

THE IMPACT OF AIR TEMPERATURE ON MORBIDITY, MORTALITY  
AND HEALTHCARE COST IN THE MEDICARE POPULATION

Garth Heutel

Georgia State University and NBER

Nolan H. Miller

University of Illinois at Urbana-Champaign and NBER

David Molitor

University of Illinois at Urbana-Champaign

July 2, 2014

**PRELIMINARY DRAFT: DO NOT CITE OR CIRCULATE**

**Abstract:** This paper merges weekly average temperature data from the National Oceanic Atmospheric Administration's National Climatic Data Center (NDCC) with Medicare claims data in order to analyze the impact of high and low temperatures on mortality, onset of new chronic conditions, and hospital spending. We find a U-shaped pattern to mortality, with high and low temperature weeks exhibiting higher mortality than a 70 degree reference week. The marginal deaths in extreme weeks are healthier than the typical person who dies in the reference week, but less healthy than the population as a whole. We find some evidence of short-term mortality displacement at moderately high temperatures, but not at extremely high temperatures and not for low temperatures, where the impact tends to grow over time. High temperatures are associated with increased onset of new chronic conditions, while low temperatures are associated with lower onset, although this result may be driven by differences in the propensity to access the health care system. In the short run, high temperatures are associated with increased Medicare hospital spending and lower temperatures are associated with decreased hospital spending, although over a one month period both high- and low-temperature weeks are associated with increased hospital spending. Using conventional figures for the value of a life year lost, we find the additional healthcare spending induced by a hot week to be about 3-6% of the mortality cost.

## I. INTRODUCTION

While it is widely recognized that environmental factors such as exposure to temperature extremes or elevated levels of air pollution impact health outcomes (e.g., morbidity or mortality), a large number of policy-relevant questions remain unanswered. Although raw mortality effects are, perhaps, most easily measured, the magnitude of the benefits of reducing mortality depends on whether the deaths due to higher pollution were from the frailest individuals – people who would have otherwise died in a few days, weeks or months, a phenomenon known as “harvesting” or “mortality displacement” – or whether those killed by pollution were otherwise healthy and would have lived for many more years in the absence of a temperature or pollution shock.

In addition to potentially affecting the number of deaths, environmental events also impact individuals who do not die. One common measure of the morbidity effects of environmental events is the number of additional hospital admissions caused by temperature or pollution events. However, displacement effects can also be a concern here, since adverse events might accelerate admissions that would have occurred in the near future or postponed discretionary admissions until after the environmental event has subsided. In order to estimate the true marginal effect of an environmental event on hospital admissions, it is necessary to net out any admissions that would have occurred even in its absence. Further, since not all hospital admissions are equally costly, it is useful to measure the cost of care in addition to counting admissions. Finally, while short term changes in morbidity and mortality may pick up the acute effects of the environmental shock, temperature events also affect individuals’ health stock, moving basically healthy individuals to a vulnerable health state, which may signal higher future health care costs and earlier death than would otherwise have occurred.

In this paper, we combine geographically detailed data on temperature from the National Oceanic Atmospheric Administration’s National Climatic Data Center (NDCC) with annual summary and detailed claims data on Medicare enrollees from the Centers for Medicare & Medicaid Services (CMS) for 1999 until 2011 in order to address some of these questions. One of the challenges in estimating the impact of extreme temperatures on health is that temperatures tend to be correlated over time, and extreme temperature events in one period affect the vulnerability of the population to future events. For example, a hot week may increase mortality, and to the extent that those who are killed were among the sicker individuals, this will result in a

healthier overall population entering the next week. The detailed, panel nature of the Medicare data allows us to measure the health status of the population in an area on a weekly basis and therefore to separate the impact of temperature events on the individuals alive in a particular period from its impact on the composition of the population over time.

We begin by looking at the relationship between temperature and mortality, where we find the U-shaped pattern common in the literature. We then use several approaches to looking at the question of harvesting and find that, contrary to conventional wisdom, the marginal person killed by an extreme temperature event is healthier than the typical person who dies but sicker than the typical person who lives. In terms of life years lost, then, these deaths are more costly than the typical death. Next, we turn to the impact of extreme temperatures on the health stock of those who do not die. We find that, in the short run, high temperatures are associated with greater onset of chronic conditions, while cold temperatures are associated with lower onset of new chronic conditions. To the extent that these differences are caused by true changes in health status, this suggests that hot weeks make the surviving population less healthy and cold temperatures protect the health of survivors. However, since onset of chronic conditions in our data is triggered by an encounter with the health system, we cannot rule out that these differences are caused by differences in the propensity to go to the hospital rather than true changes in health. Finally, we look at the association between the use of health care services and extreme weather events. We find an increase in hospital admissions and spending in hot weeks and a decrease in admissions and spending in cold weeks, both relative to the 70 degree reference week. Over the four weeks beginning with the event week, the pattern continues to hold in hot weeks, but in cold weeks we seen an increase in admissions and spending, suggesting some of the one-week reduction in admissions and spending may have been due to people delaying going to the hospital until the weather improved. Comparing the incremental hospital spending to the cost of life years lost associated with hot weeks, we find that incremental hospital costs are about 3-6% the size of the cost of life years lost.

This paper contributes to the existing literature in several ways. First, we use weekly panel data for Medicare beneficiaries for thirteen years to estimate the impact of temperature changes on mortality. We address harvesting in two ways. First, following the existing literature, we look at post-event periods ranging between one and four weeks. If the mortality effects of environmental events are due to short-term mortality displacement, we should see the estimated effects moderate

as the post-event window is expanded. Second, we construct measures intended to capture the health status of those who die following an environmental event in order to directly test whether the effects of the shock are concentrated in older and/or sicker enrollees. In particular, we construct a measure of life years lost due to the environmental shock based on the age and gender of those who lose their lives. Finally, we use data on chronic conditions in the Medicare data in order to directly assess whether incremental deaths due to environmental events are concentrated in those with more chronic conditions.

A second major contribution of the paper is in estimating the cost to Medicare due to environmental events. Part of this cost will be due to expenditures made in an attempt to save the lives of those who eventually die, part will be due to expenditures made to save the lives of people who would have died without treatment, and part will be due to costs aimed at restoring the health of people who were adversely affected by the environmental event but would not have died without treatment. Using conventional figures for the value of a life year lost, we find that the marginal expenditure to Medicare induced by temperature events is approximately 3-6% as large as the value of life years lost due to increased mortality during hot weeks, suggesting that while the increased hospital spending caused by weather events is important, the direct cost of mortality is a much larger component of the cost of temperature events.

The third major contribution of this paper is in quantifying the change in health status of those who are exposed to, but do not die from, the environmental event. For example, if exposure to a high-heat day induces a heart attack in a previously healthy person but does not kill him, this affects expected future life expectancy (i.e., there is a loss of life years even though there is no immediate loss of life) and future health care costs. Both of these are likely to be important to policymakers in evaluating the costs and benefits of climate change or pollution reduction. In this paper, we make a first pass at quantifying these effects by measuring the impact of environmental events on onset of new chronic conditions.

## II. BACKGROUND AND RELATED LITERATURE

A large literature investigates the relationship between ambient temperature and health outcomes. As surveyed in Deschenes (2013), the temperature-health relationship has been studied in both the economics and epidemiology/public health literatures. A robust finding, which we

confirm in our analysis, is that both extremely low and extremely high temperatures are associated with excess mortality, leading the temperature-mortality relationship to take on a U- or V-shape (Huynen et al., 2001; Basu and Samet, 2002; Ren, Williams and Tong, 2006; Baccini et al., 2008; Deschenes and Greenstone, 2011). However, interpreting this finding is complicated by the fact that mortality in most locations exhibits a strong seasonal pattern, with the death rate being highest in the winter and lowest in the summer. This is true even in areas with temperate climates (Curriero et al., 2002; Peng et al., 2005).

The peak in deaths during the winter – what has become known as “excess winter mortality” – may be caused by a number of factors, some of which are not directly related to temperature. Some of the excess winter deaths are due to influenza, which peaks during the winter season, while others can be attributed to increased exposure to indoor air pollution due to the use of wood fires and other combustion-based heating technologies (Moller, 2011). Nevertheless, there remains significant excess winter mortality even after accounting for these factors, which has been attributed to changes in blood chemistry and cardiovascular stress induced by exposure to the cold (Pan, Li and Tsai, 1995; Keatinge and Donaldson, 1997; Huynen et al., 2001).

As noted in the survey by Oudin Åström et al. (2011) a relatively small number of studies in the public health/epidemiology literature have investigated the role of mortality displacement in heat-related mortality. What evidence there is tends to be mixed and dependent on the location of the study. Hajat et al. (2005) find mixed results, with evidence of short-term harvesting after heat waves in London but not in Dehli. In the Czech Republic, Kyselý (2004) finds a relative rise in total mortality of 13% during heat waves, but a net mortality effect of only 1%, attributing the rest to short-term mortality displacement. Toulemon and Barbieri (2006) find evidence of modest harvesting during the 2003 heat wave in France. To the extent that a clear pattern emerges, it is that excess mortality during hot days tends to be due to short term displacement, at least to some extent, while the excess mortality during cold periods tends to build over time (Braga et al, 2001; Huynen et al., 2001; Deschenes and Moretti, 2009; Deschenes and Greenstone, 2011).<sup>1</sup>

---

<sup>1</sup> Our results confirm this general finding in the case of cold days. However, while we find some evidence that the additional mortality for moderately high temperatures is due to short-term displacement, there does appear to be an increase in mortality at the highest temperatures that is not only due to short-term harvesting. One reason why our results may differ from these other studies is that our analysis is run at the weekly level, so that mortality displacement over a very short time would be absorbed into a single week.

A small number of studies document a positive association between temperature and morbidity, although the evidence is once again limited to hospital admissions counts (Deschenes, 2013). The findings are generally mixed. For example, Schwartz et al. (2004) find that high temperatures increase hospital admissions for cardiac events, but that most of the effect is displacement from within a few days in the case of the highest temperatures. Cold weather, on the other hand, seemed to have a protective effect, reducing admissions in the short run without a significant bounce-back effect. Michelozzi et al. (2009) find mixed results in a study of twelve European cities, where they find a positive association between high temperatures and respiratory admissions but an insignificant, negative association between high temperatures and cardiovascular and cerebrovascular admissions. Barnett et al. (2005) found a positive association between cold periods and coronary events, with a stronger association in warm climates than in cold climates.

### III. DATA AND EMPIRICAL STRATEGY

Two primary data sources are used in our analysis: weather monitor readings from the National Oceanic Atmospheric Administration's National Climatic Data Center (NDCC) and Medicare administrative data.

The NOAA climatological data include daily monitor-level measures of high and low temperatures from 1992 to 2011 for many monitors throughout the nation.<sup>2</sup> The dataset contains 9190 monitors for 1992 and 25278 monitors for 2011. Quality flags are provided for each measurement by the monitor for each type of weather measurement; any observations with a flag are recoded as missing. Because the geographic level of the Medicare data observations is the zip code, we aggregate the daily monitor-level temperature data to the zip code-week level by averaging the daily high and low values over the week. Using data from the US Census on coordinates of geographic centroids of zip codes, each zip code is assigned a temperature reading by averaging readings from temperature monitors located within 20 miles of the zip code centroid, weighting the monitors by the inverse of the location to the centroid.<sup>3</sup> Each monitor is also weighted by the number of readings for the week (e.g. a monitor with 168 (24\*7) hourly readings

---

<sup>2</sup> Available from <http://www.ncdc.noaa.gov/>.

<sup>3</sup> This is the same method used by Beatty and Shimshack (2014), and also by Currie and Niedell (2005) for aggregating pollution monitor data.

is weighted seven times as heavily as one with just 24 observations, assuming they are equidistant from the zip code centroid). Not all zip codes are within 20 miles of a weather monitor. Zip codes are not assigned a temperature reading if they are not within 20 miles of a weather monitor. This affects a small number of zip codes, and an even smaller fraction of the population, since regions without a nearby weather monitor tend to be sparsely populated.<sup>4</sup>

Our measures of mortality, morbidity, and hospital utilization come from Medicare administrative data, including and beneficiary enrollment and claims records for the period 1999-2011. Our measures of mortality and morbidity come from the Master Beneficiary Summary File, which is an annual directory providing enrollment data on 100% of individuals eligible for Medicare in that calendar year. The directory includes beneficiary date of death (derived from Social Security Administration records), zip code of residence, date of birth, sex, and whether the beneficiary is enrolled in managed care (Medicare Advantage). The directory also includes chronic condition flags which indicate whether a beneficiary has been treated for 27 common chronic illnesses based on fee-for-service claims records (and thus only available for beneficiaries enrolled in traditional fee-for-service Medicare).<sup>5</sup> Finally, we measure hospital admissions and costs using the 100% Medicare Provider Analysis and Review (MedPAR) file, which includes one record for each inpatient and skilled nursing facility stay. Each record provides the beneficiary ID, the date of admission, and total payment due to the provider for the stay.

The primary unit of analysis for our study is a zip code week. There are roughly 32,000 zip codes in our sample, yielding over 21 million zip code week observations over our sample period 1999-2011. Table 1 shows summary statistics for the distribution of temperature weeks over this period. The first column shows the 12 temperature bins that correspond to our primary empirical specification described below. The second column shows the frequency with which a week in our data falls into each bin, based on the average weekly temperature in that week, where the weekly average is defined as the average of the average daily low and average daily high for the week. Our highest temperature bin corresponds to weeks where the average temperature is at least 95° F; on such weeks, which occur only 0.03% of the time, the median average low and average high are 83.6° F and 109.8° F, respectively.

---

<sup>4</sup> In both 1992 and 2011 just 2% of zip-code-week observations are missing a maximum temperature observation.

<sup>5</sup> The 27 chronic condition categories and claims-based algorithm for constructing the chronic condition flags are described at <https://www.cwdata.org/web/guest/condition-categories>.

Table 1

Distribution of Temperature				
Median Weekly				
Weekly Average Temperature (F)	Proportion of Weeks	Average Temperature	Median Weekly Average Low	Median Weekly Average High
<20	3.65%	14.8	4.7	24.2
20-30	6.96%	25.9	16.9	34.5
30-40	11.87%	35.5	26.0	44.7
40-50	16.01%	45.2	34.2	56.0
50-60	17.60%	55.1	43.3	66.9
60-70	19.57%	65.3	53.5	77.0
70-75	10.36%	72.4	61.4	83.5
75-80	8.06%	77.3	66.7	88.0
80-85	4.74%	81.9	71.7	92.4
85-90	1.00%	86.4	74.6	98.9
90-95	0.16%	91.7	78.8	104.9
95+	0.03%	96.4	83.6	109.8
Total	100%	56.6	44.7	68.5

Note: Weekly average temperature defined as the average of the weekly average low and weekly average high.

Table 2 shows the raw average weekly mortality rate for weeks in each temperature bin. Column 1 shows that mortality rates are highest in cold weeks. For example, in a week with average temperature between 20-30° F, there is an average of 112.2 deaths per 100,000 beneficiaries. Warmer weeks experience lower mortality rates. In fact, the lowest mortality rate occurs during weeks in the hottest temperature bin where average temperatures are at least 95° F. A naïve interpretation would be that swapping cooler weeks for very hot weeks lowers mortality rates. Yet this conclusion could be flawed for two important reasons. The first is that hot weeks tend to occur during the summer, potentially confounding the temperature effect with seasonality. The second is that not all regions are equally likely to experience hot weeks, and the population of individuals residing in regions where this occurs most often may differ systematically from cooler regions. In fact, column 3 of Table 2 shows that this concern is likely to be important, as the average life expectancy (based on age and sex) of the population exposed to very cold and very hot weeks is slightly older than the average population exposed to more temperate weeks. The richness of our data provide an exceptional opportunity to address these potential confounders by controlling very flexibly for both location and seasonality, as described in more detail below.

Table 2

Weekly Average Temperature (F)	Mortality and Life Expectancy		
	Mortality Per 100,000 Beneficiaries	Average Life Expectancy Per Beneficiary	Average Life Expectancy of Decedents
<20	110.4	14.9	9.9
20-30	112.2	15.0	10.0
30-40	110.6	15.2	10.2
40-50	105.7	15.3	10.4
50-60	100.4	15.3	10.4
60-70	94.6	15.3	10.5
70-75	92.7	15.3	10.5
75-80	91.9	15.4	10.7
80-85	91.9	15.3	10.7
85-90	91.4	15.3	10.8
90-95	85.8	15.1	10.6
95+	85.1	14.9	10.4
Total	99.7	15.3	10.4

Note: each column shows averages over all zip code weeks in each temperature bin, weighted by the number of beneficiaries in that zip code week.

Table 3 shows summary statistics for weekly hospital admission rates and spending in our sample and separately by average weekly temperature. Overall, on an average week, over 922 per 100,000 Medicare beneficiaries are admitted to the hospital, with each stay costing an average of \$9,533. This corresponds to an average of \$87.35 in hospital spending per beneficiary per week. While hospital admissions and total spending tend to be lower on hot weeks, the same potential confounders that arise when interpreting the raw mortality-temperature correlation are present here and motivate our empirical strategy for estimating the causal relationship between temperature and mortality and hospital spending.

Table 4 shows disease prevalence and weekly onset for the 27 chronic conditions used in our analysis, which are defined based on a beneficiary's claim history. In some of our analysis, we construct a measure of health stock defined as the total number of chronic conditions which a person does not have. This measure, which we call "chronic condition freeness", ranges from 27 for individuals with no treatment history for any of the listed chronic conditions to 0 for individuals who have been treated for all listed conditions. The average chronic condition freeness in the sample is 22.3. On average in a given week, 1,570 new chronic conditions are treated per 100,000 beneficiaries.

Table 3

Hospital Admissions and Spending				
Hospital				
Weekly Average Temperature (F)	Admissions per 100,000 Beneficiaries	Total Hospital Spending per Admission	Total Hospital Spending per Beneficiary	
<20	924.9	\$ 9,134.55	\$ 83.85	
20-30	964.9	\$ 9,400.71	\$ 90.37	
30-40	965.3	\$ 9,426.79	\$ 90.43	
40-50	938.8	\$ 9,553.76	\$ 88.81	
50-60	922.5	\$ 9,751.61	\$ 89.05	
60-70	903.8	\$ 9,660.75	\$ 86.61	
70-75	905.5	\$ 9,495.39	\$ 85.62	
75-80	899.9	\$ 9,274.67	\$ 83.38	
80-85	905.7	\$ 9,244.24	\$ 83.55	
85-90	899.9	\$ 9,926.91	\$ 88.53	
90-95	811.9	\$ 10,355.51	\$ 84.19	
95+	775.7	\$ 10,358.54	\$ 80.25	
Total	922.6	\$ 9,533.32	\$ 87.35	

Note: each column shows averages over all zip code weeks in each temperature bin, weighted by the number of beneficiaries in that zip code week.

Our empirical strategy is to regress an outcome variable, such as the weekly mortality rate, on the average temperature during the week, where average temperature is captured by a series of indicator variables that capture which of several “bins” contains the week’s average temperature. The omitted category consists of weeks with average temperature in the 70 – 75 degree range, so that the coefficient on one of the average temperature bins represents the impact on the outcome variable of exchanging a week in the bin for a week in this bin. Specifically, our regressions take the following form, where an observation is a zip code ( $z$ ) in week ( $w$ ):

$$outcome_{zw} = \sum_{b=1}^{11} \beta_b tempbin_{zw}^b + L_{zw} + X_{zw} + StYr_{zw} + ZipMo_{zw} + Wk_w.$$

In this regression, the variables  $tempbin^b$  indicate into which of 11 temperature bins (defined in Table 1) zip code  $z$  falls in week  $w$  based on the current week’s temperature, where bin 70-75 is the reference (omitted) category.  $L_{zw}$  is a set temperature lags, defined using temperature bins for each of the previous 4 weeks.  $X_{zw}$ , described below, is a set of zip code week controls aimed at controlling for the health stock of the population at the start of the week.  $StYr_{zw}$  is a full

set of state-year dummies,  $ZipMo_{zw}$  is a full set of zip code-month dummies, and  $Wk_w$  is a set of 52 calendar week dummies.

Table 4

Name	Chronic Conditions	
	Fraction of Population Disease-Free	Weekly Onset Per 100,000 Beneficiaries
Acute Myocardial Infarction	0.967	18.76
Alzheimer's Disease	0.962	25.26
Alzheimer's Disease and Related Disor	0.913	53.2
Atrial Fibrillation	0.901	44.8
Cataract	0.520	131.9
Chronic Kidney Disease	0.898	62.92
Chronic Obstructive Pulmonary Disease	0.802	77.46
Heart Failure	0.796	91.53
Diabetes	0.753	69.95
Glaucoma	0.859	39.16
Hip/Pelvic Fracture	0.976	15.53
Ischemic Heart Disease	0.635	114.29
Depression	0.799	70.68
Osteoporosis	0.880	43.93
RA/OA (Rheumatoid Arthritis/Osteoar	0.641	111.42
Stroke/Transient Ischemic Attack	0.902	46.55
Female/Male Breast Cancer	0.966	9.97
Colorectal Cancer	0.977	9.45
Prostate Cancer	0.959	13.33
Lung Cancer	0.991	9.62
Endometrial Cancer	0.995	2.02
Anemia	0.671	125.61
Asthma	0.920	27.41
Hyperlipidemia	0.504	120.16
Benign Prastatic Hyperplasia	0.882	38.57
Hypertension	0.377	153.52
Acquired Hypothyroidism	0.871	43.71
Total	22.315	1570.68

As stated above, the relationship between temperature and mortality is complex, and this informs our control strategy. Within most regions, mortality follows a seasonal pattern with deaths peaking in the winter. Although the pattern differs across regions, this is true even in climates with moderate temperatures. This suggests a seasonal pattern to mortality motivated by more than just temperature variation. In addition, mortality changes over time, as, for example, influenza

prevalence changes or air conditioning technology spreads. Finally, government policies may differ from state to state, and these differences may change over time. To address these issues, we adopt a series of geographic and temporal fixed effects. Depending on the nature of the fixed effects included in the regression, we may make greater use of time-series variation (i.e., looking at changes within zip code across years) or cross-sectional variation (i.e., looking at variations within a zip-code-month) to identify our results.

Another concern that informs our empirical strategy is the relationship between the health of the population entering a particular week and mortality during that week. All else equal, if the population exposed to a hot week is in poor health, we might expect higher (or at least different) mortality than if it is in good health. Further, we might be concerned that, if there is serial correlation in temperature over time, then if this week is hot, last week was probably hot as well, and if heat kills off unhealthy people then this would lead the population entering this week to be healthier than it otherwise would. The result would be that the impact of the current hot week would seem smaller than it would, since the population today is relatively healthy.

We address these concerns by including lags of temperature in our regressions as well as controls aimed at capturing the health stock of people entering the period. Specifically, we control for the age distribution at the start of the week by including the percent of the population at each age between 1 and 120 as well as the average number of chronic conditions present in the population at the start of the period. Thus, our regressions investigate the impact on the outcome of substituting a hot or cold week for a 70 degree week, holding fixed the distribution of the health stock going into the week and controlling for geographic and temporal differences in mortality.

## IV. ANALYSIS

### A. Mortality

We begin by investigating the relationship between temperature and mortality. Figure 1 depicts the results of our main regression where we have included fixed effects for zip code-months and state-years. The identifying variation is in weekly temperature within a zip code month but across years, controlling for differences at the state-year level. Thus we make use of cross-

sectional variation within a month but also time series variation across years. The error bars on the point estimates in the figure represent the 95% confidence interval for the estimate.<sup>6</sup>





Figure 1 clearly exhibits the U-shaped pattern common in the literature. The magnitude of the estimates captures the incremental deaths during a week as compared to a week with an average temperature of 70 degrees. So, substituting a week in the highest temperature bin (95 degrees) represents about 5 additional deaths per 100,000 people. Cold weeks also exhibit increased mortality, with a cold week exhibiting around 3.33 additional deaths per 100,000 people compared to a 70 degree week. As shown in Table 2, deaths on 70-75 degree weeks average 92.7 per 100,000 people. Thus, 5 excess deaths corresponds to a roughly 5.4% increase in the probability of death for a given individual.


Figure 2 presents the same regression using zip code-week and state-year fixed effects, a specification that relies purely on time-series variation across years but within the same zip code and week-of-the-year. The results are unchanged.

<sup>6</sup> We show unclustered standard errors for the time being. Clustering standard errors at the zipcode level takes significant computational time but does not substantially change the standard errors. Clustered standard errors will be integrated into the analysis as they become available.



A third fixed-effects strategy is to employ fixed effects for the zipcode-year-month. In this case, we would be identifying the estimates using only variation within a particular month and zipcode (e.g., January, 2002 in zip code 61820). In this case, we lose substantial variation if it is the case that more of the variation in a zip code’s temperature comes across years (e.g., some years are hot, others are cold but within years adjacent weeks are similar) than within months (e.g., this is week is hot and next week is cold). Figure 3 depicts our results for this specification. Although the results are robust for high temperatures, we see a noticeable flattening of the curve for low temperatures. Cold weeks no longer exhibit excess deaths relative to 70 degree weeks. Because we suspect that this is due to the fact that controlling for the zip-year-month discards too much useful intertemporal variation in temperatures, and because the results of the first two specifications are so similar, we will adopt the first (zip-month and state-year fixed effects) as our primary specification.<sup>7</sup>

<sup>7</sup> In fact, the results remain essentially unchanged for a wide variety of fixed effects specifications as long as the specification makes some use of across-year variation in temperature within zip code.



### B. Mortality Displacement (Harvesting)

A natural concern in the case of heat or cold related deaths is that the observed increase in deaths does not represent a true increase in deaths. Rather, it represents a moving forward in time of deaths that would have occurred in a few days or weeks even in the absence of the weather event. Implicit in this view is the idea that the event attacks the frailest people in the population (see, for example, Tolemon and Barbieri, 2008), so that those who die were likely to have few life years remaining.

We investigate the possibility of this phenomenon, which is known as “mortality displacement” or “harvesting,” in two ways. First, we expand the after-event window to include not only the week of the weather event but the following three weeks as well. If the additional deaths are due to harvesting, then an increase in deaths this week should be followed by a decrease in deaths in the coming weeks. As the post-event window becomes longer, harvesting effects should become smaller. On the other hand, if we observe that the mortality effect of a temperature event increases over time, this suggests that the result is not primarily due to harvesting.

Figure 4 depicts weekly mortality for the week of the weather event as well as the following three weeks.<sup>8</sup> On hot weeks, we see that, although there were significant increases in mortality on

<sup>8</sup> In these regressions we also control for the temperature during the three weeks after the event week in order to ensure that any increases in mortality are due to the weather in the event week and not on the fact that, for example, hot weeks tend to be followed by other hot weeks.

weeks with temperatures above 70 degrees (see Figure 1), when we look over a four week period mortality in hotter weeks is no longer statistically different from mortality in a 70 degree week. Although the standard errors of the estimates grow large for weeks above 90 degrees, at least in the case of 75 and 80 degree weeks we see a fairly precisely estimated zero, suggesting that the additional mortality during these weeks was offset by reduced mortality in the next few weeks, i.e., there was harvesting. In the case of cold weeks, we find that that the four-week mortality impact is significantly larger than the one week impact, suggesting once again that the impact is not only due to harvesting. In fact, we find that the four-week mortality impact of a 10 degree week (relative to a 70 degree week) is about 4.5 times larger than the one-week impact.<sup>9</sup>





Figure 5 depicts mortality in the three weeks following the event week, but excluding the week of the event, which clarifies the comparison of Figures 3 and 4.<sup>10</sup> In cold weeks, we see a clear pattern of additional deaths in the three weeks after the event week. In hot weeks, relative to a 70 degree week we see slightly fewer deaths during the next three weeks at temperatures in the range of 75 – 85 degrees. This “bounce back effect” suggest some near-term harvesting: some of the additional deaths on an 80 degree week are offset by fewer deaths in the next few weeks. In the case of an 80 degree week, we are unable to reject the hypothesis that all additional deaths are

<sup>9</sup> One caveat to this analysis is that, to the extent that mortality displacement takes place over periods shorter than a week (e.g., one or two days), our weekly specification will not pick this up.

<sup>10</sup> That is, if the event week is week 0, Figure 3 shows mortality in week 0, Figure 4 shows mortality in weeks 0, 1, 2 and 3, and Figure 5 shows mortality in weeks 1, 2 and 3 only.

the result of mortality displacement. In the case of very high temperatures, the estimates become imprecise due to the relatively small number of observations in these temperature bins and nonlinearity of the effect within the bin, making statistical inference difficult.



Although the absence of a “bounce-back” effect in the weeks following a temperature event is suggestive of the fact that the mortality effects are not due to harvesting, the richness of the Medicare data allows us to take a more direct approach to the question by looking at the health status of those who are killed due to a weather event. In particular, while mortality displacement is certainly of interest, it is part of a broader set of questions regarding which people are killed by extreme temperatures. Mortality displacement considers whether the weather event killed people who were likely to die shortly, anyway. In other words, mortality displacement addresses the question of whether naïve estimates of the cost of an event may be too high because although it induced a short term increase in deaths, the people who died had fewer life years remaining than a typical person. The richness of the Medicare data allows us to address this question directly. Using data on the age and health status of those alive at the start of each week and those who die as the result of a temperature event, we can investigate whether the weather event appears to kill representative individuals, or whether those who are killed tend to possess a lower health stock.

To analyze this question, we construct the following measure. Consider a particular zip-code-week observation. Let  $H_i$  be a measure of the health stock for individual  $i$ , who is alive in this zip code at the start of this week. Let  $L$  be the set of individuals who are alive at the start of the week, and let  $D$  be the set of individuals who die during the week. The average health stock

per death is  $\sum_{i \in D} H_i / |D|$ , and the average health stock per individual alive at the start of the period is  $\sum_{i \in L} H_i / |L|$ . If the death process randomly selects an individual from the population to kill, then the health status of those who die during the period should be the same as the health status of those who were alive at the start of the period.

$$\frac{\sum_{i \in D} H_i}{|D|} = \frac{\sum_{i \in L} H_i}{|L|}.$$

Simple manipulation yields the following:

$$\frac{\sum_{i \in D} H_i}{|D|} \begin{matrix} \left[ \geq \right] \\ \left[ = \right] \\ \left[ \leq \right] \end{matrix} \frac{\sum_{i \in L} H_i}{|L|} \Leftrightarrow \frac{\sum_{i \in D} H_i}{\sum_{i \in L} H_i} \begin{matrix} \left[ \geq \right] \\ \left[ = \right] \\ \left[ \leq \right] \end{matrix} \frac{|D|}{|L|}$$

That is, if the death process selects people whose average health is the same as the population, the death rate and the fraction of health status lost, should be the same. If the death process selects relatively healthy people, then the fraction of health status lost should be greater than the mortality rate, and if the death process selects relatively sick people, the fraction of health status lost should be less than the mortality rate.<sup>11</sup> In the next part of our investigation, we consider the ratio  $\sum_{i \in D} H_i / \sum_{i \in P} H_i$  for the cases where the health measure is either life years remaining or degree of chronic condition freeness, i.e., the number of chronic conditions a person does not have.

An alternative interpretation of the ratio  $\sum_{i \in D} H_i / \sum_{i \in P} H_i$  is that it represents the health-weighted mortality rate. That is, if  $H = \sum_{i \in P} H_i$ , and each individual is given weight  $H_i/H$ , then  $\sum_{i \in D} H_i / \sum_{i \in P} H_i$  is just the death rate in the zip code computed using these weights. Zip codes where the typical individual who dies is less healthy than the population as a whole will have the health-weighted mortality rate be less than the raw mortality rate and vice versa.

In the results depicted in Figure 6, we compute the life years remaining for everyone alive at the start of the week based on age and gender using life tables available from the Social Security

---

<sup>11</sup> Note that although higher health status is good, if the death process is selecting individuals who are healthier than the population, this implies that raw mortality counts underestimate the true cost of the deaths, since the individuals who die had more than average life years remaining.

Administration.<sup>12</sup> We then compute the life years lost in each week as a fraction of total life years remaining at the start of the week and use it as an outcome variable in our usual regressions, with a 70 degree week once again as the reference category. The coefficient on a temperature bin therefore represents the relative increase in the fraction of life years lost compared to a 70 degree week. Thus, if the fraction of life years lost in a temperature bin lies below the mortality rate in that bin, the people who died in that week were less healthy on average than those who die on a typical 70 degree week.




Figure 6 shows the fraction of life years lost (blue/dashed) in addition to the mortality rate (red).<sup>13</sup> Before comparing the two, note that the life years lost line also exhibits a U-shape and that the estimated fraction of life years lost at both high and low temperatures is significantly higher than at 70 degrees. This suggests that those who die at extreme temperatures are younger than those who die during the 70 degree reference week (since they have more remaining life years to lose). This is somewhat in contrast to the conventional wisdom, which suggests that extreme temperatures prey on the most fragile people in the population. In truth, these people were likely to die whether or not the extreme temperature week had occurred. It is the slightly more healthy

<sup>12</sup>See <http://www.ssa.gov/oact/STATS/table4c6.html>, accessed 6/16/2014.

<sup>13</sup> An alternative interpretation of the fraction of life years lost curve is that it is the mortality rate weighted by the life years remaining of each individual. So, it represents the life-years-remaining weighted number of people who die divided by the life-years-remaining weighted number of people alive at the start of the period. Under this interpretation, if the people who die have, on average, fewer life years remaining than the people alive at the start of the period, the weighted average mortality rate will be smaller than the raw mortality rate.

people, who would have survived in temperate weather but cannot withstand extremes, who are the marginal victims of extreme temperature events.

Comparing the mortality and fraction of life years lost lines in Figure 6, we see that the fraction of life years lost lies below the mortality curve in most temperature bins, suggesting that those who die due to hot or cold weeks have fewer life years remaining than the population as a whole. Hence, in combination with the observation, above, we see that the marginal deaths in hot or cold weeks are people who are sicker than the population as a whole, but healthier than the typical person who dies in a 70 degree week. For hot weeks, although the point estimate for life years lost lies below the point estimate for mortality rate, they are unlikely to be statistically significant. For moderate temperatures between 40 and 80 degrees, the point estimates are quite close, suggesting that the mortality process is selecting individuals whose health is similar to that of the overall population. In cold weeks, the difference between the estimates becomes larger and the standard errors of the estimates becomes smaller, suggesting that in cold weeks the people who die are systematically less healthy than the population as a whole. Consequently, in this case, raw mortality counts will tend to overstate the social cost of increased mortality.

Repeating the analysis for chronic conditions yields slightly different results, as shown in Figure 7. In order to represent the mortality rate and chronic condition prevalence in a comparable way, we consider in Figure 7 “chronic condition freeness,” i.e., the number of the 27 conditions in the Medicare data that each person does not have, rather than prevalence. This results in higher values representing healthier people, which corresponds to the way we studied life years lost above. The dashed blue line represents the total number of chronic conditions not present in the people who die, divided by the total number of chronic conditions not present in the population as a whole. As above, if this curve aligns with the mortality curve, it suggests that the people who die are similar in terms of chronic conditions to the population as a whole, while if the fraction of chronic conditions lost curve lies below the mortality curve, it suggests that those who die are less healthy than the population as a whole.

For high temperature weeks, we see that the chronic conditions curve and mortality curves align almost perfectly, suggesting that the additional people killed during hot weeks are similar in terms of health (as measured by chronic condition prevalence) to the population as a whole. This differs somewhat from what we saw when looking at life years lost in Figure 6, where we found that those killed at high temperatures were slightly older than the population as a whole. Although


the differences are not statistically significant, if one takes the point estimates seriously then those killed during hot weeks are slightly older than the population as a whole but similar in terms of chronic conditions. Again, as the chronic conditions curve is U-shaped, we see that those killed at extreme temperatures are slightly healthier than those killed during a 70 degree reference week.

At low temperatures, the chronic conditions curve is slightly below the mortality curve, but the differences seem less pronounced than they did in Figure 6. Consequently, the point estimates suggest an overall picture where those killed in cold weeks are slightly older and slightly less healthy than the population as a whole (but younger and healthier than the typical person killed in a 70 degree week).


Figures 8 and 9 provide a more direct comparison of how those who die due to a temperature event compare to people who die in a 70 degree reference week. Figure 8 plots the life years lost per decedent. Relative to a person who dies in a 70 degree week, we now see that the life years lost per decedent (i.e., average remaining life of those who die) in cold weeks is indistinguishable from the average life years of a person who dies in a 70 degree week. This suggests that the death process in cold weeks selects individuals in more-or-less the same way as the death process in 70 degree weeks. In hot weeks, on the other hand, we see some evidence that the death process selects slightly healthier individuals in hotter weeks (in particular, in the 85 degree bin) than it does in 70 degree weeks. Figure 9 shows chronic conditions per decedent. Here we see that in both hot and cold weeks, individuals who die are healthier than people who die in the reference week in the sense that they have fewer chronic conditions.<sup>14</sup>


---

<sup>14</sup> The sample in Figures 7 and 8 differs from the sample in Figure 6, since life years per decedent and chronic conditions per decedent are only defined for zip-code-weeks where someone actually dies. Consequently, Figures 7 and 8 are based on regressions that drop a portion of zip codes, generally the smaller ones.



Figures 6 - 9 paint a consistent picture. Those who die during extreme temperature events tend to be less healthy than the population as a whole, but healthier than the typical person who dies in a more temperate week.





### C. Impact on Health Stock of Survivors

Although mortality is a natural measure of the impact of a temperature event, it is not the only measure of interest. In addition to increasing mortality, temperature extremes may also harm those who do survive the event in the short term. If for example, the temperature event causes an increase in heart attacks, those who are not killed by the attack immediately might be expected to live fewer years than they would if they had not had their first heart attack.<sup>15</sup>

Table 5 depicts event-week onset of various chronic conditions for the various temperature bins. Interesting patterns emerge. Relative to a 70 degree week, we see that cold weeks are associated with significantly less onset of chronic conditions. There are several possible explanations for this finding. The first is that cold weather has a “protective effect,” and while there are significantly more deaths during cold weeks, there is less onset of new chronic diseases than during warmer weeks. Another possible explanation arises from the fact that individuals are coded as having a chronic disease based on the first or second time a Medicare claim is processed containing particular diagnosis codes. Thus, a visit to the doctor or emergency department is necessary to trigger onset of a chronic condition in our data.

If the latter is an important cause of the reduction of new chronic condition onset in cold weeks, then we should see smaller reductions in emergent conditions and larger reductions in non-

<sup>15</sup> As in the case of mortality, displacement might also be a concern here. That is, although a hot day might cause an increase in heart attacks, some of those heart attacks would occur in people who would have had a heart attack in a few days or weeks, anyway.

emergent conditions. The data support this hypothesis. For example, we see no effect on onset of new heart attacks (AMI) or hip fractures in cold weeks, consistent with the idea that patients are unable to delay going to the doctor until the weather gets better in these cases. On the other hand, we see larger reductions in chronic but not emergent conditions, such as Ischemic Heart Disease, Congestive Heart Failure, Chronic Obstructive Pulmonary Disease (COPD), Breast Cancer and Asthma.

The pattern of effects in hot weeks is substantially different. Relative to a 70 degree week, we see a small dip in chronic condition onset for temperatures between 75 and 85, followed by an increase in chronic condition onset at the highest temperatures. For the emergent conditions, Heart Attack and Hip Fracture, we see little impact on onset of new chronic conditions. Thus, for example, while hot weeks cause onset of chronic conditions, they do not appear to cause heart attacks in people who have not previously had a heart attack. On the other hand, we find evidence of new onset at high temperatures in the cases of Ischemic Heart Disease, Congestive Heart Failure, COPD and Asthma.

In the case of these latter conditions, which generally build up over time, it is unlikely that a single hot week changes an individual's health status substantially. For example, Ischemic Heart Disease (IHD) is caused by blockages of the arteries that supply blood and oxygen to the heart, and these deposits build up over time. Thus, while a particularly hot week may cause a person to manifest symptoms of IHD it is unlikely that hot weather causes IHD in the same sense that it might cause a heart attack.

#### D. Impact on Hospital Admissions and Costs

In this section, we turn to the question of how temperature events affect hospital admissions and the cost of care. This is interesting for several reasons. First, hospital admissions give another measure of the impact of weather events on health that differs from mortality. Second, hospital admissions and spending give us a sense of the cost of temperature events beyond life years lost.

Table 5: Onset of Chronic Conditions





Figure 10 shows the fraction of the population alive at the start of the week that is admitted to the hospital during the event week (blue solid line) and during the event week and the following 3 weeks (orange, dashed line), relative to admissions during a 70 degree week. Beginning with 1 week admissions, we see that, in contrast to the U-shaped pattern of mortality, there is monotonically increasing relationship. In cold weeks there are systematically fewer hospital admissions than on a 70 degree week, though more people die in cold weeks than 70 degree weeks. This could be caused either by a decrease in discretionary admissions that more than offsets an increase in non-discretionary admissions caused by the cold, or by the fact that those who are killed by cold weather do not make it to the hospital before they die. We cannot directly distinguish these two effects, although the fact that we do not see a decrease in onset of the AMI chronic disease suggests the decline might be driven by a decrease in discretionary admissions. During hot weeks, we see an increase in hospital admissions. Given that we also saw an increase in onset of chronic conditions at very high temperatures, this is likely driven by marginal admissions caused by the hot weather.

Comparing the 1 week admissions with 4 week admissions (event week and the three following weeks), we see that the same basic pattern emerges during hot weeks. Hot weeks cause a spike in immediate admissions and have little effect on admissions in subsequent weeks. During cold weeks, on the other hand, we see a starkly different story. While hospital admissions are significantly lower during the event week, by the end of the third week after the event week, we see an increase in overall admissions. Thus, while cold prevents individuals from going to the hospital immediately, there is an increase in admissions over the subsequent weeks. This is interesting for several reasons. First, to the extent that the literature has identified a “protective effect” of cold weather, this phenomenon appears to decrease as the length of the post-event period increases. Second, it is possible that, cold weather, by preventing individuals from going to the doctor hospital immediately, actually exacerbates illnesses and results in increased illness over the longer period.




While hospital admissions are one measure of the impact of temperature events on healthcare utilization, not all admissions are equally costly. Figure 11 depicts the relationship between weekly average temperature and total hospital spending, which both weights admissions by the intensity of utilization and gives us a policy-relevant parameter, the impact of the weather event on hospital spending.<sup>16</sup> The pattern that emerges is much like the pattern of hospital admissions. In hot weeks, hospital spending is higher than in a 70 degree week. Exchanging a 90 degree week for a 70 degree week increases hospital spending by about \$3.46 in the short run and \$2.32 in the long run. The point estimates for 4-week spending following hot weeks are slightly different, the differences are not statistically significant. However, if one were to take the point estimates seriously, it would suggest the immediate spike in spending is offset by lower spending over the next several weeks, evidence of short-term displacement in use of hospital services.

In cold weeks, on the other hand, we see an immediate reduction in hospital spending during the event week, while 4-week spending is significantly higher than the reference week. For example, exchanging a 40 degree week for a 70 degree week reduces hospital expenditure in the short run by about \$1.25 per beneficiary but increases long run expenditure by about \$2.00 per

<sup>16</sup> Total hospital spending is made up of Medicare spending, beneficiary out-of-pocket spending, and payments by other payers (principally Medigap or Medicaid in the case of dual eligible). Future work will examine the spending categories separately. To a first approximation, Medicare spending is 80% of total spending, since Medicare features 20% coinsurance over the beneficiary's deductible.

beneficiary. Once again, this suggests a strong rebound effect: the immediate reduction in hospital usage is offset by an even larger increase in hospital spending over the next three weeks.



#### E. Comparison of Mortality Cost and Health Care Cost


Two of the major health-related cost components associated with temperature events are the cost due to mortality and the additional health care costs generated by additional encounters with the health care system.<sup>17</sup> Figure 11 showed that a typical 90 degree week induces additional hospital spending of around \$3.46 in the event week and \$2.32 over the event week and the three following weeks.<sup>18</sup>

For comparison, Figure 12 shows the life years lost per beneficiary as a function of temperature. Here, we see that a typical 90 degree week eliminates around 0.00031 life years per beneficiary in the short run. Using a value of \$250,000 value per life year lost, this equates to a value of life years lost associated with a 90 degree week of around \$76 per beneficiary.<sup>19</sup> The four-week cost figure is slightly lower, around \$57.5 per beneficiary.

<sup>17</sup> Other major components, which are the subject of future work, are the additional loss of life years and additional health care spending due to the fact that the temperature event affected the future health of those who survive the initial event.

<sup>18</sup> Hospital spending is a major category of healthcare spending, but not the only one. Future work will characterize additional spending on physician services delivered outside of the hospital as well as other components of healthcare spending.

<sup>19</sup> Aldy and Viscusi (2008) compute an age-specific value of a statistical life year (VSLY) for individuals in their early 60s between \$150,000 and \$350,000, depending on the specification. We take the midpoint of this range, although other choices are possible.



Comparing these figures, we see that the additional hospital spending, while not insignificant, is much smaller than the mortality cost due to high-temperature weeks, with hospital cost being in the range of 3-6% of mortality cost depending on whether the impact is measured only for the event week or for the four week period starting with the event week.<sup>20</sup>

The story is slightly different for cold weeks. In the short run (including only the event week), a 40 degree week reduces life years by 0.000061 per beneficiary, which implies a value of \$15.3 per beneficiary. However, due to reduced hospital admissions, hospital costs actually go down by around \$1.25 per beneficiary, offsetting the short run mortality cost. However, if we expand the analysis to the four week period starting with the event week, we see that the mortality cost is an order of magnitude higher, increasing to around \$197 per beneficiary and incremental hospital costs become a positive \$2 per beneficiary.

The per-beneficiary numbers in Figures 11 and 12 also allow us to project the cost of exposing large areas to increased temperatures. For example, the state of Illinois had approximately 1.9 million Medicare beneficiaries in 2012.<sup>21</sup> If a heat wave were to hit the a representative state with this population, converting what would otherwise be a 70 degree week to a 90 degree week, this would generate approximately \$145 million in mortality costs (value of life years lost) and \$4.4 million in additional hospital spending. Thus, these two categories (which are important but not the only costs of high temperatures) would generate a cost of almost \$150 million

<sup>20</sup> By comparison, the ratio of the monetized benefits of mortality reductions from the Clean Air Act from 1970-1990 to all other benefits is four-to-one (EPA 1997, Table ES-4).

<sup>21</sup> Source: <http://kff.org/medicare/state-indicator/total-medicare-beneficiaries/>, accessed June 13, 2014.

in the short run. Using the four-week cost estimates, these numbers become \$109 million and \$6.6 million, for a total of around \$115 million. Again, while the cost components we have measured are important, they are by no means the only costs. In addition to the factors discussed above, it is important to emphasize that these cost estimates capture only the effect on the Medicare population, while the heat wave is sure to impact the non-Medicare population as well.

#### IV: DISCUSSION AND CONCLUSION

This paper uses the extraordinary detail of Medicare administrative data to investigate the relationship between temperature and human health. While some of our analysis confirms the findings of the literature, in particular a U-shaped relationship between temperature and mortality, we are able to go beyond the literature in several important respects. In particular, the Medicare data allow us to conduct a more detailed investigation of who is killed by extreme weather, where we find that the marginal person killed in a high- or low-temperature week is healthier than a typical person who dies during a temperate week, but less healthy than the population as a whole. This paints a somewhat more nuanced story than the conventional wisdom that extreme events kill the most vulnerable people. In truth, those people were likely to die anyway, even if the extreme weather event had not occurred. It is those who are slightly less healthy than average, but not in immediate danger of dying, who comprise the marginal deaths due to an extreme heat event. Because of this, those deaths are more costly in terms of life years lost than might otherwise be expected.

Another major way in which the Medicare data allows us to go beyond previous work is by giving us a detailed picture of health expenditures induced by heat-related events. Although there is still much to be done on this front, our initial investigation into hospital spending indicates that while hot weeks do induce additional hospital spending, these costs are much smaller in magnitude than the direct cost of life years lost (using conventional values for the value of a life year lost). In future work, we will expand this analysis to include other categories of Medicare spending, including physician and prescription drug spending. Of course, one limitation of the Medicare data is that it gives us healthcare spending for the fee-for-service Medicare population only. While the fee-for-service Medicare population is an important constituency, both because

of its size and its vulnerability to temperature events, we are unable to directly look at the Medicare Advantage (i.e., Medicare managed care) population or at younger individuals.

This paper also moves beyond the existing literature in that it begins to investigate the impact of extreme weather events on those who are not immediately killed by the event. Our investigation of onset of chronic conditions shows that extreme temperature events do affect those who survive the initial incident. To the extent that a person who has, for example, a heart attack as a result of a hot week has a reduced life expectancy, the loss of life years this entails should be factored into the cost of the hot week. In future work, we will build age and chronic condition specific life tables that will allow us to estimate the life years lost due to the change in health status of those who survive following an extreme weather event. Since mortality rates are rather low, even small impacts on the health of those who do not die can generate large costs by virtue of the fact that the group of survivors is so large.

## References

- Aldy, Joseph E., and W. Kip Viscusi. "Adjusting the value of a statistical life for age and cohort effects." *The Review of Economics and Statistics* 90, no. 3 (2008): 573-581.
- Baccini, Michela, Annibale Biggeri, Gabriele Accetta, Tom Kosatsky, Klea Katsouyanni, Antonis Analitis, H. Ross Anderson et al. "Heat effects on mortality in 15 European cities." *Epidemiology* 19, no. 5 (2008): 711-719.
- Barnett, Adrian G., Annette J. Dobson, Patrick McElduff, Veikko Salomaa, Kari Kuulasmaa, and Susana Sans. "Cold periods and coronary events: an analysis of populations worldwide." *Journal of Epidemiology and Community Health* 59, no. 7 (2005): 551-557.
- Barreca, Alan I. "Climate change, humidity, and mortality in the United States." *Journal of Environmental Economics and Management* 63, no. 1 (2012): 19-34.
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro. *Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the 20th Century*. No. w18692. National Bureau of Economic Research, 2013.
- Beatty, Timothy KM, and Jay P. Shimshack. "Air pollution and children's respiratory health: A cohort analysis." *Journal of Environmental Economics and Management* 67, no. 1 (2014): 39-57.
- Braga, Alféio Luís Ferreira, Antonella Zanobetti, and Joel Schwartz. "The time course of weather-related deaths." *Epidemiology* 12, no. 6 (2001): 662-667.
- Curriero, Frank C., 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, no. 1 (2002): 80-87.
- Deschenes, Olivier. "Temperature, human health, and adaptation: A review of the empirical literature." *Energy Economics* (2013).
- Deschênes, Olivier, and Michael Greenstone. "Climate change, mortality, and adaptation: evidence from annual fluctuations in weather in the US." (2007).
- Deschenes, Olivier, and Enrico Moretti. "Extreme weather events, mortality, and migration." *The Review of Economics and Statistics* 91, no. 4 (2009): 659-681.
- Environmental Protection Agency. "Benefits and Costs of the Clean Air Act – Retrospective Study." (1997). Available at <http://www.epa.gov/air/sect812/copy.html>.

- Hajat, Shakoor, Ben G. Armstrong, Nelson Gouveia, and Paul Wilkinson. "Mortality displacement of heat-related deaths: a comparison of Delhi, Sao Paulo, and London." *Epidemiology* 16, no. 5 (2005): 613-620.
- Huynen, MMTE-Martens, Pim Martens, Dieneke Schram, Matty P. Weijnenberg, and Anton E. Kunst. "The impact of heat waves and cold spells on mortality rates in the Dutch population." *Environmental health perspectives* 109, no. 5 (2001): 463.
- Keatinge, W. R., and G. C. Donaldson. "Cold exposure and winter mortality from ischaemic heart disease cerebrovascular disease respiratory disease and all causes in warm and cold regions of Europe." *Lancet* 349, no. 9062 (1997): 1341-6.
- Kysely, Jan. "Mortality and displaced mortality during heat waves in the Czech Republic." *International Journal of Biometeorology* 49, no. 2 (2004): 91-97.
- Michelozzi, Paola, Gabriele Accetta, Manuela De Sario, Daniela D'Ippoliti, Claudia Marino, Michela Baccini, Annibale Biggeri et al. "High temperature and hospitalizations for cardiovascular and respiratory causes in 12 European cities." *American journal of respiratory and critical care medicine* 179, no. 5 (2009): 383-389.
- Moller, Peter. "Excess winter mortality, wood fires and the uncertainties associated with air pollutants." *NZ Med J* 124, no. 1330 (2011): 58-65.
- Oudin Åström, Daniel, Forsberg Bertil, and Rocklöv Joacim. "Heat wave impact on morbidity and mortality in the elderly population: a review of recent studies." *Maturitas* 69, no. 2 (2011): 99-105.
- Pan, Wen-Harn, Lung-An Li, and Ming-Jan Tsai. "Temperature extremes and mortality from coronary heart disease and cerebral infarction in elderly Chinese." *The Lancet* 345, no. 8946 (1995): 353-355.
- Peng, Roger D., Francesca Dominici, Roberto Pastor-Barriuso, Scott L. Zeger, and Jonathan M. Samet. "Seasonal analyses of air pollution and mortality in 100 US cities." *American journal of epidemiology* 161, no. 6 (2005): 585-594.
- Ren, Cizao, Gail M. Williams, and Shilu Tong. "Does particulate matter modify the association between temperature and cardiorespiratory diseases?." *Environmental health perspectives* 114, no. 11 (2006): 1690.
- Schwartz, Joel, Jonathan M. Samet, and Jonathan A. Patz. "Hospital admissions for heart disease: the effects of temperature and humidity." *Epidemiology* 15, no. 6 (2004): 755-761.
- Toulemon, Laurent, and Magali Barbieri. "The mortality impact of the August 2003 heat wave in France: investigating the 'harvesting' effect and other long-term consequences." *Population studies* 62, no. 1 (2008): 39-53.



## APPENDIX A: CONCEPTUAL FRAMEWORK

### A1: Contribution of the Paper

In evaluating the desirability of a policy, policymakers must compute the costs and benefits of the policy. For a concrete example, consider a policy aimed at reducing extreme heat events or pollution through reducing carbon and particulate matter emissions through regulating power plants. To properly evaluate such a policy, the regulator must compute the cost associated with adverse events such as greater prevalence of hot or high-pollution weeks. The costs of the adverse event can be decomposed into the health impact, the increased costs incurred as a result of the event, and the decreased consumption utility of individuals who live through the event. That is:

$$\text{Cost of Adverse Event} = \Delta\text{LY} * \text{V} + \Delta\text{Cost} + \Delta\text{Utility},$$

where LY is life years remaining in the population and V is the value assigned to a life year. This expression can be further decomposed:

$$\text{CAE} = [\Delta\text{LY|D} + \Delta\text{LY|ND}] * \text{V} + \Delta\text{Healthcare Cost} + \Delta\text{Other Cost} + \Delta\text{Utility},$$

where LY|D is the life years of people who die due to the event and LY|ND is the life years of people who do not die.

The majority of work on the impact of temperature or pollution events has focused on raw mortality counts. However, the cost of a death depends on how long that person would have lived in the absence of the event. The mortality displacement literature is a crude attempt to address this issue by attempting to determine whether the incremental deaths due to the event would have occurred in a few days or weeks regardless of the event, in which case they should be ascribed a much lower value than the death of a person who would otherwise have lived many years. Other work has looked at whether the mortality effects of events are concentrated in “vulnerable populations” such as the sick or elderly. In this paper, we address this issue by explicitly estimating the expected life years lost due to the event, using detailed data on the age and gender distribution of the individuals alive at the start of the event week. Future work will integrate information on health history (e.g., chronic conditions, recent use of health services) into this analysis.

This second component of cost is the impact on the life years remaining of those who do not die as a result of the event, but whose health is affected by it nonetheless. For example, if a weather event causes new heart attacks, which in turn damages heart muscle and reduces life expectancy even in those who survive the initial heart attack, this is a cost of the event that should

factor into regulators' decisions, but is often ignored. Because the Medicare data give us weekly information on prevalence of chronic conditions, we are able to identify whether the event increases onset of chronic conditions such as heart attack or ischemic heart disease that might affect an individual's life expectancy even if it does not immediately kill him. While at this point we focus on onset counts, future work will use the Medicare data to develop explicit models of life expectancy based on medical history as a tool for estimating the life expectancy impact of weather events on those who survive the initial event.

A third component of cost that should be of interest to regulators is the change in expenditure induced by the event. Some of these expenditures may result from individual decisions such as whether to run an air conditioner. Of particular interest to the government is the additional expenditure by programs such as Medicare to care for those whose health is affected by the event. Previous work on this topic has focused on hospital admissions counts. However, not all hospital admissions are equal. Using the Medicare data, we are able to compute the additional hospital expenditure induced by an event.

## A2: Harvesting, Selection and the Death Process

Whether the incremental deaths caused by temperature events are due to harvesting, i.e., short term displacement of deaths that would have occurred in a few days or weeks even in the absence of the temperature event, is a special case of a more general group of questions dealing with how particular events change the death process. To help make ideas clear, we can suppose that on any given day, there is a distribution of fragility (e.g., ill health) in the population and a death process that assigns a probability of dying to each level of fragility. An example is shown in Figure A1.

A temperature event, such as a high temperature week, can have several effects on who dies. First, the probability of death might be different at low temperatures and high temperatures. Figure A1 depicts probability-of-death curves at low (blue, dashed) and high (red, dotted) temperatures. In each case, the probability-of-death curves feature a steep segment for highly fragile people, dividing the population into low fragility people, who face a low probability of death, and high-fragility people, who face a significant probability of death. Drawn this way, the probability-of-death curve features a "cut-off" below which death is a fairly rare event. A temperature event such as a hot week can both move the steep segment – i.e., the cut-off level –

to the left, giving people a significant probability of death who previously had a rather low probability of dying. However, it can also move the segment farthest to the right upward, corresponding to an increase in the probability of death for those already in the frailest of health.

Either of these effects would increase the death rate, but they have somewhat different interpretations. The case where the temperature event increases the probability of death for those who already have a high probability of death seems to correspond to what most people think about when considering harvesting (see, for example, Toulemon and Barbieri). These deaths have a relatively low cost in terms of life years lost since high-frailty people have fewer life years remaining, anyway.

Deaths that arise due to the other kind of shift, which increases the likelihood of death for those who previously had a low probability of death, involve higher cost in terms of life years lost since these people generally are in better health. The type of deaths we find in our analysis more closely correspond to this sort of death, since we find that the incremental deaths induced by temperature events involve people who are sicker than the population as a whole, but healthier than the people who die during more temperate temperatures.




Figure A1: Fragility and the Probability of Death

A high-temperature event can also affect those who die by altering the distribution of fragility going into the next week. Here, there are two possible effects. First, since the death process works primarily on those in the worst health, the fragility distribution of those who survive the temperature event is likely to be less fragile (i.e., healthier) than the distribution was before the health event. However, there is also a countervailing factor, and that is the fact that high-heat weeks can make previously healthy people sick, and this will tend to worsen the fragility distribution the following week. Since both factors are likely to be at play, and the relative strength of the two is likely to depend on the fragility distribution going into the week and the exact nature of the temperature shock, we address this issue by controlling for the fragility of the population by including the distribution of age in the population and average number of chronic conditions in the zip code at the start of each week.