developing economies less than 10%, suggesting that the effects cannot be totally attributable to decreases in agricultural yield (Table 6).

Table 6

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

Another concern is the potential for spurious correlation arising from secular but heterogeneous time trends in the temperature data. If some countries were warming (cooling) faster than others during the period of interest, we may incorrectly attribute secular changes in 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).²¹

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 7).

Table 7

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.

²¹We discuss this issue in more detail in the Appendix.

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 measure of 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.

Using this data, we stratify the sample based on whether the country had below or above median air conditioning penetration per capita in 1980. Table 8 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 above-median air conditioning penetration 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 (26), (28), (30), and (32) – the subset countries with above-median air conditioning – is shallower than that implied by the coefficients in columns (25), (27), (29), and (31) – which represent the subset of countries with below-median air conditioning.

Table 8

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 7.6). 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.

6. EMPIRICAL RESULTS II: HOUSEHOLD HEATING AND COOLING EXPENDITURES

There are clear limitations in using country-level aggregates to illuminate what is mostly an individual-level phenomenon (thermal stress leading to diminished productivity). And while much more research must be done to establish a clear causal picture, here we present a first-pass at further “micro-foundation” of the broader hypothesis. We use household data on heating and cooling expenditures to provide more support for the overarching research hypothesis at hand: that direct thermal stress of the human body may be driving part of the observed temperature-GDP relationship. We use data on US households from the Residential Energy Consumption Survey (RECS) to show that individuals not only suffer from direct

performance decline under thermal stress (as documented in the medical literature), but that they also respond in the form of changes in consumption patterns – notably on household heating and cooling. Importantly, we highlight the differential in willingness to pay that arises from physiological factors (e.g. basal metabolic rate), which are plausibly orthogonal to other omitted confounders.

We first estimate the relationship between thermal stress and willingness to pay for thermoregulation using repeated cross-sections. We then illustrate how this relationship depends on a key parameter of the model presented in Section 4: basal metabolic rates (BMR), which fall monotonically with age. We find that young people whose BMR is higher, and as such need to dispose of more heat, are willing to pay more per degree of cooling and less willing to pay per degree of heating, controlling for income, electricity prices, and detailed efficiency characteristics of housing and heating/cooling equipment.

Insofar as engaging in productive work raises the effective cost of thermal stress and/or raises one’s level of exertion, we would expect willingness to pay for thermoregulation to rise when working. As a crude first-pass at testing this hypothesis, we use data on self-reported work status to suggest that people who use their homes as places to work are willing to pay more for AC than those who are at home for similar amounts of time but not working, correcting for income differences.

While all of these results are robust to controls for a wide range of geographic and socioeconomic correlates including income, local energy prices, housing characteristics such as the degree of insulation or the age of heating/cooling equipment, and household race and employment status, the cross-sectional nature of the analysis implies that we cannot rule out the possibility of omitted variable bias. Once again, we consider these results to be suggestive and indicative of the need for further research.

6.1. Empirical Framework and Statistical Model.

We focus on the case of US households where, by and large, there is affordable access to thermoregulatory capital.²² The basic intuition is that if there is a link between thermal stress and welfare, we might observe evidence for this connection in consumer behavior in the market.

A general cross-sectional model of thermoregulatory demand may be written as:

$$X_{it} = f(T_{it}, P_{it}, M_{it}, Z_{it}) + \epsilon_{it}$$

where X_{it} is the level of heating or cooling expenditure for household i in year t , T_{it} is a vector of heating and cooling degree days experienced by that household, interpolated from readings of the weather stations nearest to that household in that year, P_{it} is a vector of physiological determinants of thermoregulatory demand, M_{it} is a vector of mechanical determinants of the per unit effectiveness of expenditure on thermoregulation (level of housing unit insulation, efficiency of AC or heating equipment installed, etc.), Z_{it} is a vector of socioeconomic characteristics and other possible determinants of demand (income, race, working status, local price of electricity and energy inputs etc.), and $f(\cdot)$ is a “nonparametric” conditional mean function of X_{it} . Thus, $f(\cdot)$ and Z_{it} are defined so that $E(\epsilon|T, P, E, Z) = 0$.

The ideal research design would consist of a controlled experiment using a panel of households for which temperature shocks of varying levels are randomly assigned.

²²As of 2009, 87% of US households are equipped with air conditioning. See Figure (7.3) in Data Appendix.

Random assignment ensures that thermal stress is independent of other factors, particularly those which may be unobserved.

However, data that replicates such an experiment does not exist. Thus we use four waves of a nationally representative cross-section of US households to estimate this demand function (EIA [2005]). To our knowledge, this is the first attempt at estimating demand for “thermoregulation” as a flow good.²³The repeated cross-section comprises roughly analogous data from years 1997, 2001, 2005, and 2009.

Before proceeding, we note once again the obvious potential for several types of bias in our estimation. First, any omitted variables that simultaneously affect the level of temperature stress and the demand for thermoregulation (for example, self-selection of individuals of different heat and cold “tolerance” into warmer or cooler regions) may bias the results. Second, insofar as the demand for thermoregulation and the market price for thermoregulation – notably the price of electricity and energy inputs – are simultaneously determined, we will not be estimating the true demand curve from the cross-section. The fact that energy prices are highly regulated, and price-changes (particularly for consumers) relatively infrequent and administered by regulatory agencies, suggests that this bias may be limited. Finally, insofar as the expenditure data represents a snapshot of spending in time (one year), it likely omits some of the fixed costs associated with heating and cooling. At the same time, we note that our aim is to elucidate the plausibility of a particular causal mechanism – thermoregulation – and not necessarily the precise estimation of demand elasticities or welfare impacts associated with particular policies.

6.2. Data.

Our data comes from the Residential Energy Consumption Survey (RECS), a nationally representative sample of a cross-section of housing units from the US, administered by the Energy Information Agency (EIA). Survey data on housing unit characteristics, fuel and electricity usage patterns, and household demographics is combined with data from energy suppliers to these homes to estimate energy costs and usage for heating, cooling, appliances and other end uses. Here we focus on the 2005 survey, which collected data from 4,380 households in housing units statistically selected to represent the 113 million housing units that are occupied as a primary residence. We replicate the following analysis for years 1997 and 2001, and find nearly identical results.²⁴

The RECS is a rich data set, including key household-specific variables that allow us to estimate the relationship between thermal stress and individual behavior. Notably, the data records annual cooling and heating degree days associated with

²³While what we call thermoregulation is undoubtedly related to the “climate amenity” (for example, Maddison [2003]), the demand for which has indeed been estimated before, we believe that the amenity value of climate as historically estimated overlooks the core physiological feature of thermoregulation as we have defined it. For instance, most hedonic estimates of climate amenities ignore the variable costs associated with heating and cooling homes in their regressions.

²⁴Parts of the 2009 survey are currently available and also yield similar results.

each of the households surveyed.²⁵ Combined with data for annual dollar expenditures on air conditioning and heating, and fairly detailed controls for the efficiency characteristics of the heating and cooling systems in each household, this allows us to estimate a mapping between the degree of external thermal stress experienced by a specific household and the demand for thermoregulation. The use of cooling and heating degree data, which represents a measure of the cumulative heat or cold stress experienced by a region over a certain amount of time (in this case, one year), gives us a more accurate measure of temperature stress than using annual average temperatures as explanatory variables, though in theory even more disaggregated measurements are possible (Auffhammer et al. [2013]).

We run two sets of standard OLS regressions, one with annual household cooling expenditures as the dependent variable of interest, the other with heating expenditures as the dependent variable. The central idea is to estimate a (static) demand for thermoregulation, controlling for important covariates, and then to examine the extent to which this demand depends on thermoregulatory factors such as body metabolism, which, as our model shows, may be important in interpreting individuals' decisions under thermal stress.

6.2.1. *Summary Statistics.*

We report key summary statistics for the 2005 data in Table 9.²⁶

As a whole, the US experiences more heating degree days than cooling degree days, by roughly a factor of 3 to 1, though clearly there is vast regional variation. The average cooling expenditures for US households in 2005 was \$256, with a standard deviation of \$250, and a right-skewed distribution. Average heating expenditures are roughly \$550 per household, but vary considerably more than cooling expenditures.

There is considerable variation in the average age of households, a feature that we utilize to test the significance of the physiological mechanism through which thermal stress affects demand. While not presented in the table above, it is worth noting that roughly 47% (2,076) of households report someone residing at home for the majority of the day, and approximately 5% (224) report that someone works from home for the majority of the day.

6.3. Results.

Table 10 highlights the intuition that willingness to pay for thermoregulatory capital will depend on the level of thermal stress. A key hypothesis that emerges from the model in section 4 is that willingness to pay for thermoregulatory capital will rise with thermal stress in either direction away from the biological optimum. We take CDD and HDD as measures of the accumulated amount of thermal stress

²⁵Cooling degree days (CDD) are a measure of how hot a location was over a period of time, relative to a base temperature. In the RECS data, the base temperature is 65° Fahrenheit, and the period of time is one year. The CDD for a single day is the difference between that day's average temperature and the base temperature if the daily average is greater than the base; it is zero if the daily average temperature is less than or equal to the base temperature. The number of CDD's for a longer period of time is the sum of the daily cooling degree-days for the days in that period. Note that some studies compute CDD's using a base that is higher or lower than 65. The computation is performed in an analogous manner for heating degree days (HDD).

²⁶We note that a non-trivial proportion of the data recorded zero entries for heating and cooling expenditure, which may represent misreported data. The table above reports summary statistics after dropping these observations. The full dataset is reported in the appendix.

over the course of the year. Controlling for the complete vector of available physiological, mechanical, and socioeconomic determinants of demand, we find that both cooling and heating demand depend significantly on the amount of thermal stress.

Table 10

US households spend on average 19 cents per cooling degree day per year on cooling, and roughly 8 cents per heating degree day per year. While both regressions include the complete vector of controls available in our data set, we suppress them to highlight the relationship between thermal stress and average expenditure.

Of course, raw expenditure on heating and cooling can be misleading, if not corrected for energy prices or the mechanical efficiency of the thermoregulatory capital used. For a given level of thermal stress, one would expect households facing higher energy prices to have lower optimal heating and cooling demands – and vice versa. Ideally, we would be able to correct for differences in electricity and fuel oil prices across individual households. Household-specific electricity prices, however, were not readily available. We have, instead, constructed dummy variables for each of the 13 census divisions into which the surveyed households are classified, noting the significant similarities within divisions in terms of electricity prices and climatic conditions. These are included in the regression results throughout the analysis, though they are suppressed from the tables below.²⁷

Similarly, for a given level of thermal stress, larger, less efficiently insulated houses will require more energy (and thus greater total heating and cooling expenditures) to bring them to a particular target temperature. The same would be true for houses with old and inefficient air conditioners and heating systems. We control for these energy efficiency characteristics by including a suite of variables including the number of rooms, age of housing unit, age of AC unit, and the type of heating or cooling system, which is recorded at the household level (Figures (7.4) and (7.5)).

It is clear is that demand for thermoregulation depends strongly on the degree of heat and cold stress experienced locally. We now turn to the issue of whether or not this connection can be attributed to human physiology, as our model and much of the medical literature suggests.

6.3.1. *Establishing Causal Significance of Physiology.*

One way to explore whether or not the thermoregulatory mechanism is driving this relationship is to test whether heating and cooling expenditures vary across levels of basal metabolic rates (BMR), controlling for other relevant observables such as income or local climate.

It is well-established in the medical literature that the relationship between BMR and age is systematic in nature and substantial in magnitude, falling monotonically with age by on average 40% between ages 1 and 60 (Mitchell [1962]: Figure 6.1).

While direct data on BMR is not available as part of the EIA data set, we observe the ages of members of each household, making it possible to use age variables as proxies for BMR. The hypothesis is that older individuals will on average optimally choose to spend less (more) on thermoregulation for a given level of heat stress (cold stress)²⁸. So we expect cooling expenditure to fall, and heating expenditure

²⁷We include these variables in our reporting of robustness checks in the data appendix.

²⁸This assumes that the thermostat is set to correspond to the mean preference among members of the household – we could relax this assumption and allow for it to be set by the youngest or




FIGURE 6.1. Basal Metabolic Rate (BMR) declines with age (Source: Mitchell [1962])

to rise, with average age, as older people need to dispose of less heat. In our model, this corresponds to systematic heterogeneity in the $g(A)$ function (which reflects metabolic rate, which in turn affects how ambient temperatures map onto core body temperature).

Indeed this is what we find. There is a clear relationship between household thermoregulatory demand and implied BMRs. Table 11 reports the results of regressing heating and cooling expenditures on average household age and the full vector of controls. We present two different specifications of the relevant age variable: average age, which captures the average BMR of the household, and the age of the youngest and oldest members of the household, which we might expect matter more for cooling and heating respectively, as these will be the members who need to cool or heat most in response to high or low temperatures. Columns (1) and (3) present results for average age, and (2) and (4) present results using youngest and oldest members' ages.

Table 11

As we expect, households with lower average BMRs, as proxied by average age of its members, spend significantly less on air conditioning per degree of thermal stress, and significantly more on heating. The magnitude is non-trivial. A household with an average age of 20 spends roughly 15% (\$28) more per year on AC and 12% (\$54) less on heating than an otherwise equivalent household with an average age of 60, assuming both are exposed to the same level of thermal stress throughout the year. If we use youngest and oldest ages as the relevant proxy for BMR, we find a similar result. It is worth noting that we observe this effect even while controlling for the experienced climate of each household; that is, we account for the possibility that older individuals may tend to live in milder climates. While this does not prove that thermal stress is affecting individuals' consumption decisions with respect to

oldest members of the household, i.e. those with the most intense preferences with respect to thermoregulation.

thermoregulation, we believe it suggests that there is an important physiological component at play.

Insofar as metabolic rates affect the overall energy balance of the housing unit, one might also expect that the number of people in the house would matter. Each person generates of the order of 100 watts, so that a family of four can contribute to the warming of their home, increasing demand for cooling. The data is suggestive on this front as well. Households with more individuals spend significantly more (roughly \$14 or 7% more per degree day per person) for cooling. The correlation between number of household members and heating expenditures is not significantly different from zero.

6.3.2. *The Effects of Thermal Stress on Labor Performance.*

We construct a dummy variable denoting whether the survey respondent stated that a member of the household worked from home during the day, and one denoting whether at least one individual is at home during the day but not necessarily for work. Our model suggests that expenditure on both AC and heating should be higher for households with someone at home who is working than for those with someone at home but not working, assuming that idiosyncratic preferences with respect to thermoregulation do not vary systematically across individuals who stay at home without working and those who work from home. Here, we report the results using 2009 data, which documents work-from-home status for a larger sample of households.

Table 12

As shown in Table 12, demand for thermoregulation per degree day is significantly higher in households that have someone at home during the day. Those with someone working from home have higher expenditures for heating and cooling than those that have someone who is at home during the day but not working, and by a significantly larger amount for cooling. It appears that individuals might be taking into account the impact of thermal stress on labor productivity – and, as consequence, income – in their decisions with regard to purchasing thermoregulatory services.

6.3.3. *Cultural or Genetic Adaptation?*

Some studies have suggested that sensitivity to thermal stress may vary substantially according to race, ethnicity and/or cultural origin. Hsiang et al. [2012] find that mathematics scores for Caucasian children are dramatically more temperature sensitive than for those of Hispanic or African American children. While social scientists have historically balked at suggestions of causal connections between genetic or physiological differences and economic behavior, due perhaps to a legacy of pseudo-scientific justifications for imperialism (for example, see Ridgeway, 1908), recent medical literature suggests that genetic adaptation plays a significant role in thermoregulation. Ruff [1994] finds that absolute body breadth matters significantly for thermoregulation in humans as in other mammals. Caucasians, who generally tend to have narrower body breadth, having been genetically adapted to colder climes, would in principle have more difficulty thermoregulating under heat stress.

In our sample, controlling for income, housing unit characteristics, and household size and composition, households identifying as Hispanic or Native Hawaiian or Pacific Islander spend significantly less on cooling per degree day than the average

household, suggesting a higher tolerance for hot temperatures (Table 13). We find less racial variation in willingness to spend per degree day on heating (Table 14).

Table 13

Table 14

Of course, these regressions are merely associative. Yet they suggest the need for future research. Insofar as the distributional consequences of future climate change depend in large part on who is affected where, understanding more clearly the links between even the most primitive (e.g. racial and ethnic) factors and welfare impacts of heat stress may be worth investigating.

7. CONCLUSION

Three main conclusions emerge from our analysis.

Firstly, there is a clear and significant relationship between temperature and human performance of a range of economically significant tasks. An extensive physiological literature attests to this point, with recent economic studies consistent with these findings.

Secondly, cross-country panel data suggests a single-peaked and significant relationship between output per capita and temperature. While we are far from being able to quantify the exact magnitude of the causal impact of temperature on income, our results suggest that the physiological impact of thermal stress on labor productivity likely plays a causal role. The magnitude of these temperature impacts is large, and suggest that many simulation-based estimates of the costs of climate change may be downward-biased. These results also suggest that integrated assessment models should include the direct effect of climate change on labor productivity, among other direct impacts of thermal stress on the human organism.

Finally, people are willing to pay significant sums for thermoregulation. Furthermore, willingness to pay for thermoregulation varies with basal metabolic rate in the direction that our theory would suggest: younger individuals have a higher willingness to pay for cooling and a lower willingness to pay for heating than older individuals, suggesting that physiological factors significantly influence economic behavior. Using one's home for income generation increases the willingness to pay for thermoregulation (controlling for income, climate, and electricity prices), consistent with the idea that people want to avoid the impact of heat stress on their productivity.

The bottom line appears to be that temperature affects economic performance at the micro and macro levels enough to be a significant explanatory variable in cross-country comparisons. This suggests a novel and under-emphasized mechanism through which climate change may affect economic activity.

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.
- EIA. Residential energy consumption survey: Energy information agency, 2005.
- 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.
- Alan Heston, Robert Summers, and Bettina Aten. Penn world table version 7.1., Technical report, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, November 2012.
- J.K. Horowitz. The income-temperature relationship in a cross-section of countries and its implications for global warming. Technical report, 2001.
- Solomon M. Hsiang. Temperatures and cyclones strongly associated with economic production in the caribbean and central america. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372, 2010.
- Solomon M. Hsiang, 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.
- 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 Press, 1997.
- R.S. Kovats and S. Hajat. Heat stress and public health: A critical review. *Ann Rev Public Health*, 21:41–55, 2008.
- David Maddison. The amenity value of the climate: the household production function approach. *Resource and Energy Economics*, 25(2):155–175, 5 2003. doi: [http://dx.doi.org/10.1016/S0928-7655\(02\)00024-6](http://dx.doi.org/10.1016/S0928-7655(02)00024-6). URL <http://www.sciencedirect.com/science/article/pii/S0928765502000246>.
- 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.
- H. Mitchell. *Comparative Nutrition of Man and Domestic Animals*, volume 1. Academic Press, 1962.
- Baron de Montesquieu. *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.
- C.B. Ruff. Morphological adaptation to climate in modern and fossil humans. *American Journal of Physical Anthropology*, 1994.
- 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 hippocampal 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.

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*** (0.012)	0.076*** (0.013)	0.067*** (0.013)	0.052*** (0.012)
TEMPERATURE SQUARED	-0.004*** (0.000)	-0.003*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)
PRECIPITATION (MM)	-0.004*** (0.001)	-0.002 (0.002)	-0.003* (0.001)	-0.004*** (0.001)
HUMAN CAPITAL INDEX	0.081** (0.040)	0.067* (0.041)	0.060 (0.043)	0.071 (0.046)
LOG CAPITAL STOCK	0.338*** (0.022)	0.327*** (0.022)	0.293*** (0.024)	0.234*** (0.027)
LAGGED TEMP		0.065*** (0.013)	0.041*** (0.014)	0.034*** (0.013)
LAGGED TEMP^2		-0.002*** (0.000)	-0.001*** (0.000)	-0.001** (0.000)
LAGGED PRECIP		-0.004** (0.001)	-0.002 (0.002)	-0.003* (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 (0.040)	0.034 (0.041)	0.017 (0.044)	0.031 (0.047)
LOG CAPITAL STOCK	0.334*** (0.022)	0.323*** (0.022)	0.287*** (0.024)	0.228*** (0.028)
VH	-0.121*** (0.021)	-0.098*** (0.024)	-0.055** (0.022)	-0.042** (0.021)
H	-0.074*** (0.012)	-0.051*** (0.014)	-0.033** (0.015)	-0.030** (0.015)
M	0.002 (0.018)	0.002 (0.019)	0.003 (0.017)	0.014 (0.014)
C	0.046*** (0.015)	0.030* (0.016)	0.023 (0.016)	0.030* (0.016)
VC	0.045*** (0.008)	0.030*** (0.008)	0.032*** (0.008)	0.024*** (0.007)
PRECIPITATION	-0.005*** (0.001)	-0.003* (0.002)	-0.003* (0.002)	-0.004*** (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.

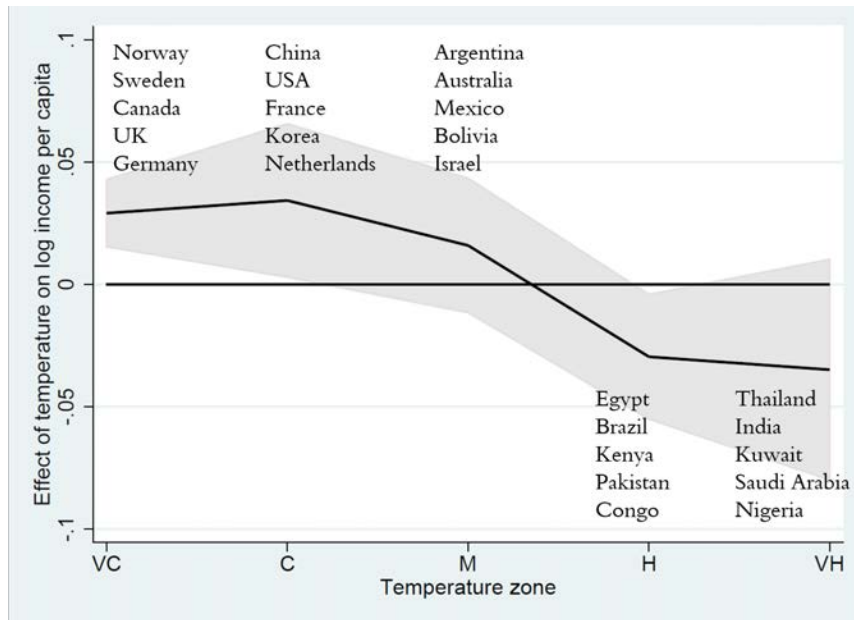


FIGURE 7.1. A 1°C warmer year results in negative output shocks in warmer countries, positive output shocks in cooler countries. Shaded bands denote 95% confidence intervals. Representative countries for each zone.


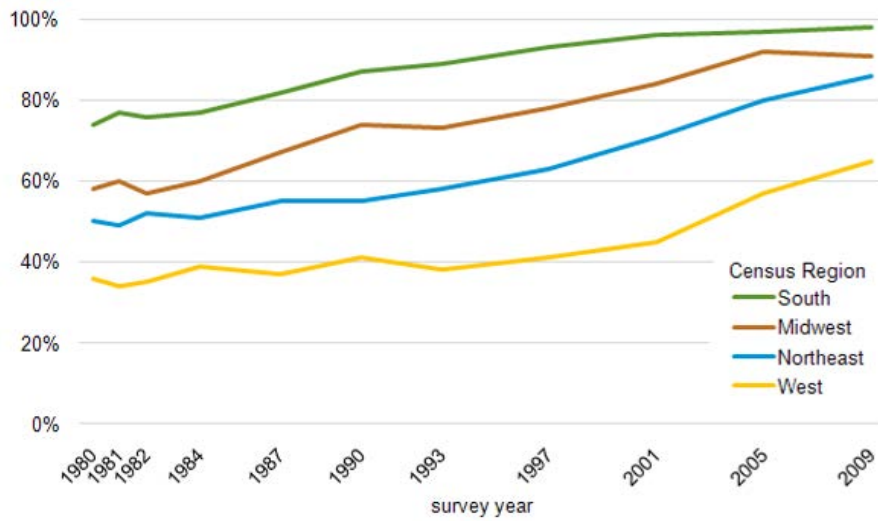


FIGURE 7.2. Percentage of adults who are physically inactive (2011)



Source: U.S. Energy Information Administration, 2009 Residential Energy Consumption Survey

FIGURE 7.3. Air conditioning penetration over time in US households (source: EIA, 2009)

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, indeterminate 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)

VARIABLES	5-BIN CLASSIFICATION				3-BIN CLASSIFICATION			
	no lag (13)	1-lag (14)	5-lags (15)	10-lags (16)	no lag (17)	1-lag (18)	5-lags (19)	10-lags (20)
	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 7. 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 WITH AIR CONDITIONING PER CAPITA (BELOW/ABOVE MEDIAN IN 1980)								
	below (25)	above (26)	below(lag1) (27)	above(lag1) (28)	below(lag5) (29)	above(lag5) (30)	below(lag10) (31)	above(lag10) (32)
VARIABLES	lpincome	lpincome	lpincome	lpincome	lpincome	lpincome	lpincome	lpincome
HOT	-0.086*** (0.012)	-0.079*** (0.025)	-0.070*** (0.014)	-0.052* (0.030)	-0.048*** (0.014)	-0.016 (0.026)	-0.038*** (0.014)	0.004 (0.027)
MILD	-0.050** (0.020)	0.076*** (0.026)	-0.031 (0.023)	0.044* (0.026)	-0.005 (0.019)	0.025 (0.025)	0.007 (0.016)	0.028 (0.022)
COLD	0.042*** (0.015)	0.042*** (0.010)	0.029* (0.016)	0.028*** (0.011)	0.017 (0.017)	0.036*** (0.011)	0.039** (0.017)	0.022** (0.010)
HUMAN CAPITAL	0.307*** (0.058)	-0.204*** (0.057)	0.291*** (0.058)	-0.222*** (0.057)	0.326*** (0.064)	-0.249*** (0.059)	0.409*** (0.074)	-0.230*** (0.061)
LOG CAPITAL STOCK	0.258*** (0.023)	0.644*** (0.037)	0.247*** (0.024)	0.629*** (0.038)	0.214*** (0.025)	0.578*** (0.041)	0.160*** (0.029)	0.509*** (0.046)
WPRE								
L1WIEM_H			-0.041*** (0.014)	-0.049 (0.030)	-0.026* (0.014)	-0.015 (0.028)	-0.016 (0.014)	-0.009 (0.029)
L1WIEM_M			-0.037 (0.024)	0.061** (0.027)	-0.024 (0.020)	0.026 (0.025)	-0.006 (0.015)	0.011 (0.024)
L1WIEM_C			0.031* (0.017)	0.032*** (0.011)	0.021 (0.018)	0.021* (0.011)	0.023 (0.017)	0.028*** (0.010)
OBSERVATIONS	2,209	1,154	2,137	1,119	1,861	979	1,529	804
R-SQUARED	0.984	0.979	0.984	0.979	0.987	0.982	0.990	0.986

TABLE 8. 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. Below/above denotes whether or not countries were below or above median per capita AC expenditure in 1980.

Variable	Mean	SD	Min	Max	N	Description
HDD	4214.69	2041.13	37.00	9419.00	3418	heating degree days: base 65F
CDD	1619.91	908.47	0.00	5059.00	3418	cool degree days: base 65F
Cooling	256.52	249.71	1.00	1897.00	3418	electric AC use (estimated) in dollars
Heating	549.47	479.89	12.00	4516.00	3418	total space heat use in dollars
Avgage	42.32	20.34	7.67	95.00	3418	average age of people living in the household

TABLE 9. Summary statistics of EIA RECS data

	(1)	(2)
	Cooling	Heating
CDD	0.187***	
	(0.006)	
HDD		0.079***
		(0.003)
Observations	2,321	4,230
R-squared	0.655	0.483
Robust standard errors in parentheses	*** p<0.01, ** p<0.05, * p<0.1	

TABLE 10. Spending on air conditioning and heating per household depends strongly on local cooling and heating degree days

2005 RECS SURVEY RESULTS				
VARIABLES	(1)	(2)	(3)	(4)
	AC	AC	HEATING	HEATING
	expenditure	expenditure	expenditure	expenditure
CD65	0.193***	0.181***		
	(0.011)	(0.010)		
HD65			0.067***	0.047***
			(0.006)	(0.008)
HHINCOME	0.000***	0.000***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
AVGAGE	1.139***		0.870*	
	(0.251)		(0.523)	
CDD_AVGAGE	-0.001***			
	(0.000)			
HDD_AVGAGE			0.000	
			(0.000)	
AGEHHMEMY		0.524**		
		(0.254)		
CDD_YOUNGEST		-0.001***		
		(0.000)		
OLDEST				-0.575
				(0.568)
HDD_OLDEST				0.001***
				(0.000)
OBSERVATIONS	4,380	3,367	4,230	4,230
R-SQUARED	0.664	0.671	0.489	0.491
ROBUST STANDARD ERRORS IN PARENTHESES				

*** P<0.01, ** P<0.05, * P<0.1

TABLE 11. Results indicate a strong relationship between household BMR (proxied by age) and willingness to pay for thermoregulatory spending.

**VARIATION IN COOLING & HEATING EXPENDITURE PER
DEGREE DAY BY WORK-FROM-HOME STATUS (RECS 2009)**

VARIABLES	(1)	(2)
	COOLING EXPEND	HEATING EXPEND
CD65	0.156***	
	(0.006)	
HD65		0.048***
		(0.004)
WORKATHOME	-49.148***	-61.446*
	(14.457)	(31.606)
CDD_WORKATHOME	0.074***	
	(0.014)	
HDD_WORKATHOME		0.020**
		(0.010)
ONLYATHOME	-1.438	-0.364
	(5.737)	(10.915)
CDD_ONLYATHOME	0.012**	
	(0.006)	
HDD_ONLYATHOME		0.004
		(0.003)
OBSERVATIONS	12,082	11,494
R-SQUARED	0.548	0.466

TABLE 12. Whether or not someone works from home raises willingness to pay per degree day for both heating and cooling, but by a significantly larger amount for cooling.

**VARIATION IN COOLING EXPENDITURE PER DEGREE
DAY BY RACE/ETHNICITY**

VARIABLES	(8) COOLING EXPEND
CD65	0.181***
	(0.010)
CD65 * "BLACK OR AFRICAN-AMERICAN"	-0.012
	(0.015)
CD65 * "AMERICAN INDIAN OR ALASKA NATIVE"	-0.075*
	(0.042)
CD65 * "OTHER"	-0.004
	(0.021)
CD65 * "HISPANIC"	-0.034**
	(0.017)
CD65 * "ASIAN"	-0.039
	(0.025)
CD65 * "NATIVE HAWAIIAN OR OTHER PACIFIC ISLANDER"	-0.074**
	(0.031)
OBSERVATIONS	3,367
R-SQUARED	0.671

TABLE 13. Hispanic and Native Hawaiian or Pacific Islander families tend to spend less on cooling per degree day, suggesting possible genetic or cultural differences in sensitivities to thermal stress

**VARIATION IN HEATING EXPENDITURE PER DEGREE
DAY BY RACE/ETHNICITY**

VARIABLES	(11) HEATING EXPEND
HD65	0.067*** (0.006)
HD65 * "BLACK OR AFRICAN-AMERICAN"	0.010 (0.009)
HD65 * "AMERICAN INDIAN OR ALASKA NATIVE"	-0.003 (0.035)
HD65 * "OTHER"	-0.007 (0.011)
HD65 * "HISPANIC"	-0.008 (0.010)
HD65 * "ASIAN"	-0.024* (0.014)
HD65 * "NATIVE HAWAIIAN OR OTHER PACIFIC ISLANDER"	-0.022 (0.017)
OBSERVATIONS	4,230
R-SQUARED	0.489

TABLE 14. Racial variation in willingness to spend per degree day on heating is less pronounced than variation in cooling expenditures.




FIGURE 7.4. Distribution of average year constructed for US households in EIA dataset (2005)

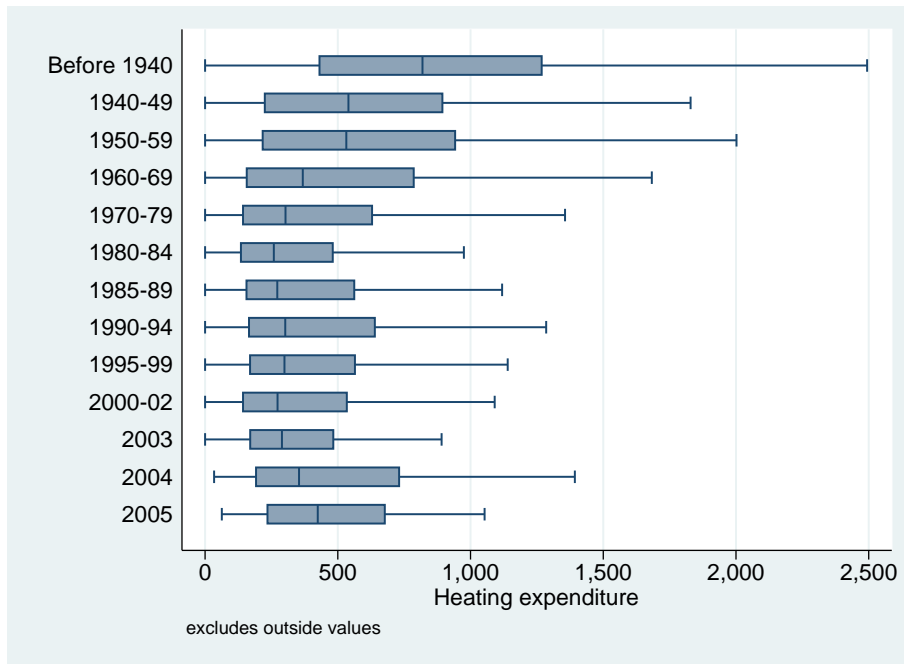


FIGURE 7.5. Heating expenditure depends on housing characteristics such as insulation, which is correlated with age of housing unit

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

(A) Maddala and Wu (1999) Panel unit root test (MW)				(B) Pesaran (2007) Panel unit root test (CIPS)			
Specification without trend				Specification without trend			
Variable	lags	chi_sq	p-value	variable	lags	Zt-bar	p-value
wtem	0	1985.117	0.000	wtem	0	-32.823	0.000
wtem	1	1087.542	0.000	wtem	1	-19.487	0.000
wtem	2	531.137	0.000	wtem	2	-8.681	0.000
wtem	3	380.707	0.000	wtem	3	-5.952	0.000
wtem	4	280.994	0.281	wtem	4	-0.550	0.291
Specification with trend				Specification with trend			
Variable	lags	chi_sq	p-value	variable	lags	Zt-bar	p-value
wtem	0	2558.759	0.000	wtem	0	-36.870	0.000
wtem	1	1551.344	0.000	wtem	1	-24.024	0.000
wtem	2	763.353	0.000	wtem	2	-12.467	0.000
wtem	3	621.323	0.000	wtem	3	-8.719	0.000
wtem	4	463.470	0.000	wtem	4	-3.160	0.001

Null for MW and CIPS tests: series is I(1).
 MW test assumes cross-section independence.
 CIPS test assumes cross-section dependence is in form of a single unobserved common factor.

-multipurt- uses Scott Merryman's -xtfisher- and Piotr Lewandowski's -pescadf-.

FIGURE 7.7. Unit root tests suggest stationarity of population-weighted average temperatures within countries and time span surveyed

8. DATA APPENDIX

8.1. Air Conditioning Data. The AC imports data comes from WITS, the World Integrated Trade Solution data set (<http://wits.worldbank.org/wits/>). More specifically, it comes from 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-entered – 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.

8.2. Robustness Checks for International Panel Results. 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 (Figure 9.6). 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? How do we think about global time trends in the presence of year fixed effects? In the analysis above, we have presented the results without de-trending the temperature data, noting that, if anything, the coefficients relating temperature and income would likely be biased downward (lower magnitude) for hotter countries, given the mechanical correlation between higher income years and hotter years (both of which tend to occur later in the series). Allowing for country-specific temperature trends (in addition to global trends) leads to the results presented in Table 5. However, it is unclear whether this is the correct specification for the question at hand.

As DJO (2013) 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 (which we think is the case, although we haven't described it in those terms before) 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, actually, seem to be quite true, given the results we get back with the "detrended" data. The fact that the impact on hot countries goes away 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).

8.3. 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-meanded 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].