PRELIMINARY AND INCOMPLETE
NOT FOR ATTRIBUTION

Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US*

Olivier Deschênes
University of California, Santa Barbara

Michael Greenstone
MIT and NBER

April 2007

* We thank the late David Bradford for initiating a conversation that motivated this project. Our admiration for David's brilliance as an economist was only exceeded by our admiration for him as a human being. We are grateful for especially valuable criticisms from Maximillian Auffhammer, David Card, Hank Farber, Ted Gayer, Jon Guryan, Solomon Hsiang, David Lee, Enrico Moretti, Jesse Rothstein, Bernard Salanie, and Matthew White and Catherine Wolfram. We are also grateful for comments from seminar participants at British Columbia, Columbia, MIT (CEEPR), Pennsylvania, Princeton, Stanford, UC-Berkeley, UC-Irvine, UC-Los Angeles, UC-San Diego, UC-Santa Barbara (Bren), Washington, and Toronto. Elizabeth Greenwood provided truly outstanding research assistance. We also thank Ben Hansen for exemplary research assistance. We are very grateful to the British Atmospheric Data Centre for providing access to the Hadley Centre data. Greenstone acknowledges the University of California Energy Institute and the Center for Labor Economics at Berkeley for hospitality and support while working on this paper.

Introduction

The climate is a key ingredient in the earth's complex system that sustains human life and well being. There is a growing consensus that emissions of greenhouse gases due to human activity will alter the earth's climate, most notably by causing temperatures, precipitation levels, and weather variability to increase. According to the UN's Intergovernmental Panel on Climate Change (IPCC) Third Assessment Report (2001), climate change is likely to affect human health directly through changes in temperature and precipitation and indirectly through changes in the ranges of disease vectors (e.g., mosquitoes) and other channels. The design of optimal climate change mitigation policies requires credible estimates of the health and other benefits of reductions in greenhouse gases, however current evidence on the magnitudes of the direct and indirect impacts is considered insufficient for reliable conclusions (WHO 2003).[1]

Conceptual and statistical problems have undermined previous efforts to develop estimates of the health impacts of climate change. The conceptual problem is that the canonical economic models of health production predict that individual will respond to climate changes that threaten their health by purchasing goods that mitigate the health damages (Grossman 2000). In the extreme, it is possible that individuals would fully "self-protect" such that climate change wouldn't affect measured health outcomes. In this case, an analysis that solely focuses on health outcomes would incorrectly conclude that climate change had zero impact on welfare.

On the statistical side, there are at least three challenges. First, there is a complicated, dynamic relationship between temperature and mortality, which can cause the short-run relationship between temperature and mortality to differ substantially from the long run one (Huynen et al. 2001; Deschênes and Moretti 2005).[2] Second, individuals' locational choices---which determine exposure to a climate---are related to health and socioeconomic status, so this form of selection makes it difficult to uncover the

[1] See Tol (2002a and 2002b) for overall estimates of the costs of climate change, which are obtained by summing costs over multiple areas including agriculture, forestry, species/ecosystems, sea level rise, and human health. Deschenes and Greenstone (2006) provides evidence on the impacts on the US agricultures sector. Also, see Schlenker, Hanemann, and Fisher (2006).

[2] For example, Deschênes and Moretti (2005) document the importance of forward displacement or "harvesting" on hot days. Specifically, they find that hot days can lead to an immediate increase in mortality but a decline over the subsequent days such that 30 days after a hot day there is virtually no increase in mortality.

causal relationship between temperature and mortality. Third, the relationship between temperature and health is highly nonlinear and likely to vary across age groups and other demographic characteristics.

This paper develops measures of the welfare loss associated with the direct risks to health posed by climate change in the US that confront these conceptual and statistical challenges. Specifically, the paper reports on statistical models for demographic group by county annual mortality rates and for annual state-level energy consumption (perhaps primary form of protection against high temperatures) that model temperature semi-parametrically. The mortality models include county and state by year fixed effects, while the energy ones include state and Census-division by year fixed effects. Consequently, the temperature variables are identified from the unpredictable and presumably random year-to-year variation in temperature so concerns about omitted variables bias are likely to be unimportant.

We combine the estimated impacts of temperature and predicted changes in climate to develop estimates of the impact of climate change on mortality and energy consumption. The preferred mortality estimates indicate that climate change will lead to roughly 35,000 more deaths per year by the end of the century, which is roughly a 1.3% increase in the annual fatality rate. However, these estimated overall impacts are statistically indistinguishable from zero and the 95% confidence interval ranges from a decline of 23,000 fatalities to an increase of 93,000 per year. The estimates for some subgroups are more precise, as there are statistically significant increases in mortality rates predicted for male and female infants and some older age groups.

The energy results suggest that by the end of the century climate change will cause total US residential energy consumption to increase by 4 - 6 quadrillion (i.e., 10^{15}) British thermal units (Btus) or 25% - 35% of average annual consumption in the 1970-2003 period. This estimated increase is statistically significant and when valued at the average energy prices from the 1991-2000 period, it implies that there will be an additional $30 -$45 billion (2006$) per year of US energy consumption. It seems reasonable to assume that the mortality impacts would be larger without this compensatory response. More generally, the analysis suggests that a substantial portion of the adjustment to climate change will occur through adaptation or changes in consumption patterns, rather than increased mortality.

There are a few important caveats to these calculations and, more generally, the analysis. The

estimated impacts likely overstate the mortality and adaptation costs, because the analysis relies on inter-annual variation in weather and less expensive adaptations (e.g., migration) will be available in response to permanent climate change. On the other hand, the estimated welfare losses fail to include the impacts on numerous other determinants of welfare (e.g., morbidities) that may be affected by climate change so in this sense they are an underestimate. Additionally, the effort to project outcomes at the end of the century requires a number of strong assumptions, including that the climate change predictions are correct, relative prices (e.g., for energy and medical services) will remain constant, the same energy and medical technologies will prevail, and the demographics of the US population (e.g., age structure) and their geographical distribution will remain unchanged. These assumptions are strong, but their benefit is that they allow for a transparent analysis that is based on the available data, rather than unverifiable assumptions.

The analysis is conducted with the most detailed and comprehensive data available on mortality, energy consumption, weather, and climate change predictions for fine US geographic units. The mortality data come from the Compressed Mortality Files (CMF), energy data are from the Energy Information Administration (EIA), and the weather data are from the thousands of weather stations located throughout the US. We focus on end of century (i.e., 2070-2099) climate change predictions based on the Hadley Centre's 3rd Ocean-Atmosphere General Circulation Mode and the A1F1 scenario that predicts large increases in temperature. This scenario most closely approximates a "business-as-usual" or no carbon tax case.

Finally, it is notable that the paper's approach mitigates or solves the conceptual and statistical problems that have plagued previous research. First, the availability of data on energy consumption means that we can measure the impact on mortality and self-protection expenditures. Second, we demonstrate that the focus on annual mortality, rather than daily, mitigates concerns about harvesting/ forward displacement and delayed impacts. Third, the county fixed effects adjust for any differences in unobserved health across locations due to sorting. Fourth, we model temperature with 20 separate variables that measure the number of days in a year that the daily mean temperature falls in 5 degree Fahrenheit ranges or bins. Consequently, we don't have to rely on functional form assumptions to infer

the impacts of the hottest and coldest days on mortality. Fifth, we estimate separate models for 16 demographic groups (8 age categories and male/female), which allows for substantial heterogeneity in the impacts of temperature.

The paper proceeds as follows. Section I briefly reviews the patho-physiological and statistical evidence on the relationship between weather and mortality. Section II provides the conceptual framework for our approach. Section III describes the data sources and reports summary statistics. Section IV presents the econometric approach and Section V describes the results. Section VI assesses the magnitude of our estimates of the effect of climate change and discusses a number of important caveats to the analysis. Section VII concludes the paper.

I. Background on the Relationship between Weather and Mortality

Individuals' heat regulation systems enable them to cope with high and low temperatures. Specifically, exposure to both high and low temperatures generally triggers an increase in the heart rate in order to increase blood flow from the body to the skin, leading to the common thermoregulatory responses of sweating in hot temperatures and shivering in cold temperatures. These responses allow individuals to pursue physical and mental activities without endangering their health within certain ranges of temperature. Temperatures outside of these ranges pose dangers to human health and can result in premature mortality. This section provides a brief review of the mechanisms and the challenges for estimation.

Hot Days. An extensive literature documents a relationship between extreme temperatures (usually during heat waves) and mortality (e.g., Klineberg 2003; Huynen 2001; Rooney et al. 1998). These excess deaths are generally concentrated among causes related to cardiovascular, respiratory, and cerebrovascular diseases. The primary mechanism in these periods is the additional stress imposed on the cardiovascular and respiratory systems by the need for body temperature regulation. In the context of specific indicators of body operations, elevated temperatures are associated with increases in blood viscosity and blood cholesterol levels. It is not surprising that numerous studies have shown that access to air conditioning greatly reduces mortality on hot days (CITATIONS TO COME).

An important feature of the relationship between heat and mortality is that the number of deaths immediately caused by a period of very high temperatures is at least partially compensated for by a reduction in the number of deaths in the period immediately subsequent to the hot day or days (Basu and Samet 2001; Deschênes and Moretti 2005). This pattern is called forward displacement or "harvesting," and appears to occur because heat affects individuals that were already very sick and would have died in the near term. Since underlying health varies with age, these near-term displacements are more prevalent among the elderly.

Cold Days. Cold days are also a risk factor for mortality. Exposure to very cold temperatures causes cardiovascular stress due to changes in blood pressure, vasoconstriction, and an increase in blood viscosity (which can lead to clots), as well as levels of red blood cell counts, plasma cholesterol, and plasma fibrinogen (Huynen et al. 2001). Further, susceptibility to pulmonary infections may increase because breathing cold air can lead to bronchoconstriction. Interestingly, there is some evidence that cold weather can affect driving fatalities. In particular, Eisenberg and Warner (2005) found that on days with snow accumulation there is a decline in fatal motor vehicle accidents relative to dry days, although there is an increase in nonfatal ones.

Deschênes and Moretti (2005) provide the most comprehensive evidence on the impacts of cold days on mortality. They find "evidence of a large and statistically significant effect on mortality within a month of the cold wave. This effect appears to be larger than the immediate effect, possibly because it takes time for health conditions associated with extreme cold to manifest themselves and to spread" (Deschênes and Moretti 2005). Thus, in the case of cold weather, it may be that there are delayed impacts and that the full effect of a cold day takes a few weeks to manifest itself. Further, they find that the impact is most pronounced among the young and elderly and concentrated among cardiovascular and respiratory diseases.

Implications. The challenge for this study and any study focused on substantive changes in life expectancy is to develop estimates of the impact of temperature on mortality that are based on the full long-run impact on life expectancy. In the case of hot days, the previous literature suggests that this task requires purging the temperature effects of the influence of harvesting or forward displacement. In the

case of cold days, the mortality impact may accumulate over time. In both cases, the key point is that the full impact of a given day's temperature may take numerous days to manifest fully.

Our review of the literature suggests that the full mortality impacts of cold and hot days are evident within 30 days (Huynen et al. 2001; Deschênes and Moretti 2005). The below econometrics section outlines a method that allows the mortality impacts of temperature to manifest themselves over long periods of time. Further, the immediate and longer run effects of hot and cold days are likely to vary across the populations with larger impacts among relatively unhealthy subpopulations. One important determinant of healthiness is age, with the old and young being especially sensitive to environmental insults. Consequently, we conduct separate analyses for 16 demographic groups defined by the interaction of gender and 8 age categories.

II. Conceptual Framework

This paper's goal is to estimate part of the health related welfare impact of climate change. To understand the welfare effects of an exogenous increase in temperature, it is instructive to consider a simple 1-period model of utility maximizing individuals that formalizes the decision to invest in human health (see Grossman 2000). For this section's expositional purposes, we assume that climate change leads to an increase in temperatures in the summer only when higher temperatures are harmful for health.

We assume individuals consume a jointly aggregated consumption good, x_C. Their other consumption good is their mortality risk, which leads to a utility function of

(1) $U = U[x_C, s]$,

where s is the survival rate. The production function for survival is expressed as:

(2) $s = s(x_H, T)$,

where T represents temperature. Consequently, survival is a function of x_H, which is a private good that affects the probability of survival, and ambient temperature. For example, energy consumption is an example of x_H since energy is used to run air conditioners, which affect survival on hot days. We define x_H such that $\partial s/\partial x_H > 0$. Since this section focuses on summers where higher temperatures are dangerous, we assume that $\partial s/\partial T < 0$.

The individual faces a budget constraint of the form:

(3) $I - x_C - px_H = 0$,

where I is exogenous earnings or income. We assume the prices of x_C and x_H are 1 and p, respectively.

The individual's problem is to maximize (1) through her choices of x_C and x_H, subject to equations (2) and (3). The solution of this maximization problem reveals that the input demand equations for x_C and x_H are functions of prices, income, and temperature. In equilibrium, the ratio of the marginal utilities of consumption of the two must be equal to the ratio of the prices: $[(\partial U/\partial s)\cdot(\partial s/\partial x_H)]/[\partial U/\partial x_C] = p$.

Now, consider an exogenous increase in temperature due to climate change, which reduces the survival probability. If individuals are not allowed to adjust their consumption bundles, then $[(\partial U/\partial s)\cdot(\partial s/\partial x_H)]/[\partial U/\partial x_C] > p$. In this case, their expenditures on health production are suboptimally low and they are literally dying to increase consumption of x_H.

To move back into equilibrium after the temperature increase, consumers will increase their consumption of x_H at the expense of consumption of x_C. Since the effective price of survival has increased, the chosen survival rate, s, is likely to decline. In fact, it is possible that there would be a large change in the consumption of x_H and little change in s. However, the key point for this paper's purposes is that the full welfare effect of the exogenous change in temperature is reflected in changes in the survival rate and the consumption of x_H.[3]

The subsequent empirical work will develop estimates of the impact of climate change on mortality rates and on energy consumption. We will monetize the changes in mortality and energy consumption to develop a partial measure of the welfare loss associated with climate change. This will only be a partial measure of the welfare loss, because climate change may affect other health outcomes

[3] A richer model would allow individuals to choose the temperatures they face through their location choice (see, e.g., Rosen 1974). The advantage of such a model is that in principle, it is possible to capture the full welfare effects of a change in temperature through the land market since land is a fixed factor. If there is imperfect information about the change in temperatures and its health effects, then this hedonic method will not produce reliable estimates of the welfare effects of a change in temperature. An advantage of this paper's health production function approach is that it does not rely on perfect information assumptions. Its disadvantages are that for data reasons it may not be possible to measure all the potential health effects (e.g., changes in morbidity rates) and it cannot capture any non-health benefits (e.g., individuals may prefer higher temperatures for amenity reasons).

(e.g., morbidity rates). Further, although energy consumption likely captures a substantial component of health preserving expenditures, climate change may induce other forms of adaptation (e.g., substituting indoor exercise for outdoor exercise or changing the time of day when one is outside). These other outcomes are unobservable in the data files that we have compiled for this paper, so our welfare estimates will be incomplete.

Finally, this one-period model misses an issue, which may be especially relevant in light of our empirical strategy that relies on inter-annual fluctuations in weather to learn about the welfare consequences of permanent climate change. Specifically, individuals cannot engage in the full set of adaptations in response to a single year's weather realization. It is easy to turn the thermostat down and use more air conditioning on hot days, and it is even possible to purchase an air conditioner in response to a single year's heat wave. A number of adaptations, however, cannot be undertaken in response to a single year's weather realization. For example, permanent climate change is likely to lead to some migration (presumably to the North), which will be missed with our approach.

The limited set of adaptations available in the short-run means that our approach is likely to overstate the mortality and energy consumption-related costs of climate change. This is because the longer-run adaptations will only be undertaken if they lead to smaller reductions in consumption of x_C, than the reductions that occur when only the short-run adaptations are available. So, this study only examines a subset of the outcomes likely to be affected by climate change and in this respect it underestimates the costs, but it is likely to provide an upper bound on the costs among the outcomes it measures (i.e., mortality and energy consumption).

III. Data Sources and Summary Statistics

To implement the analysis, we collected the most detailed and comprehensive data available on mortality, energy consumption, weather, and predicted climate change. This section describes these data and reports some summary statistics.

A. Data Sources

Mortality and Population Data. The mortality data is taken from the Compressed Mortality Files (CMF) compiled by the National Center for Health Statistics. The CMF contains the universe of the 72.3 million deaths in the US from 1968 to 2002. Importantly, the CMF reports death counts by race, sex, age group, county of residence, cause of death, and year of death. In addition, the CMF files also contain population totals for each cell, which we use to calculate all-cause and cause-specific mortality rates. Our sample consists of all deaths occurring in the continental 48 states plus the District of Columbia.

Energy Data. The energy consumption data comes directly from the EIA's State Energy Data System. These data provide state-level information about energy price, expenditures, and consumption from 1968 to 2002. The data is disaggregated by energy source and end use sector. All energy data is given in British Thermal Units, or BTU.

We used the database to create an annual state-level panel data file for total energy consumption by the residential sector, which is defined as "living quarters for private households." The database also reports on energy consumption by the commercial, industrial, and transportation sectors. These sectors are not a focus of the analysis and they don't map well into the health production function model outlined in Section II. Further, factors besides temperature are likely to be the primary determinant of consumption in these sectors.

The measure of total residential energy consumption is comprised of two pieces-- "primary" consumption, which is the actual energy consumed by households, and "electrical system energy losses." The latter accounts for about 2/3 of total residential energy consumption; it is largely due to losses in the conversion of heat energy into mechanical energy to turn electric generators, but transmission and distribution and the operation of plants also account for part of the loss. In the 1968-2002 period, total residential energy consumption increased from 7.3 quadrillion (quads) British thermal units to 21.2 quads and the mean over the entire period was 16.6 quads.

Weather Data. The weather data are drawn from the National Climatic Data Center (NCDC) Summary of the Day Data (File TD-3200). The key variables for our analysis are the daily maximum and

minimum temperature, as well as total daily precipitation.[4] These data are taken from the 7,380 stations that met our sample selection rule for at least one year during the 1968-2002 period. The acceptable station-level data is then aggregated at the county level by taking the simple average of the measurements from all stations within a county.

To ensure the accuracy of the weather readings, we developed a weather station selection rule. Specifically, we dropped all weather stations at elevations above 7,000 feet since they were unlikely to reflect the weather experienced by the majority of the population within a county. Among the remaining stations, we considered a year's readings valid if the station operated at least 363 days. The average annual number of stations with valid data in this period was 3,879. The county by years with acceptable weather data accounted for 53.4 of the 72.3 million deaths in the US from 1968 to 2002.

Climate Change Prediction Data. Climate predictions are taken from the Hadley Centre's 3rd Coupled Ocean-Atmosphere General Circulation Model, which we refer to as Hadley 3 (T. C. Johns et al. 1997, Pope et al. 2000). This is the most complex and recent model in use by the Hadley Centre and more broadly is considered a state of the art global climate model. This is because Hadley 3 is a coupled atmospheric-ocean general circulation model. This type of model is considered the most complex because it considers the interplay of several earth systems and is therefore considered the most appropriate for climate predictions.

Predictions of climate change in the Hadley 3 model are available for several emission scenarios, corresponding to 'storylines' describing the way in which the world (population, economies, etc.) will develop over the next 100 years. We consider the two extreme scenarios: A1F1, which predicts the largest increase in global mean temperature and B1, which predicts the smallest.

Throughout the paper, we focus on the Hadley 3 A1F1 scenario because we believe it most accurately represents the "business-as-usual" scenario against which carbon taxes or other interventions should be judged. Both the A1F1 and B1 scenarios assume that there will be rapid economic growth

[4] Other aspects of daily weather such as humidity and wind speed could influence mortality, both individually and in conjunction with temperature. Importantly for our purposes, there is little evidence that using wind chill factors (a non-linear combination of temperature and wind speed) perform better than simple temperature levels in explaining daily mortality rates (Kunst et al. 1994).

during the current century, especially in the developing world which leads to economic convergence. The two scenarios diverge on the issue of fossil fuel usage. The A1F1 scenario assumes that fossil fuels will continue to be the primary source of energy, while the B1 scenario assumes that energy comes mostly from alternative (i.e., non-fossil fuel) sources (IPCC 1996). Given the abundant supply of inexpensive coal and other fossil fuels, a switch to alternative sources is unlikely without carbon taxes or the equivalent. Thus, our view is that A1F1 is the proper benchmark scenario because it doesn't reflect a policy intervention and we emphasize the results based on this scenario.

We use the results of the application of these scenarios to the Hadley 3 model to develop regional estimates of the impact of climate change within the US. The Hadley 3 predictions are available for grid points spaced at 2.5° (latitude) x 3.75° (longitude). There are 153 grid points in total, but we use the 89 located on land to develop the regional estimates. Six states do not have a grid point, so we developed daily Census division-level predictions for the 9 US Census divisions using the land grid points. More details on the temperatures predictions are available in the data appendix.

B. Summary Statistics

Mortality Statistics. Table 1 reports the average annual mortality rates per 100,000 by age group and gender using the 1968-2002 CMF data. It is reported separately for all-causes of death and for deaths due to cardiovascular disease, neoplasms, respiratory disease, and motor-vehicle accidents.[5] These four categories account for roughly 75% of all fatalities, though the relative importance of each cause varies by sex and age. Importantly, these entries are consistent with the average annual mortality rates reported by the National Center for Health Statistics.

The all cause and all age mortality rates for women and men are 804.4 and 939.2, respectively, but there is tremendous heterogeneity in mortality rates across age and gender groups. For all-cause mortality, the female and male infant mortality rates are 1,031.1 and 1,292.1. After the first year of life,

[5]In terms of ICD-9 Codes, the causes of deaths are defined as follows: Neoplasms = 160-250, Cardiovascular Diseases = 320-490, Respiratory Diseases = 500-580, Motor Vehicle Accidents = XXXX. We focus on motor-vehicle accidents, rather than all external causes since climate change could change the number of days with unsafe driving conditions.

mortality rates don't approach this level again until the 55-64 category. The annual mortality rate starts to increase dramatically at older ages, and in the 75-99 age category it is 8.0% for women and 9.4% for men. The higher annual fatality rates for men at all ages are striking and explain their shorter life expectancy.

As is well-known, mortality due to cardiovascular disease is the single most important cause of death in the population as a whole. The entries indicate that cardiovascular disease is responsible for 48.4% and 43.6% of overall female and male, mortality. It is noteworthy that the importance of the different causes of death varies dramatically across age categories. For example, motor vehicle accidents account for 22.1% (23.8%) of all mortality for women (men) in the 15-24 age group. In contrast, cardiovascular disease accounts for 59.6% (53.7%) of all mortality for women (men) in the 75-99 category, while motor vehicle accidents are a negligible fraction. More generally, for the population aged 55 and above---where mortality rates are highest---cardiovascular disease and neoplasms are the two primary causes of mortality.

Weather and Climate Change Statistics. We take advantage of the richness of daily weather data and climate change predictions data by using the information on daily minimum and maximum temperatures. Specifically, we calculate the daily mean temperatures at each weather station as the mean of the each day's minimum and maximum temperature. The county-wide mean is then calculated as the unweighted average across all stations within a county. The climate change predictions are calculated analogously, except that we take the average of the daily predicted mean temperature across the grid points within the relevant geographic units.

The "Actual" column of Table 2 reports the weighted average of the daily mean temperatures across counties from 1968-2002, where the weight is the population between ages 0 and 99 in the relevant year. The average daily mean is 56.6° F.[6] The entries for the four Census regions confirm that the South is the hottest part of the country and the Midwest and Northeast are the coldest ones.[7] Since people are

[6] The average daily mean and all other entries in the table (as well as in the remainder of the paper) are calculated across counties that meet the weather station sample selection rule described above.

[7] The states in each of the Census regions are: Northeast-- Connecticut, Maine, Massachusetts, New Hampshire, Vermont, Rhode Island, New Jersey, New York, and Pennsylvania; Midwest-- Illinois, Indiana, Michigan, Ohio, Wisconsin, Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, and South Dakota; South-- Delaware,

more familiar with daily highs and lows from newscasts, the table also documents the average daily maximum and minimums.[8] The average daily spread in temperatures is 21.2° F, indicating that hottest and coldest hours can differ substantially from the mean. It is important to emphasize that the weighted means, as well as the subsequent predictions, depend on the distribution of the population across the US so systematic migration (e.g., from South to North) would change these numbers even without any change in the underlying climate.

Figure 1 depicts the variation in the measures of temperature across 20 temperature bins in the 1968-2002 period. The leftmost bin measures the number of days with a mean temperature less than 0° F and the rightmost bin is the number of days where the mean exceeds 90° F. The intervening 18 bins are all 5° F wide. These 20 bins are used throughout the remainder of the paper as they form the basis for our semi-parametric modeling of temperature in equations for mortality rates and energy consumption.

The figure depicts the mean number of days that the typical person experiences in each bin; this is calculated as the weighted average across county by year realizations, where the county by year's population is the weight. The average number of days in the modal bin of 70° - 75° F is 38.2. The mean number of days at the endpoints is 0.8 for the less than 0° F bin and 1.6 for the greater than 90° F bin.

The remaining columns of Table 2 report on the predicted changes in temperature from the B1 and A1F1 scenarios and the Hadley 3 model for the 2070-2099 period.[9] Each set of predictions is based on a single run of the Hadley 3 model. Recall, the B1 scenario predicts the smallest greenhouse gas accumulation and temperature increase, whereas A1F1 predicts the largest increases and as the business as usual scenario is our focus.

The B1 scenario predicts a change in mean temperature of just 0.5° F. Interestingly, there is substantial heterogeneity with mean temperatures expected to increase by 5.0° F in the Midwest and to

---

District of Columbia, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, West Virginia, Alabama, Kentucky, Mississippi, Tennessee, Arkansas, Louisiana, Oklahoma, and Texas; and West-- Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, Wyoming, Alaska, California, Hawaii, Oregon, and Washington.

[8] For counties with multiple weather stations, the daily maximum and minimum are calculated as the average across the maximums and minimums, respectively, from each station.

[9] For comparability, we follow much of the previous literature on climate change and focus on the temperatures predicted to prevail at the end of the century.

decline by 2.7° F in the West. The A1F1 scenario predicts a gain in mean temperature of 6.5° F. The increases in the Midwest and South exceed 9° F, while there is virtually no predicted change in the West.

Figure 2 provides an opportunity to understand how the full distributions of mean temperatures are expected to change under the A1F1 scenario. One's eye is naturally drawn to the last two bins. The typical person will experience 18.9 additional days per year where the mean daily temperature is between 85° F and 90° F. Even more amazing, the mean daily temperature is predicted to exceed 90° F 43.8 extra days per year.[10] To put this in perspective, given the current distribution of the population across the country, the average person experiences just 1.6 days per year where the mean exceeds 90° F.

An examination of the rest of the figure highlights that the increase in these very hot days is not being drawn from the entire year. For example, the number of days where the maximum is expected to be between 55° F and 80° F declines by 55.5 days. Interestingly, the mean number of days where the minimum temperature will be below 30° F is predicted to fall by just 3.8 days, further underscoring that the temperature increases are not spread evenly throughout the year.

IV. Econometric Strategy

This section describes the econometric models that we employ to learn about the impact of temperature on mortality rates and residential energy consumption.

A. Mortality Rates.

We fit the following equations for county-level mortality rates of various demographic groups:

$$(4)\quad Y_{ctd} = \sum_j \theta_{dj}^{TMEAN} TMEAN_{ctj} + \sum_l \delta_{dl}^{PREC} PREC_{ctl} + X_{ct}\beta_d + \alpha_{cd} + \gamma_{std} + \varepsilon_{ctd}.$$

Y_{ctd} is the mortality rate for demographic group d in county c in year t. In the subsequent analysis, we use 16 separate demographic groups, which are defined by the interaction of 8 age categories (0-1, 1-14, 15-24, 25-44, 45-54, 55-64, 65-74, and 75+) and gender. X_{ct} is a vector of observable time varying determinants of fatalities measured at the county level. The last term in equation (4) is the stochastic

[10] At the risk of insulting the reader, we want to underscores that a mean daily temperature of 90° F is very hot. For example, a day with a high of 100° F would need a minimum temperature greater than 80° F to qualify.

error term, ε_{ctd} .

The variables of interest are the measures of temperature and precipitation, and we have tried to model these variables with as few parametric assumptions as possible while still being able to make precise inferences. Specifically, they are constructed to capture the full distribution of annual fluctuations in weather. The variables $TMEAN_{ctj}$ denote the number of days in county c and year t where the daily mean temperature is in one of the 20 bins used in Figures 1 and 2. Thus, the only functional form restriction is that the impact of minimum temperature is constant within 5F degree intervals. This degree of flexibility and freedom from parametric assumptions is only feasible because we are using 35 years of data from the entire US. Since extreme high and low temperatures drive most of the health impacts of temperature, we tried to balance the dual and conflicting goals of allowing the impact of temperature to vary at the extremes and estimating the impacts precisely enough so that they have empirical content. The variables $PREC_{ctl}$ are simple indicator variables denoting annual precipitation equal to "l" in county c in year. These intervals correspond to 2 inches bins.

The equation includes a full set of county fixed effects, α_{cd} . The appeal of including the county fixed effects is that they absorb all unobserved county-specific time invariant determinants of the mortality rate for each demographic group. So, for example, differences in permanent hospital quality or the overall healthiness of the local age-specific population won't confound the weather variables. The equation also includes state by year indicators, γ_{std} , that are allowed to vary across the demographic groups. These fixed effects control for time-varying differences in the dependent variable that are common within a demographic group in a state (e.g., changes in state Medicare policies).

The validity of any estimate of the impact of climate change based on equation (4) rests crucially on the assumption that its estimation will produce unbiased estimates of the θ_{dj}^{TMEAN} and δ_{dl}^{PREC} vectors. The consistency of the components of each θ_{dj} requires that after adjustment for the other covariates the unobserved determinants of mortality do not covary with the weather variables. In the case of the mean temperatures, this can be expressed formally as $E[TMEAN_{ctj}\ \varepsilon_{ctd} \mid X_{ct}, \alpha_{cd}, \gamma_{std}] = 0$. By conditioning on the county and state by year fixed effects, the θ_{dj} 's are identified from county-specific deviations in weather about the county averages after controlling for shocks common to all counties in a state. Due to

the unpredictability of weather fluctuations, it seems reasonable to presume that this variation is orthogonal to unobserved determinants of mortality rates. The point is that there is reason to believe that the identification assumption is valid.

A primary motivation for this paper's approach is that it may offer an opportunity to identify weather-induced changes in the fatality rate that represent the full impact on the underlying population's life expectancy. Our review of the literature suggests that the full effect of particularly hot and cold days is evident within approximately 30 days (Huynen et al. 2001; Deschênes and Moretti 2005). Consequently, the results from the estimation of equation (4) that uses the distribution of the year's daily temperatures should largely be free of concerns about forward displacement and delayed impacts. This is because a given day's temperature is allowed to impact fatalities for a minimum of 30 days for fatalities that occur from February through December. An appealing feature of this set-up is that the θ_{dj}^{TMEAN} coefficients can be interpreted as reflecting the full long-run impact of a day with a mean temperature in that range.

The obvious limitation is that the weather in the prior December (and perhaps earlier parts of the year if the time frame for harvesting and delayed impacts is longer than 30 days) may affect current year's mortality. To assess the importance of this possibility, we also estimate models that include a full set of temperature variables for the current year (as in equation (4)) and the prior year. As we demonstrate below, our approach appears to purge the estimates of fatalities of people with relatively short life expectancies.[11]

There are two further issues about equation (4) that bear noting. First, it is likely that the error terms are correlated within county by demographic groups over time. Consequently, the paper reports standard errors that allow for heteroskedasticiy of an unspecified form and are clustered at the county level.

[11] A daily version of equation (4) would be very demanding of the data. In particular, there would be a tension between our flexibility in modeling temperature and the number of previous days of temperature to include in the model. Equation (4) models temperature with 20 variables, so a model that includes 30 previous days would use 600 variables for temperature, while one with 365 days would require 7300 temperature variables. Further, daily mortality data for the entire US is only available from 1972-1988 and there may not be sufficient variation in temperature within this relatively short period of time to precisely identify some of the very high and very low temperature categories.

Second, it may be appropriate to weight equation (4). Since the dependent variable is demographic group-specific mortality rates, we think there are two complementary reasons to weight by the square root of demographic group's population (i.e., the denominator). First, the estimates of mortality rates with large populations will be more precise than the estimates from counties with small populations, and this weight corrects for the heteroskedasticity associated with the differences in precision. Second, the results can then be interpreted as revealing the impact on the average person, rather than on the average county.

Residential Energy Consumption. We fit the following equation for state-level residential energy consumption:

$$(5)\ \ln(C_{st}) = \sum_j \theta_j^{TMEAN} TMEAN_j + \sum_l \delta_l^{PREC} PREC_l + X_{st}\beta + \alpha_s + \gamma_{dt} + \varepsilon_{st}.$$

C_{st} is residential energy consumption in state s in year t and d indexes Census Division. The modeling of temperature and precipitation is identical to the approach in equation (4). The only difference is that these variables are measured at the state by year level—they are calculated as the weighted average of the county-level versions of the variables, where the weight is the county's population in the relevant year. The equation also includes state fixed effects (α_s) and census division by year fixed effects (γ_{dt}) and a stochastic error term, ε_{st}.

The challenge for the successful estimation of this equation is that there has been a dramatic shift in the population from the North to the South over the last 35 years. If the population shifts were equal within Census divisions, this wouldn't pose a problem for estimation but this has not been the case. For example, Arizona's population has increased by 223% between 1968 and 2002, compared to just 124% for the other states in its Census Division and due to its high temperatures it plays a disproportionate role in the identification of the θ_j's associated with the highest temperature bins.[12] The point is that unless we correctly adjust for these population shifts, the estimated θ_j's may confound the impact of higher

[12] For example, we estimated state by year regressions for the number of days where the mean temperature was in the > 90° F bin that adjusted for state fixed effects and census division by year fixed effects. The mean of the annual sum of the absolute value of the residuals for Arizona is 3.6 but only 0.6 in the other states in its Census Division. The other states in Arizona's Census Division are Colorado, Idaho, New Mexico, Montana, Utah, Nevada, and Wyoming.

temperatures with the population shifts.

As a potential solution to this issue, the vector X_{st} includes the ln of population and gross domestic product by state as covariates. The latter is included since energy consumption is also a function of income. The below analysis demonstrates the importance of controlling for these covariates and in mitigating confounding associated with the population shifts out of the Rust Belt and to warmer states.

Finally, we will also report the results from versions of equation (5) that model temperature with heating and cooling degree days. We follow the consensus approach and use a base of 65° F to calculate both variables.[13] Specifically, on a given day, the number of cooling degree days equals the day's mean temperature (i.e., the average of the minimum and maximum) minus 65° F for days where the mean is above 65° F and zero for days when the mean is below 65° F. Analogously, a day's heating degree days is equal to 65° F minus its mean for days where the mean is below 65° F and zero otherwise. So, a day with a mean temperature of 72° F would contribute 7 cooling degree days and 0 heating degree days, while a day with a mean of 51° F would contribute 0 cooling degree days and 14 heating degree days.

To implement this alternative method for modeling a year's temperature, we sum the number of heating and cooling degree days separately over the year. We then include the number of heating and cooling degree days and their squares in equation (5) instead of the $TMEAN_{ctj}$ variables.

V. Results

This section is divided into three subsections. The first explores the extent of variation in the temperature variables in the context of the rich statistical models that we employ. The second provides estimates of the impact of predicted climate change on overall mortality. It also implements a series of specification tests and assesses whether these effects are concentrated among particular causes of death.

[13] Electrical, natural gas, power, heating, and air conditioning industries utilize heating and cooling degree calculations to predict demand (http://www.fedstats.gov/qf/meta/long_242362.htm). Further, the National Oceanic and Atmospheric Administration recommends using a base of 65° F for both heating and cooling degree days (http://www.cpc.ncep.noaa.gov/products/analysis_monitoring/cdus/degree_days/ddayexp.shtml). Further, an examination of the figures suggests in Engle et al.'s seminal paper on relationship between temperature and electricity sales suggests that 65° F is a reasonable base for both cooling and heating degree days (Engle et al. 1986).

The third examines the impact of predicted climate change on residential energy consumption.

A. How Much Variation is there in Temperature?

As we discussed above, our preferred specifications model temperature with 20 separate variables. For this method to be successful, it is important that there is substantial inter-annual variation in county temperature after adjustment for these county and state by year fixed effects in the mortality equations. If this is the case, the predicted health impacts of climate change will be identified from the data, rather than by extrapolation due to functional form assumptions.

Figure 3 provides an opportunity to assess the extent of inter-annual variation in temperature. For each daily mean temperature bin, we create a data file where the observations are from all county by year observations with valid weather data between 1968 and 2002. We then regress the annual realization of the number of days that the relevant county had a daily mean in the temperature bin against state-by-year and county fixed effects. For each county by year, we sum the absolute value of the residuals. The figure reports the mean of this number across all county by year observations. The resulting figures can be interpreted as the average number of days in a county by year that are available to identify the parameter associated with that temperature bin after adjustment for the fixed effects.

An inspection of the figure demonstrates that there is substantial variation in temperatures, so it should be possible to obtain relatively precise estimates of the impacts of most of the temperature bins. Notably, due to the large data file, there are still many days available to estimate the impact of even the extreme bins. For example, the mean of the absolute value of the residuals for the bin for the > 90° F bin is 0.7 days. Although this may seem small, the size of our data file helps greatly. Since there are XXXX county by year observations, this means that there are roughly XXXX county by days to help identify the impact of a day in this bin. The analogous figure for the 85° - 90° F bin is XXXX days.

B. Estimates of the Impact of Climate Change on Mortality

All Cause Mortality Results. Figure 4 provides an opportunity to better understand the paper's approach and understand the basis of the results for male infants. It plots the estimated θ_j's from the

estimation of equation (4) for male infants. In this version of the equation, we dropped the $TMEAN_j$ variable associated with the 65° - 70° F bin so each θ_j reports the estimated impact of an additional day in bin j on the infant mortality rate (i.e., deaths per 100,000) relative to the mortality rate associated with a day where the temperature is between 65° - 70° F. The figure also plots the estimated θ_j's plus and minus one standard error of the estimates so that the precision of each of these estimates is evident.

The most striking feature of this graph is that the response function is generally flat so temperature has little influence on male infant mortality rates, except at the highest and coldest temperatures. Recall, the Hadley 3 A1F1 results predict that the changes in the distribution of temperature will be concentrated among days where the mean temperatures exceeds 50° F, so the estimated θ_j's in this range are most relevant for this paper's exercise. If the estimates are taken literally, it is evident that the shift of days into the last bin will lead to an increase in infant mortality. For example, the results suggest that the shift of a day from the 70° - 75° F bin (estimated θ = -0.78) to the > 90° F bin (estimated θ = 0.92) would lead to 1.7 more infant deaths per 100,000 births.

It is also important to highlight that the estimated θ_j's have associated sampling errors. Among the most relevant θ_j's, the largest standard error is in the highest bin which is due to the relatively small days with a mean temperature exceeding 90° F. The imprecision of the estimated impact of this bin poses a challenge for making precise inferences about the impact of the predicted changes in temperature on mortality rates. The estimated θ_j's at the lowest temperatures are also imprecise, but they play an inconsequential role in this exercise.

We now turn to Table 3, which summarizes the results from the estimation of separate versions of equation (4) for the 16 gender by age groups using the Hadley 3 A1F1 scenario. These versions include all twenty $TMEAN_j$ variables. Estimates for females and males are reported in the left and right panels, respectively. Columns (1a) and (2a) report the predicted change in annual mortality for each demographic group and its estimated standard error. For a given county and demographic groups, these impacts are calculated as follow:

$$(6)\quad M_{cd} = POP_{cd} \times \sum_j \hat{\theta}_{dj}^{TMEAN} \Delta TMEAN_{cj}$$

That is, we multiply the predicted change in the number of days in each temperature cell from the Hadley 3 A1F1 predictions ($\Delta TMEAN_{cj}$) by the corresponding demographic-group specific impact on mortality ($\hat{\theta}_{dj}^{TMEAN}$) and then sum these products. This sum is then multiplied by the average population for that demographic in that county (POP_{cd}) over the sample period. Finally, the impacts for a given demographic group are summed over all counties and this is the national demographic group-specific estimate of the change in annual mortality. It is straightforward to calculate the standard error, since the estimated mortality change is a linear function of the estimated parameters.

Columns (1b) and (2b) report the change in the annual mortality rate. This is calculated as the ratio of the change in the demographic group's mortality rate due to predicted climate change to the group's overall mortality rate. Since fatalities at relatively young ages are likely associated with larger changes in life expectancies, columns (1c) and (2c) report the change in life years due to predicted climate change for each age category. This entry is the product of the predicted increase in annual fatalities and the residual life estimate for each age group (evaluated in the middle of the age range) and sex, taken from the 1980 Vital Statistics.[14] Negative values correspond to loss of life-years, while positive entries correspond to gains in life-years. We note that this calculation may overstate the change in life years, because affected individuals are likely to have shorter life expectancies than the average person. Nevertheless, it provides a way to capture that fatalities at young ages may have greater losses of life expectancy than those at older ages.

The entries in columns (1d) and (2d) report p-values from F-tests of the hypothesis that the twenty estimated θ_j's are equal. This test is not directly informative about the mortality impacts of predicted climate change, but it provides a summary of the impact of temperature on mortality in the US. A failure to reject the null is consistent with the view that in the US individuals are able to easily adapt to changes in temperature that could pose risks to mortality.

We begin by returning to infant mortality, which is reported in the first row. These entries indicate that predicted climate change will increase the number of female and male infant deaths by

[14] Starting with infants and progressing towards the oldest age category, the residual life estimates for females are: 78.1, 72.1, 59.4, 45.8, 30.9, 22.4, 14.8, and 6.3. The corresponding estimates for males are: 70.7, 64.8, 52.4, 39.5, 25.2, 17.5, 11.3, and 5.0.

roughly 1,000 and 1,800 per year. The female estimate borders on statistical significance at conventional levels, while the male estimate is substantially more precise. These estimates are equivalent to increases of 5.5% (female) and 7.8% (male) infant mortality rates. The life years calculation suggests that these extra fatalities would lead to a loss of about 200,000 life years of life expectancy every year. This finding of higher temperatures leading to increased rates of infant mortality is consistent with the medical evidence that infants' thermoregulatory systems are not fully developed.

In the remainder of the table, there is evidence of mortality impacts for some sub-groups. The most substantial impacts are concentrated among 75-99 females. The entries suggest that there would be an additional 11,500 fatalities per year in this demographic group and their annual mortality rate would increase by roughly 2.0%. Due to their age, the total loss of life years is comparable to the loss for infants even though the increase in fatalities is 11 times larger. There is also evidence of an increased mortality rates for 1-14 year olds and for men in the 45-54 and 55-64 age categories. However, there is little evidence of an increase in the mortality rate for many of the demographic groups; for example, the null of zero increase in fatalities cannot be rejected at even the 10% level for 9 of the 16 demographic groups. Overall, these differences in the results across age, gender, and age by gender categories underscore the importance of estimating these models in such a disaggregated manner.

The bottom row of Table 3 reports the aggregate impacts, which are simply the sum of the impacts for each demographic group. For both females and males, annual mortality is predicted to increase by approximately 17,500 deaths per year. This excess mortality corresponds to increases in the annual mortality rate of 1.3% for both genders. The predicted loss of life years is about 290,000 for women and 490,000 for men. The marked difference between males and females reflects the differences in the distribution of the mortality impacts across age categories, especially the heavy concentration in the 75-99 age category for women.

To understand the source of these aggregate estimates, it is instructive to examine the regression coefficients (i.e., $\hat{\theta}_{dj}^{TMEAN}$) that drive the overall estimates. Figures 5A and 5B plot the population-weighted average of these parameters across age groups scenario for female and males. Each data point represents the impact on the annual mortality rate (per 100,000) of an additional day in the relevant

temperature bin, relative to the 70° - 75° F bin. The figure also plots the estimated θ_j's plus and minus one standard error of the estimates. The y-axes are scaled identically so that the response functions can be compared easily.

Both figures suggest that mortality risk is highest at the colder and hotter temperatures. It is evident that trading some days in the 50° - 80° F bin range for hotter days as is predicted in the Hadley 3 A1F1 scenario will lead to mortality increases. The figures underscore the lack of precision of the mortality estimates at the highest temperature bins and this explains the lack of precision in the Table 3 estimates. Cold days appear to be relatively harmful for men, but the imprecision of these estimates mean that they should be interpreted cautiously.

The approach of modeling temperature with 20 separate variables and allowing the impacts to vary by demographic group is very demanding of the data, so we assessed whether making some restrictions would help to allow for more precise inferences and generally concluded that the answer is no. For example, we estimated models that (XXXX INSERT LIST). None of these alternative specifications helped to reduce the standard errors substantially. Ultimately, the problem is that our predicted impacts of climate change rely so heavily on the mortality impacts at the highest temperature bins where the available data is more limited than is ideal.

Robustness Analysis. Table 4 reports on the estimated aggregate mortality impact for females and males from a series of alternative models. The top row repeats the overall estimate from Table 3 so that it can serve as the basis for comparisons. The second row details the results when the climate predictions come from the Hadley 3 B1 Scenario which predicts smaller temperature increases. This scenario produces mortality estimates that are smaller, roughly 20,000 extra fatalities per year. The estimated mortality impacts are essentially unchanged when the state by year effects are replaced with year fixed effects (row 3). The next two rows report the largest and smallest mortality impacts from the 30 years in the 2070-2099 period.[15]

The specifications in the remaining rows lead to the same qualitative conclusions. In rows (6)

[15] In all other parts of the paper, we use the predicted distribution of temperature averaged over the 2070-2099 period.

and (7) we use 40 separate temperature variables: (6) uses the same 20 temperature bins for the daily maximum and for the daily minimum temperatures and (7) uses the 20 temperature bins for the current years temperature and the previous year's temperature to allow for the possibility that we have not adequately accounted for the dynamics of the mortality-temperature relationship. Row (8) reports on the predicted impact of climate change based on $\hat{\theta}_{dj}^{TMEAN}$'s that are obtained using post-1980 data only. The intuition is that air conditioning ownership rates and medical technologies are more advanced in the later years. Surprisingly, the overall impact of predicted climate change is essentially unchanged.

Accounting for the Dynamics Relationship Between Temperature and Mortality. Figure 6 provides an opportunity to assess the paper's success at modeling the complicated dynamic relationship between temperature and mortality to address the issues of harvesting/forward displacement and delayed impacts. The figure replicates the daily analysis of Deschênes and Moretti (2005) and was constructed with the Multiple Cause of Death Files (MCOD) for the 1972-1988 period. The key difference with the CMF is that the MCOD files contain the exact date of death between 1972 and 1988.

We use these data to estimate daily and annual versions of equation (4). In both versions, males and females are combined. In the first, the dependent variable is the daily age-adjusted mortality rate across all age categories.[16] This equation includes county fixed effects, state by year fixed effects, state by month fixed effects, and the 20 temperature variables. The figure then reports the estimated parameters on the temperature variables.

In the annual version, the analysis is conducted exactly as described above, except that for comparison purposes we combine males and females in each age category and estimate separate models for the 8 age categories. The figure reports the population weighted averages of the $\hat{\theta}_{dj}^{TMEAN}$'s across the age categories, just as was done in Figures 5A and 5B (except here it is for both genders at once).

The figure reveals the shortcomings of the daily model that is prevalent in the previous literature (Cite TK). This is most evident at the coolest temperatures; the daily approach suggests that relationship between mortality rates are equal from the <0° F bin through the 70° - 75° F bin. A comparison with the

[16] We are currently constructing samples that will permit the estimation of the daily models separately for different age categories.

annual approach highlights that the delay impacts hypothesis is correct as the mortality impacts are substantially greater at the colder temperatures. For example, the average of the estimated $\hat{\theta}_j^{TMEAN}$ for the 7 bins representing temperatures 30° F and below is .02 which suggests that an extra day in that temperature range is associated with .02 additional deaths per 100,000 population. The analogous calculation from the annual approach is 0.22, which is more than 10 times larger.

The paper's primary purpose is to learn about the likely impacts of climate change and there are important differences between the estimated $\hat{\theta}_{dj}^{TMEAN}$ at the higher temperatures too. Here, the estimated coefficients from the daily model overstate the mortality impact of a hot day; for example the estimated impact of days in the 85° - 90° F and > 90° F bins are 0.14 and 0.13 larger, respectively, in the daily model. The result is that the predicted increase in mortality with the daily approach is 45,452 but only 19,138 with the annual approach.

Furthermore, these results suggest that the reports of the extremely elevated risk of mortality associated with hot days overstate the mortality impacts of these episodes (CITE Chicago and Paris Studies). This is because the individuals that die on these days appear to have had little life expectancy remaining, just as is predicted by the harvesting/forward displacement hypothesis. In this respect, the results confirm the Deschênes and Moretti (2005) findings although we have done so with a much more blunt approach; their paper traces out the precise dynamics of the mortality-temperature relationship on hot days.

Results for Cause-Specific Mortality. Table 5 reports on the estimated impacts of the Hadley 3 A1F1 scenario on deaths due to cardiovascular diseases, neoplasms, respiratory diseases, and motor-vehicle accidents. It is evident that the largest mortality impacts are concentrated among cardiovascular and respiratory diseases. Interestingly, the respiratory results are statistically different from zero at the 10% level for men 45 and older. The large decline in motor vehicle fatalities for males 15-24 would seem to be related to a reduction in dangerous driving days. Overall, this by-cause exercise is demanding of the data and it isn't possible to make precise inferences for many of the causes or demographic groups.

Geographic Variation in the Estimated Impacts. To Come.

C. Estimates of Adaptation from Energy Consumption

We now turn to an analysis of the effect of inter-annual fluctuations in temperature on residential energy consumption. Specifically, this subsection fits versions of equation (5) to the state by year data on residential energy consumption from the EIA. As a reminder, the annual mean of residential energy consumption in this period is 16.6 quads.

Figure 7 plots the estimated θ_j's from the specification that includes the familiar 20 temperature variables. The coefficients report the estimated impact of an additional day in bin j on annual energy consumption, relative to energy consumption on a day where the temperature is between 65° - 70° F. The estimates are adjusted for the ln of population and state gross domestic product, their squares and interaction. The figure also plots the estimated θ_j's plus and minus one standard error of the estimates so that the precision of each of these estimates is evident.

The response function has a U-shape, indicating that that energy consumption is highest in cold and hot days. Notably, the function turns up sharply at the last three bins. So, for example, an additional day in the > 90° F bin is associated with an extra 0.11 quads of energy consumption. The response function is very flat and precisely estimated for temperatures between 45 – 80° F; these seven estimated θ_j's all range between -0.01 and 0.01. In fact, the shape of this function undermines the convention in the literature of modeling heating and cooling degree days linearly with a base of 65 because fitting a line through these points will overstate consumption in the flat range and understate it at the extremes of the temperature distribution.

Table 6 presents the predictions of the impact of climate change on annual residential energy consumption from estimation of versions of equation (6). All specifications include state fixed effects and census division by year fixed effects. The specification details are noted in the bottom rows. For each specification, we reports results from modeling temperature with the 20 separate variables and with cooling and heating degree days (and their squares). The rows in the specification analysis panel provide results from some robustness checks.

The baseline specification in the first row reports compelling evidence that predicted climate change will cause a sharp increase in energy consumption. Specifically, the estimates suggest an increase

in residential energy consumption in the range of 4-6 quads or 25%-35%. All of these estimates would be judged to be statistically different from zero at conventional significance levels.

The specification checks generally support the validity of the findings in the baseline specification. The predicted impact of climate change is about 40% smaller when the post-1980 only data is used to obtain the estimated θ_j's. One interpretation is that the proliferation of energy saving technologies (e.g., better insulation and greater efficiency of appliances) was increasing faster than air conditioning. The modeling of temperature with 20 separate variables for the daily minimum and maximum leaves the estimates unchanged. A model that replace the census division by year fixed effects with year fixed effects leads to smaller impacts when temperature is modeled with 20 variables and larger ones with heating and cooling degree days. Hadley 3 B1 predictions lead to smaller increases in energy consumption with the 20 variables and precisely estimated declines with heating and cooling degree days.

Finally, an alternative approach is to estimate the relationship between energy consumption and temperatures in the cross-section (Mansur, Mendelsohn, and Morrison 2007; Mendelsohn 2006). In the absence of specification error, this approach will reveal the equilibrium relationship between energy consumption and temperature. However, this relationship isn't informative about the social costs associated with obtaining the presumably greater levels of energy efficiency in the hotter parts of the country; consequently it is not especially useful in determining the predicted costs of climate change. In principle, it would be possible to use land markets to measure the full costs of these adaptations, however as Deschenes and Greenstone (2007) have demonstrated the cross-sectional hedonic for land prices appears to suffer from substantial misspecification or omitted variables bias, making this approach less appealing from a practical perspective.[17]

Overall, these results imply that predicted climate change will lead to substantial increases in energy consumption in the residential sector. This finding is consistent with predicted increases in energy consumption from a study of California (Franco and Sanstad 2006). To the best of our knowledge, these

[17] Despite the limitations of this cross-sectional approach, we implemented it with our national data on residential energy consumption by estimating versions of equation (5), except these versions all dropped the state fixed effects. Thus, this is a pooled cross-section. This approach produced inconsistent results with the 20 variable specification yielding negative and statistically significant declines in energy consumption, while the heating and cooling degree day approach led to increases in energy consumption that are similar to those in the baseline specification in Table 6.

estimates on energy consumption are the first ones based on data from the entire country. In addition to being useful for policy purposes, they should help climate modelers who have not yet incorporated feedback effects from higher energy consumption into their models.

VI. Interpretation

We now use the above findings to calculate a partial measure of the economic costs of climate change. The Hadley 3 A1F1 results in Table 3 suggest that climate change would lead to a loss of roughly 750,000 life years annually. A valuation of a life year at about $100,000 is roughly consistent with Ashenfelter and Greenstone's (2004) estimate of the value of a statistical life. So if we ignore the sampling variability, the Hadley 3 A1F1 results suggest that the direct impacts of climate change on mortality will lead to annual losses of roughly $75 billion in the 2070-2099 period.[18]

The social cost of the additional energy consumption should be added as a welfare loss. The estimates from Table 6 imply an increase in consumption of 4-6 quads. The average cost of a quad in 2006$ between 1990 and 2000 is $7.6 billion, so this implies that there would be an additional $30-$45 billion of energy consumption at the end of the century.

There are a number of caveats to these calculations and to the analysis more generally that bear noting. First, the life-year calculation assumes that the individuals' whose lives are affected by the temperature changes had a life expectancy of 78.6 for women and 71.2 for men. It is certainly possible that our efforts to purge the influence of harvesting and delayed impacts were not entirely successful and, in this case, the estimated impact on life years would be smaller.

Second, it is likely that these calculations do not reflect the full impact of climate change on health. In particular, there may be increases in the incidence of morbidities due to the temperature increases. Additionally, there are a series of indirect channels through which climate change could affect human health, including greater incidence of vector borne infectious diseases (e.g., increased incidence of malaria and dengue fever). Further, it is possible that the incidence of extreme events would increase and

[18] It is also possible to make a similar calculation using estimates of how the value of a statistical life varies over the life cycle (Murphy and Topel 2006).

these could affect human health (Kerry 200X). This study is not equipped to shed light on these issues.

Third, the brief theoretical section highlighted that our estimates likely overstate the increase in mortality and energy consumption due to climate change. This is because the higher temperatures are likely to cause individuals to increase expenditures on goods that protect themselves from the changes in temperature. Our identification strategy relies on inter-annual fluctuations in weather, rather than a permanent change, and individuals can only engage in a limited set of responses or adaptations in response to the weather being unseasonably hot for a few days. For example, as the load results indicate, individuals can use more air conditioning on hot days and some people might even be motivated to buy an air conditioner.

However, there are a number of adaptations that cannot be undertaken in response to a single year's weather realization. For example, permanent climate change is likely to lead to some migration (presumably to the North) and this will be missed with our approach. Although these adaptations may be costly, individuals will only undertake them if they are less costly than the alternative. For this reason, our approach is likely to overstate the costs of climate change.

VII. Conclusions

To Come.

Data Appendix

I. Hadley 3 Census Division-Level Predictions

We downloaded the Hadley Climate Model 3 (HadCM3) data from the British Atmospheric Data Centre (http://badc.nerc.ac.uk/home/), which provides a wealth of atmospheric data for scientists and researchers. Hadley Centre data appears on BADC thanks to the Climate Impacts LINK Project, a distributor of archived climate model output to researchers. Daily climate predictions generated by the Hadley 3 model are available for all future years from the present to 2099 and for several climate variables – we downloaded the predicted maximum and minimum temperatures and precipitation levels for each day during the years 2070-2099.

The HadCM3 grid spans the entire globe; latitude points are separated by 2.5°, and longitude points are separated by 3.75°. We use the 89 gridpoints that fall on land in the contiguous United States to develop climate predictions for the 9 US Census Divisions. At the Census Division level, each day's mean temperature is calculated as the simple average across all grid points within the Division. The data used in this paper was originally generated by the Hadley Centre for the International Panel on Climate Change's (IPCC) Special Report on Emissions Scenarios (SRES).

II. EIA Energy Consumption Data

The consumption data is derived from several different reports and forms depending on energy source. Coal consumption data for most sectors comes from the EIA's Annual Coal Report; electric power sector coal use is the exception, coming instead from forms EIA-906 "Power Plant Report" and EIA-920 "Combined Heat and Power Plant Report". Natural gas consumption data comes from the EIA's Natural Gas Annual. Most petroleum data is the "product supplied" data found in EIA's Petroleum Supply Annual, with the exception again of electric power sector use, which is reported on EIA-906 and EIA-920. Solar, wind, geothermal, and most biomass energy use data is also reported on those forms. Residential and commercial use of biomass is reported on forms EIA-457 "Residential Energy Consumption Survey" and "Commercial Buildings Energy Consumption Survey". Nuclear electric power and other electricity data comes from the EIA Electric Power Annual. Finally, system energy losses and interstate flow are estimated in the State Energy Data System.

References (incomplete)

Deschenes, Olivier and Michael Greenstone, "The Economic Impacts of Climate Change: Evidence from Agricultural Profits and Random Fluctuations in Weather," *American Economic Review*, (forthcoming) 2006.

Deschenes, Olivier and Enrico Moretti. "The Impact of Extreme Weather on Mortality and Longevity," Mimeograph. 2005

Energy Information Administration, Form EIA-826, "Monthly Electric Sales and Revenue Report with State Distributions Report." July 2006.

Franco, Guido and Alan H. Sanstad. "Climate Change and Electricity Demand in California." 2006

Huynen, Maud M.TE., Pim Martents, Dieneke Schram, Matty P. Weijenberg, and Anton E. Kunst. "The Impact of Heat Waves and Cold Spells on Mortality Rates in the Dutch Population." *Environmental Health Perspectives* vol. 109(5) May 2001. pp. 463-70.

Houghton, J.T., L.G. Meira Filho, B.A. Callender, N. Harris, A. Kattenberg, and K. Maskell, eds. (1996): *Climate Change 1995: The Science of Climate Change. Contribution of Working Group I to the Second Assessment of the Intergovernmental Panel on Climate Change*, Cambridge, England: Cambridge University Press.

Intergovernmental Panel on Climate Change (1990): *Climate Change: The IPCC Scientific Assessment* (J. T. Houghton, G. J. Jenkins, and J. J. Ephraums, eds.), Cambridge: Cambridge University Press.

International Panel on Climate Change (IPCC) Special Report on Emissions Scenarios (SRES). 1996.

IPCC 2001. Third Assessment Report….

Mansur, Mendelsohn, and Morrison. "Climate Change Adaptation: A Study of Fuel Choice and Consumption in the Energy Sector." 2007

Mendelsohn, R. 2006. "Energy Impacts in J. Smith and R. Mendelsoh (eds.) The Impacts of Climated Change on Regional Systems: A compreshensive Analysis of California. Edward Elgar Publishing, Northampton, MA.

Moore, Thomas Gale, "Health and Amenity Effects of Global Warming," *Economic Inquiry* 36: 471-488 1998.

Newell, Richard and Billy Pizer. 200Y. To come

Rooney C., A.J. McMichael, R.S. Kovats, and MP Coleman, "Excess Mortality in England and Wales, and in Greater London, During the 1995 Heatwave," *Journal of Epidemiological Community Health* 52:482-486 1998

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher (2006): "The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions," *Review of Economics and Statistics*, 88(1): 113-125.

Tol, Richard S.J., "Estimates of the Damage Costs of Climate Change, Part I Benchmark Estimates,"

Environmental and Resource Economics 21: 47-73 2002a.

Tol, Richard S.J., "Estimates of the Damage Costs of Climate Change, Part II Dynamic Estimates," *Environmental and Resource Economics* 21: 135-160 2002b.

Weitzman, Martin L. (2001): "Gamma Discounting," *American Economic Review*, 91: 260-71.

World Health Organization. 2003. Climate Change and Human Health – Risks and Responses.

Table 1: Average Annual Mortality Rates per 100,000 Population, 1968-2002

	All Causes of Death		Cardiovascular Disease		Neoplasms		Respiratory Disease		Motor-Vehicle Accidents	
Age Group:	**Females**	**Males**	**Females**	**Males**	**Females**	**Males**	**Females**	**Males**	**Females**	**Males**
Infants	1,031.1	1,292.1	21.1	25.4	4.6	4.7	45.2	58.6	---	---
1-14	29.9	41.8	1.5	1.7	3.9	4.9	1.6	1.8	3.1	4.7
15-24	53.2	151.6	3.1	4.5	4.9	7.2	1.4	1.8	11.8	36.1
25-44	68.0	148.5	10.8	24.3	18.0	15.8	2.2	3.1	4.8	16.5
45-54	386.1	699.1	100.1	269.0	158.1	165.0	14.3	22.1	7.0	20.3
55-64	917.3	1,696.0	315.8	757.6	361.4	501.4	46.2	82.3	7.1	18.4
65-74	2,108.0	3,754.8	942.6	1,789.2	644.8	1,073.0	136.3	272.8	9.3	19.6
75-99	7,571.8	9,982.7	4,512.1	5,357.6	1,088.3	1,942.1	578.0	975.6	13.3	34.4
All Ages	804.4	939.2	389.7	409.8	174.5	206.3	53.2	64.4	8.1	21.4

Table 2: Population-Weighted Averages of Daily Mean across Counties, 1968-2002

	Actual	Hadley 3, B1		Hadley 3, A1F1	
		Level	Difference	Level	Difference
Average Daily Mean					
All Counties	56.6	57.1	0.5	62.6	6.0
(Std Dev)	(8.1)	(3.6)		(7.4)	
Northeast Region	50.9	54.0	3.1	56.5	5.6
Midwest Region	49.9	55.0	5.1	59.1	9.2
South Region	64.4	62.4	-2.0	73.6	9.2
West Region	58.6	55.4	-3.2	58.3	-0.3
Average Daily Maximum					
All Counties	46.0	47.2	1.2	52.5	6.5
(Std Dev)	(8.1)	(3.7)		(7.7)	
Average Daily Maximum					
All Counties	67.2	66.9	-0.3	72.7	5.5
(Std Dev)	(8.8)	(3.7)		(7.5)	

Table 3: Estimates of the Impact of Climate Change on Annual Mortality, Hadley 3-A1F1 Scenario

	FEMALES				MALES			
Age Group:	**Annual Fatalities** (1a)	**Change in Mortality Rate** (1b)	**Annual Change in Life Years** (1c)	**F-test (p-values)** (1d)	**Annual Fatalities** (2a)	**Change in Mortality Rate** (2b)	**Annual Change in Life Years** (2c)	**F-test (p-values)** (1d)
Infants	995.9 (518.7)	5.5%	-77,302.5	0.180	1,835.3 (681.7)	7.8%	-128,671.4	0.468
1-14	578.4 (320.6)	8.0%	-41,476.0	0.873	693.8 (414.4)	6.6%	-44,643.9	0.658
15-24	-579.2 (498.0)	-6.0%	34,723.6	0.038	263.4 (1408.3) 0.187025345	0.9%	-13,903.8	0.079
25-44	860.1 (952.6)	2.2%	-39,124.3	0.015	(3096.1)	3.6%	0.0	0.001
45-54	947.9 (907.3)	1.8%	-29,926.6	0.207	2935.6 (1674.7)	3.3%	-75,796.9	0.316
55-64	1098.7 (1263.0)	1.1%	-25,369.0	0.133	4,372.5 (2166.6)	2.7%	-79,360.6	0.003
65-74	2,166.8 (2547.5)	1.2%	-33,021.8	0.133	314.2 (3398.0)	0.1%	-3,735.2	0.001
75-99	11,528.6 (5444.6)	2.0%	-78,279.5	0.370	4,158.6 (4128.8)	1.0%	-22,581.4	0.009
Aggregate Impact	**17,597.2** (12452.3)	**1.3%**	**-289,776.1**		**14,573.3** (16968.5)	**1.3%**	**-368,693.3**	

Table 4: Alternative Estimates of the Impact of Climate Change on Annual Mortality

	Females	Males
1. Baseline	**17,597.2** (12452.3)	**17,566.8** (16968.5)
2. Baseline, Hadley 3 B1 Scenario	**11,140.0** (5779.4)	**8,989.0** (7983.4)
3. Year Effects Only, H3-A1F1	**19,746.5** (17092.8)	**15,270.9** (17163.1)
4. Smallest Annual Impact H3-A1F1, 2070-2099	**9,862.8** (9104.6)	**9,621.8** (20899.3)
5. Largest Annual Impact H3-A1F1, 2070-2099	**28,161.6** (15333.3)	**28,698.4** (23734.5)
6. Controls for Minimum and Maximum Instead of Mean (H3-A1F1)	**15,847.9** (12275.8)	**25,551.4** (14255.2)
7. Includes lagged temperature	**9,200.5** (20286.5)	**14,753.6** (24017.6)
8. Post-1980 Data Only	**21,928.7** (11733.9)	**13,688.5** (16410.7)

Table 5: Estimates of the Impact of Climate Change on Annual Cause-Specific Mortality, Hadley 3-A1F1 Scenario

	Cardiovascular Diseases		Neoplasms		Respiratory Diseases		Motor-Vehicle Accidents	
	Females	**Males**	**Females**	**Males**	**Females**	**Males**	**Females**	**Males**
Age Group:								
Infants	121.5 (62.7)	140.8 (71.9)	-70.3 (30.3)	-35.9 (42.8)	25.3 (151.5)	123.7 (242.1)	---	---
1-14	93.8 (59.6)	65.9 (70.8)	7.2 (130.2)	-46.1 (113.0)	65.8 (94.0)	7.6 (81.0)	-195.2 (99.5)	-193.2 (138.3)
15-24	63.2 (93.5)	95.9 (110.6)	59.0 (108.5)	28.4 (115.1)	-88.6 (72.3)	90.0 (65.2)	-468.5 (216.7)	-1055.3 (443.8)
25-44	70.3 (386.3)	159.5 (473.4)	-480.2 (320.4)	144.3 (350.1)	-25.3 (130.6)	145.7 (196.5)	-194.9 (250.1)	609.3 (596.5)
45-54	305.9 (441.3)	511.1 (880.3)	510.7 (573.4)	142.9 (496.8)	138.3 (177.7)	484.9 (221.2)	44.4 (115.5)	107.5 (253.4)
55-64	243.4 (820.8)	2136.3 (1224.8)	-954.4 (782.8)	-838.1 (1283.2)	-309.8 (309.2)	639.0 (328.7)	174.2 (147.7)	101.3 (188.5)
65-74	1,336.1 (1685.6)	-732.3 (1949.0)	-2,178.7 (1135.6)	-2,577.1 (1230.2)	324.5 (599.5)	1528.2 (790.3)	53.4 (113.6)	-263.4 (156.7)
75-99	1,700.4 (3484.8)	2,050.6 (2882.2)	1,421.0 (1495.6)	-2,834.8 (1276.7)	1,342.4 (1500.2)	2323.6 (1234.0)	-167.6 (104.0)	-15.0 (152.6)
Aggregate Impact: Std Error	**3,934.6** (7034.7)	**4,427.7** (7663.0)	**-1685.6** (4576.9)	**-6,016.4** (4907.9)	**1,472.6** (3035.1)	**5342.8** (3158.9)	**-754.2** (1047.1)	**-708.8** (1929.9)
Total Years Lost (-) or Gained (+)	**-69,235.3**	**-77,790.5**	**52,963.7**	**52,698.9**	**-10,547.4**	**-72,001.2**	**48,807.1**	**51,387.5**

Table 6: Estimates of the Impact of Climate Change on Annual Energy Consumption, Hadley 3-A1F1 Scenario

	(1)	(2)	(3)	(1)	(2)	(3)
Baseline Model	**4.1**	**5.0**	**4.9**	**5.1**	**6.1**	**6.0**
	(1.9)	(1.9)	(1.9)	(1.5)	(1.5)	(1.5)
Specification Analysis:						
Fixed-Effects Model	**2.6**	**3.1**	**3.0**	**2.8**	**3.7**	**3.6**
(post-1980)	(1.3)	(1.1)	(1.1)	(0.9)	(1.0)	(1.0)
Fixed-Effect Models	**4.7**	**5.1**	**4.9**	**---**	**---**	**---**
20 Cells for Min and Max	(1.6)	(1.6)	(1.6)			
Fixed-Effect Models	**2.0**	**2.9**	**2.7**	**5.1**	**7.9**	**7.8**
Plus Year Effects	(1.9)	(1.9)	(1.8)	(1.4)	(1.9)	(1.9)
Fixed-Effect Models	**1.1**	**1.4**	**1.4**	**-2.8**	**-2.7**	**-2.8**
Hadley 3-B1	(0.5)	(0.5)	(0.5)	(0.5)	(0.5)	(0.5)
Controls for Log(POP)	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
Controls for Log(GDP)	**No**	**Yes**	**Yes**	**No**	**Yes**	**Yes**
Interactions	**No**	**No**	**Yes**	**No**	**No**	**Yes**
Model Temperature with 20 Variables	**Yes**	**Yes**	**Yes**	**No**	**No**	**No**
Heating and Cooling DD (and Squares)	**No**	**No**	**No**	**Yes**	**Yes**	**Yes**

Figure 1: Distribution of Annual Daily Mean Temperatures (F), 1968-2002

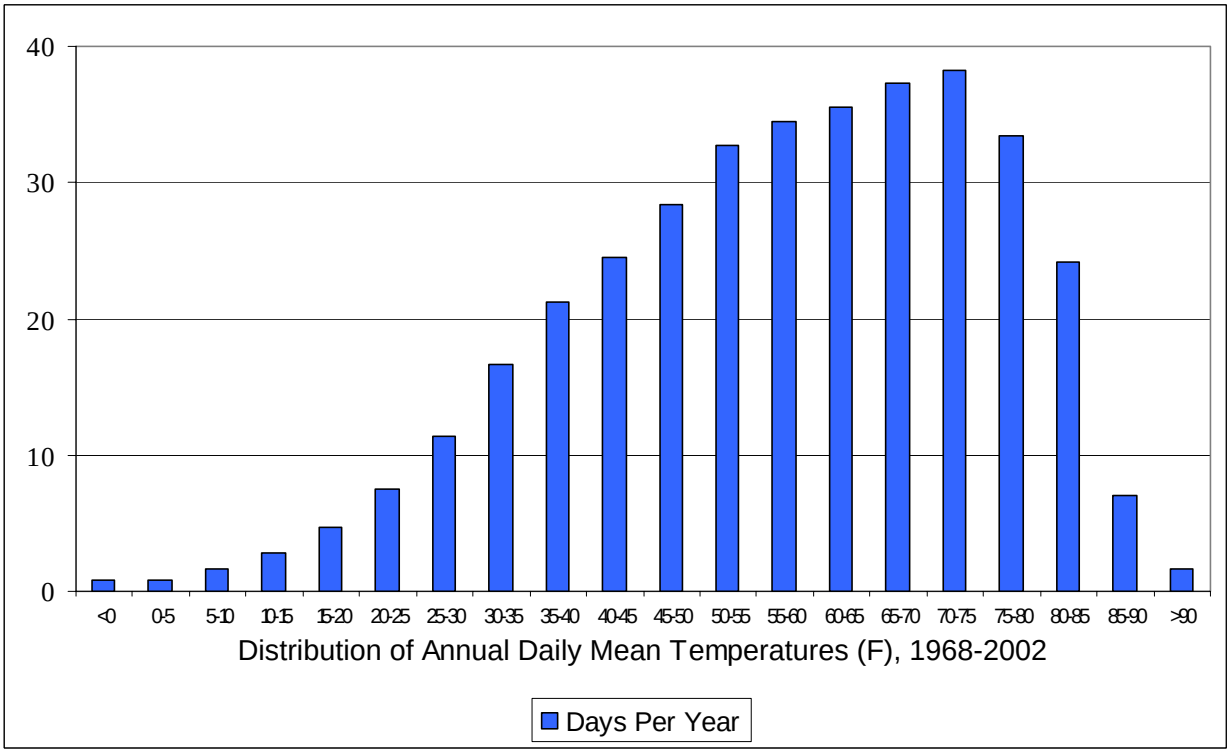


Figure 2: Changes in Distribution of Annual Daily Mean Temperatures (F) Under 3 Climate Change Scenarios

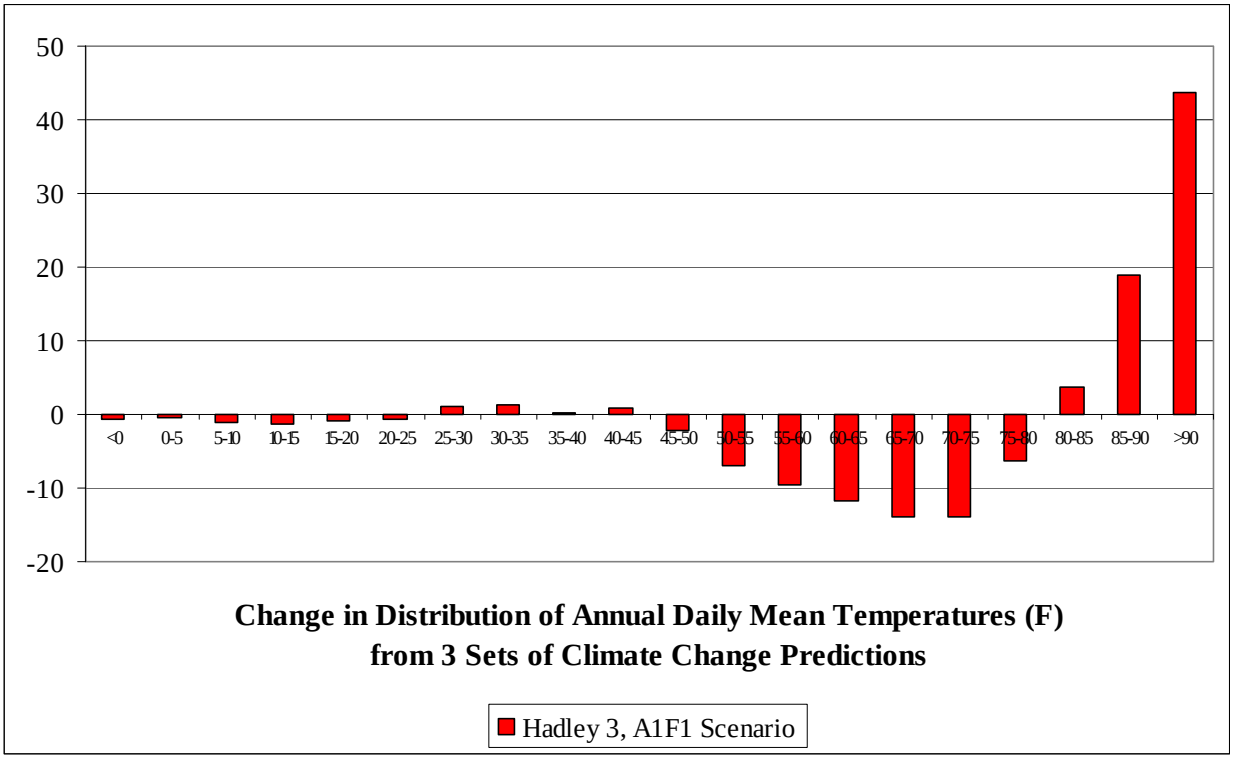


Figure 3: Residual Variation in Annual Daily Mean Temperatures (F), 1968-2002

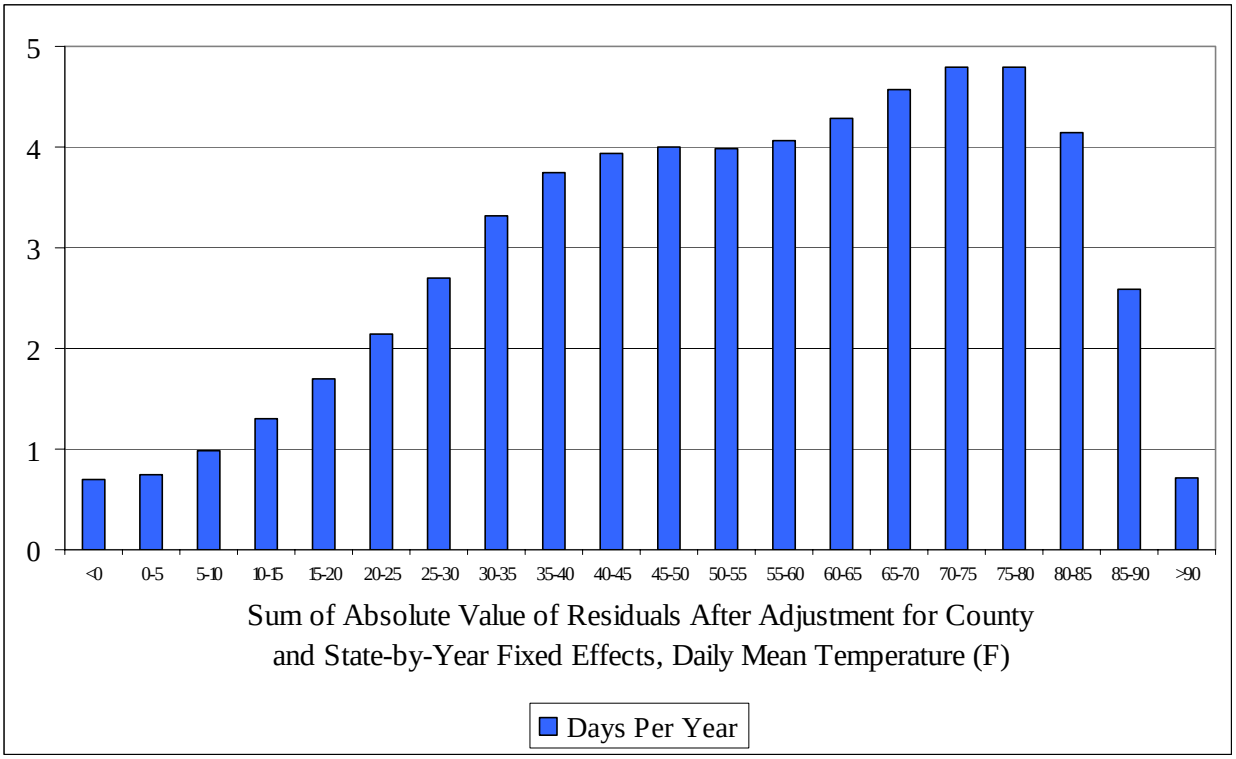


Figure 4: Estimated Regression Coefficients, Male Infants (relative to temperature cell 65-70)

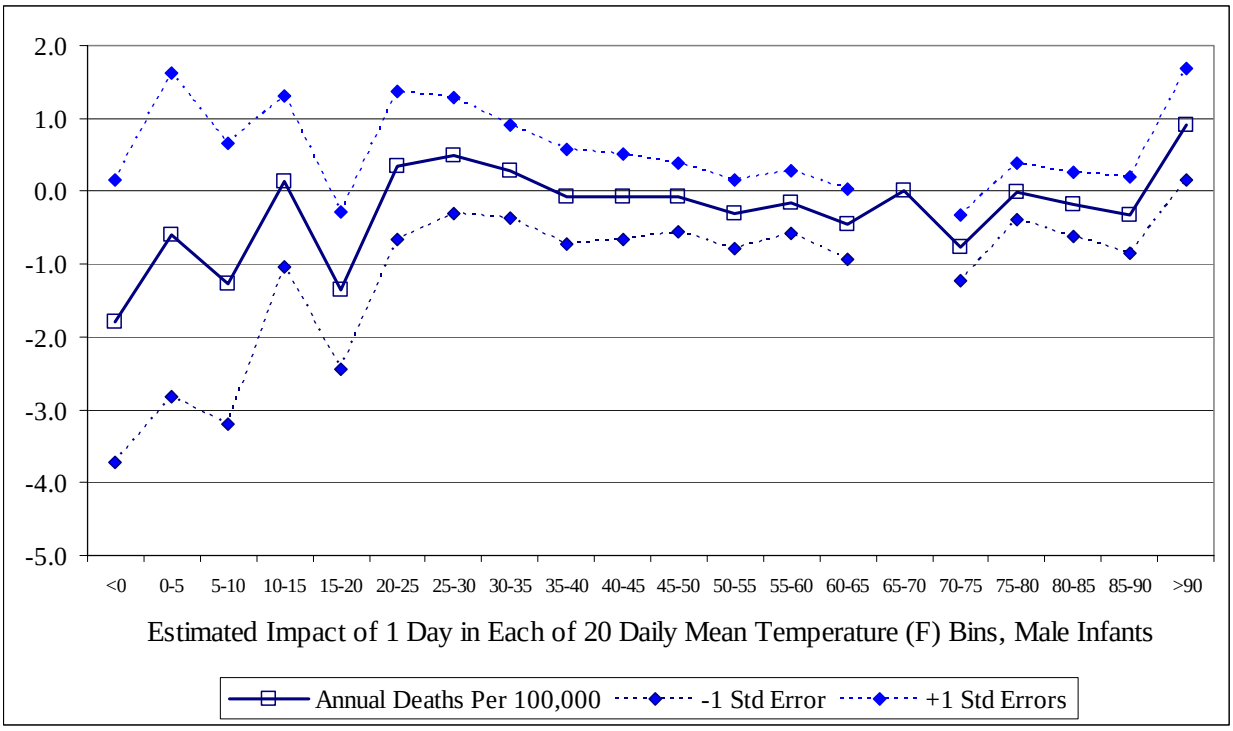


Figure 5A: Population-Weighted Average Regression Estimates Across Age Groups, Females (relative to temperature cell 65-70)

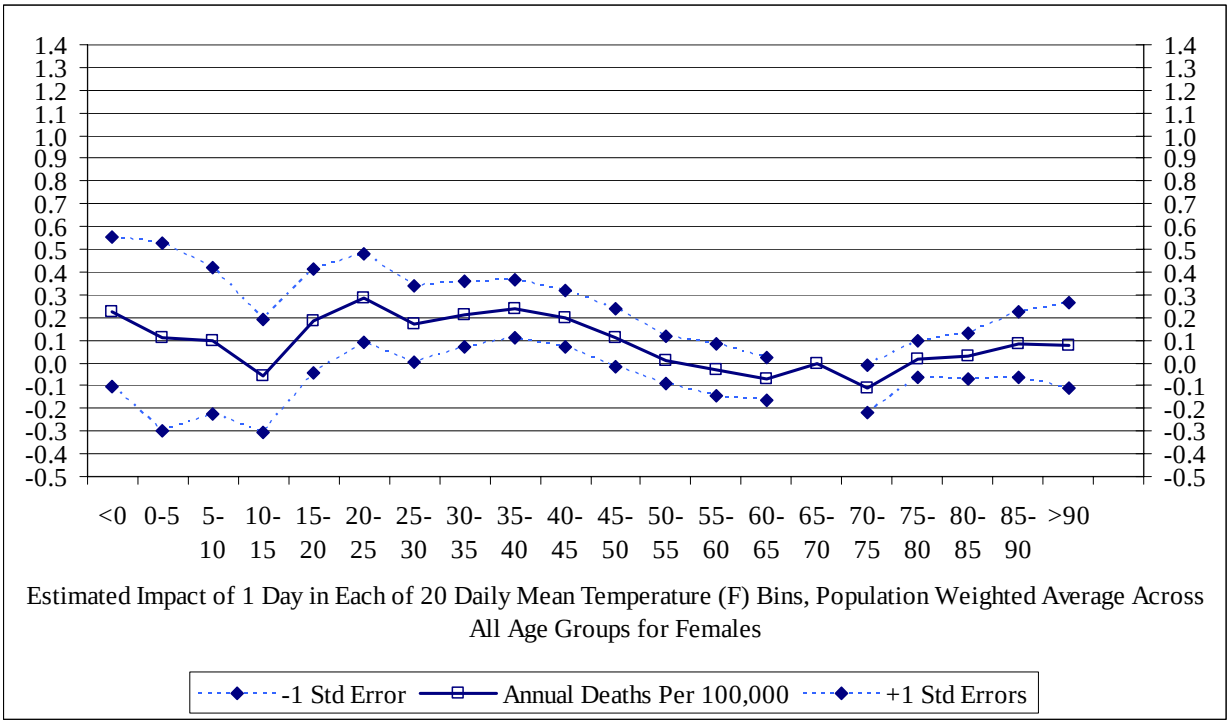


Figure 5B: Population-Weighted Average Regression Estimates Across Age Groups, Males (relative to temperature cell 65-70)

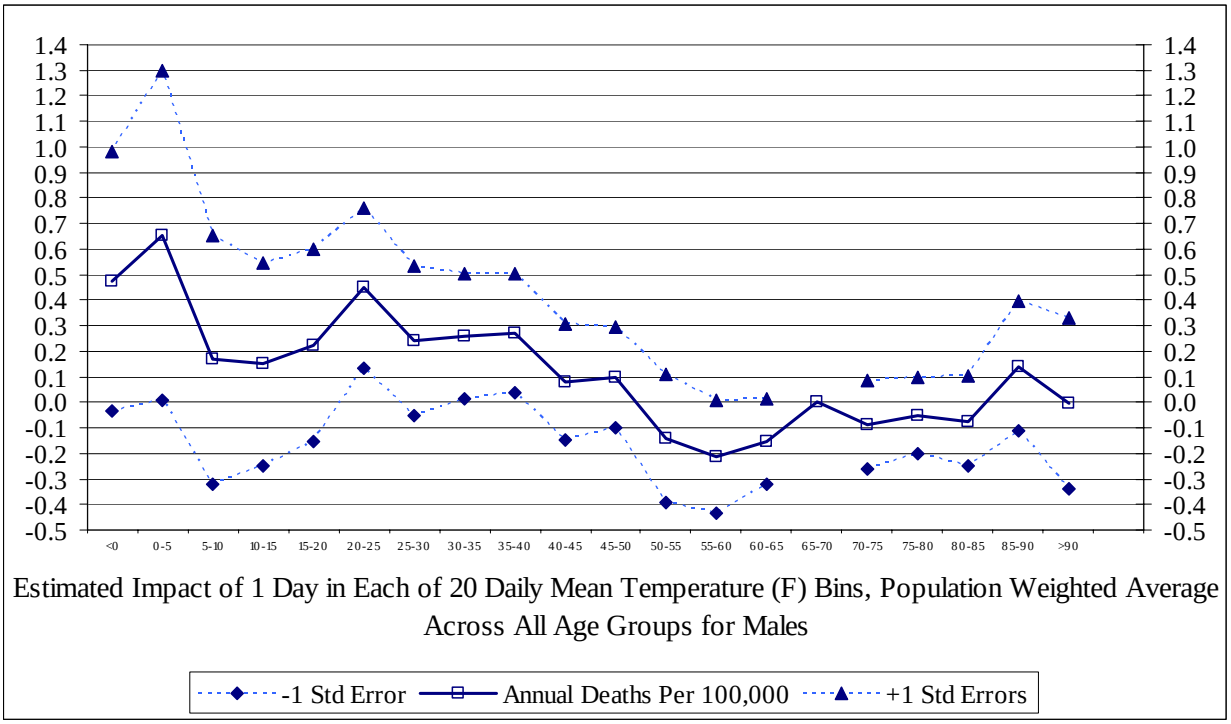


Figure 6: Population-Weighted Average Regression Estimates Across Age Groups, for Daily and Annual Approaches (relative to temperature cell 65-70)

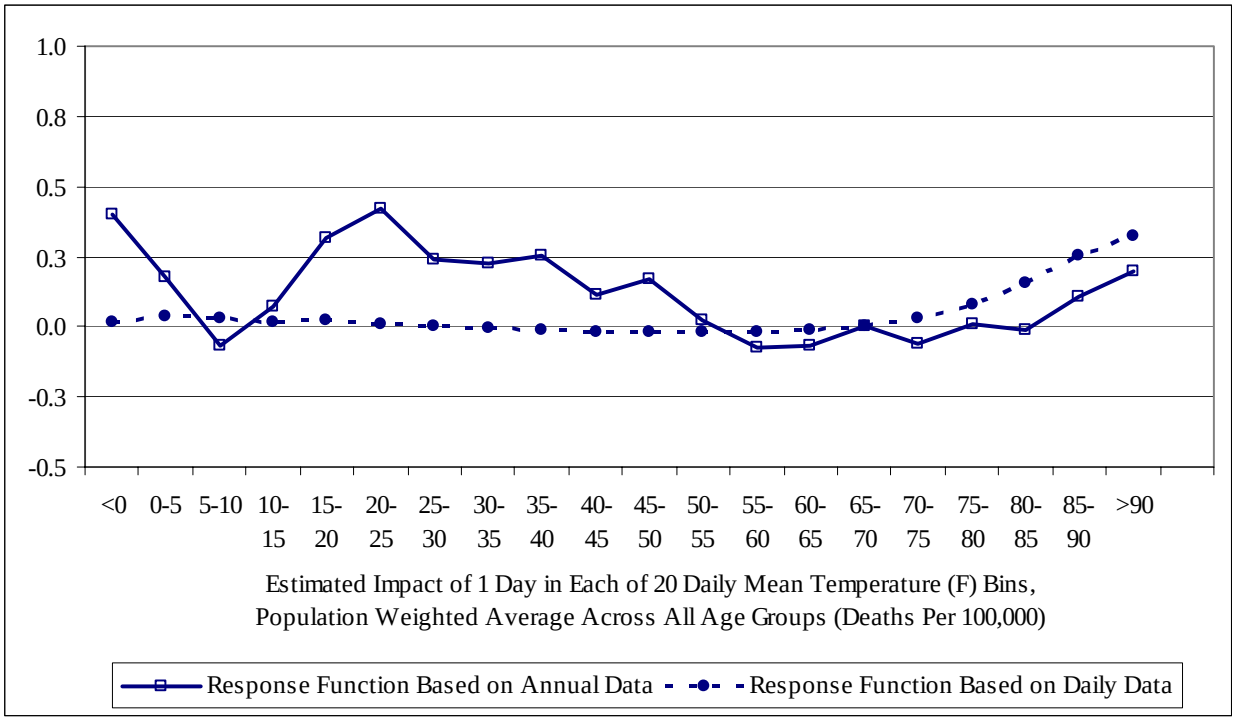


Figure 7: Estimated Impact on Total Energy Consumption in the Residential Sector

