The Dynamic Relationship Between Temperature and Morbidity *

Corey White
University of California, Santa Barbara
cdwhite@umail.ucsb.edu

This version: March 1, 2017

Abstract

This paper examines the relationship between temperature and hospital usage with a focus on the role of behavioral responses to temperature. I use high-frequency data on the near universe of hospital and emergency department (ED) visits in California between 2005-2014 to estimate the effects of temperature on hospital usage patterns. I find that a day with mean temperature under 40°F leads to a 6.1% *decrease* in ED visits on the day of the event, but that total net visits *increase* by approximately 11.0% above the daily mean after accounting for visits in the weeks that follow. Additionally, I find that a day over 80°F is associated with a same-day increase in ED visits of 3.5%, and a total net increase of 5.1%. For both cold and hot temperatures, I provide evidence of the mechanisms – whether biological or behavioral – that explain these patterns.

JEL: I10, I12, I18, Q50, Q51, Q54
Keywords: Health, Weather, Temperature, Climate

*This paper has benefitted greatly from numerous conversations with Olivier Deschênes, Peter Kuhn, Heather Royer and Maya Rossin-Slater, comments from Matthew Neidell and two anonymous referees, and from discussions with Jennifer Purcell Deacon, Robert Deacon, Kyle Meng, participants at the UCSB Environmental Lunch Seminar Series, the Broom Center for Demography Seminar Series, All California Labor Economics Conference Poster Session, the Association of Environmental and Resource Economists Summer Conference and the Western Economic Association International Annual Meetings. I would like to thank the Broom Center for Demography and the UCSB Department of Economics for providing funding for this research. Finally, I would like to thank Betty Henderson-Sparks for facilitating access to restricted data through California's Office of Statewide Planning and Development. All errors are my own.

1 Introduction

The threat of climate change has spurred a large literature in recent years that has shown a remarkable sensitivity of numerous outcomes to weather and climate (see Dell et al. (2014) for a review). One of the primary outcomes of interest in this literature is health, and several studies document the effects of weather on mortality. The effects of weather on health outcomes that do not result in mortality (i.e., morbidity) have for the most part been overlooked. Analysis of such lower level health outcomes can not only provide a more complete picture of the overall effects of weather on health, but can provide an opportunity to develop a deeper understanding of the channels through which these impacts operate. In addition to the biological relationship between weather exposure and health, it is possible that the behavioral responses in which individuals engage act as an important mechanism in this relationship. An understanding of the specific channels through which health is affected by weather is important in a variety of settings. For instance, such an understanding can guide actions toward mitigating the damages associated with climate change. Alternatively, if such behavioral responses are costly, failing to account for them would lead to an underestimate of the total cost of a given weather event.

In this paper, I estimate the relationship between temperature and hospital usage in California over the period 2005 to 2014; using daily data I focus on the dynamics of this relationship in an attempt to explore the role of behavioral responses. Specifically, I exploit random daily variation in weather within zip-codes in California to document how emergency department (ED) visit rates respond to temperature on the day of the event as well as the period that follows, controlling for seasonal factors using zip-by-week-of-year fixed effects and annual factors using county-by-year fixed effects.

I find that a day with mean temperature under 40°F is associated with a 6.1% *decrease* in ED visits on the day of the event, relative to a day in the 60-65°F range (henceforth, the contemporaneous effect). In the 30 days that follow, there are a series of small increases that result in a total net *increase* of 11.0% above the mean daily visit rate (henceforth, the cumulative effect).[1]

Since it is unlikely that exposure to cold temperatures improves health through a biological mechanism, the contemporaneous decrease in ED visits is likely driven by behavioral factors. Possible behavioral drivers include a decreased willingness to seek treatment for a medical condition or behavioral changes that lead to a decrease in the actual incidence of disease. I show that the contemporaneous drop is no smaller for diseases that are not deferrable in nature (e.g., heart attacks), suggesting that this decrease is at least partially due

[1] The "cumulative effect" is inclusive of the "contemporaneous effect".

to a decrease in the incidence of disease. While the contemporaneous effect of cold weather is common to nearly all disease types, the cumulative increase is driven primarily by visits for respiratory diseases and other disease categories with a high proportion of communicable diseases.

For hot temperatures, I find that a day over 80°F is associated with a 3.5% contemporaneous increase in ED visits; through increases in visits in each of the days that immediately follow the weather event, this effect grows to a 31-day cumulative effect of 5.1%. The delayed visits may be attributable to either delayed symptoms or to deferred treatment. I find that the delayed visits are driven by conditions that are more deferrable in nature, providing suggestive evidence that it is not the symptoms that are delayed, but the decision to seek treatment.

For both cold and hot temperatures, the deleterious impacts are driven by children. This is a significant finding as studies of temperature and mortality have found that the elderly are more sensitive to weather-related mortality. The implication is that studies focusing exclusively on mortality do not capture the potentially large impacts experienced by groups who are unlikely to die from weather-related illnesses.

I find that these weather events result in substantial increases in health care costs.[2] For every 100,000 persons exposed to a day under 40°F or a day over 80°F, I find that the cost of providing hospital services increases by $12,176 or $7,994, respectively.[3] Scaled to large population, these costs are significant: if the population of the greater Los Angeles metropolitan area (18.7 million) were exposed to single day over 80°F, these estimates imply a $1.49 million increase in hospital costs. Furthermore, these estimates only reflect the cost of providing hospital services and should thus be interpreted as lower bound estimates of the total cost of temperature-induced morbidity.

The relationship between weather and health has been studied extensively by both epidemiologists (see Basu and Samet (2002); Ye et al. (2012) for reviews) and economists (Deschênes and Moretti, 2009; Deschênes and Greenstone, 2011; Barreca, 2012; Barreca et al., 2016; Isen et al., 2015; Karlsson and Ziebarth, 2016). The literature in economics has focused on adaptive responses to weather and climate in both the short- and long-term. In the long term, Deschênes and Moretti (2009) argue that migration from colder to warmer climates is responsible for a portion of the increase in life expectancy since 1970. In the shorter term, a number of papers show a substantial increase in residential electricity usage, likely due to increases in the usage of air conditioning, when temperatures rise (Deschênes and Greenstone,

[2]Note that the measure of heath care costs does not represent patient charges, but the cost of providing health services.

[3]Note that only the estimates for hot temperatures are statistically different from zero.

2011; Barreca, 2012; Barreca et al., 2016). Taking this a step further, Barreca et al. (2016) show that the adoption of air conditioning is likely responsible for much of the improvement in the heat-mortality relationship over the past century in the U.S.

A small literature studies the relationship between temperature and behavior more generally. Graff-Zivin and Neidell (2014) use the American Time-Use Survey to show a decrease in hours devoted to labor for weather-exposed workers during hot temperatures, and additionally demonstrate a sizeable shift from time spent outdoors to time spent indoors in response to both hot and cold temperatures. These are examples of behavioral responses in which an individual's expected health is unlikely to be the only factor in their decision, and are the types of responses that may drive some of the patterns observed in the current paper. It is possible that such responses can simultaneously be utility-enhancing yet damaging – rather than protective – to health in expectation.

This paper makes four primary contributions to the existing literature. The first contribution is to introduce the hypothesis that behavioral responses to weather (and environmental conditions more generally) need not be protective in nature. The assumption that behavioral responses to environmental conditions are protective implies that such responses mitigate the biological damages associated with exposure to extreme weather. This conceptual shift in thinking regarding the nature of behavioral responses to environmental conditions introduces substantial ambiguity in terms of the predicted health impacts of weather.

Second, I depart from much of the previous literature on weather and health in focusing on measures of morbidity rather than mortality. While the temperature-mortality relationship has been well documented using large-scale panel fixed-effects methods, the literature documenting the temperature-morbidity relationship has primarily been confined to studies of specific weather phenomena or specific cities over time (de Pablo et al., 2009; Knowlton et al., 2009; Ye et al., 2012). One exception is Karlsson and Ziebarth (2016), who study the impacts of temperature on morbidity and mortality in Germany. A primary reason for this gap in the literature is simply a lack of data: there exists no nationally-representative hospitalization data for the U.S. that identifies location of residence at a spatial or temporal scale fine enough for this type of analysis.[4]

I focus on California, which since 2005 has collected data on all inpatient hospital admissions and ED visits with high temporal and spatial detail (daily data identifying the patient's zip-code of residence). An important contribution of this research is to develop reliable reduced-form estimates for the effects of temperature on measures of morbidity for

[4] The Healthcare Cost and Utilization Project's (HCUP) Nationwide Emergency Department Sample (NEDS) does not contain state identifiers. Further, HCUP does provide state-level data for an increasingly large number of states, but these data sources do not include the exact date of admission.

a large region in the U.S. Due to a lack of research in the area, the effects of temperature on non-fatal health outcomes have been left out of most climate impacts calculations (Portier et al., 2013; Houser et al., 2014). The estimates I present here indicate a substantial response to temperature in terms of hospital usage, highlighting the importance of this omission. In terms of relative risk, however, the estimates presented here are quite similar in size to those reported in the temperature-mortality literature. For instance, my estimates indicate that a day over 80°F is associated with a 5.1% increase above the mean daily visit rate, while the estimates from Deschênes and Moretti (2009) indicate an approximate 4% increase (>80°F) and the estimates from Barreca et al. (2016) indicate an approximate 10% increase (>90°F).

In any case, my estimates do indicate a sizeable change in hospital usage in response to temperature and thus there is potential for climate change to have a substantial impact on this outcome. That being said, I evaluate the predicted impacts of climate change on hospital usage in California and find very small effects. This is due to the fact that both hot and cold temperatures are associated with increases in hospital usage, so the effect of a predicted increase in the number of hot days per year is countered by the effect of a predicted decrease in the number of cold days.

The third contribution of this research is that the focus is not only on how weather affects health directly, but treatment-seeking behavior as well. When an individual is observed in the data, they are not necessarily observed at the time of the health shock, but at the time when they seek treatment for their illness. By examining the timing of health shocks relative to the associated treatment, this analysis not only furthers our understanding of the temperature-health relationship, but speaks to how individuals respond to health shocks more generally.

Finally, this paper contributes to a long-standing debate on the causes of winter seasonality of infectious diseases such as influenza. There are a variety of theories regarding the cause of infectious disease seasonality that, while not necessarily mutually exclusive, can roughly be broken into two groups (see Lofgren et al., 2007 for a review of these theories). One group of theories posits that seasonality of infectious disease is driven by cyclical factors that vary smoothly throughout the year.[5] The other group of theories posit that seasonality is driven either directly or indirectly by cold weather itself. By showing that hospital visits for respiratory diseases are highly responsive to cold weather, conditional on seasonal controls, the evidence presented here strongly supports the latter group of hypotheses. The mechanisms underlying the relationship between cold weather and infectious disease could be

[5] One such theory for which there is biological support argues that though infectious disease pathogens may be present year-round, immune-system strength varies seasonally and is at least in part a function of exposure to sunlight (Dowell and Ho, 2004).

either biological or behavioral; while the mechanisms cannot be disentangled in the context of this study, prior research suggests that both may play a role.[6]

In the Online Appendix, I present several additional results, including estimates indicating that respiratory diseases are highly responsive to low absolute humidity. These results are consistent with research relating low absolute humidity and cold temperatures to influenza mortality (Barreca and Shimshack, 2012).

The rest of the paper is organized as follows: Section 2 provides a simple conceptual framework for thinking about how behavioral responses to temperature serve to modify the temperature-health relationship. Section 3 describes the data used in the analysis, and Section 4 outlines the econometric strategy. In Section 5, I discuss the main results as well as several extensions. In Section 6, I calculate the cost of providing hospital services associated with temperature shocks and discuss implications for climate change in California. Finally, Section 7 provides a summary discussion and concludes. In the Online Appendix, I provide several robustness checks as well as a number of additional results.

2 Conceptual Framework

Previous research has considered how behavioral responses can shape the relationship between environmental conditions and health. (Neidell, 2009; Graff-Zivin and Neidell, 2009; Deschênes and Greenstone, 2011; Moretti and Neidell, 2011; Barreca et al., 2016; Graff-Zivin and Neidell, 2013). Such responses have most often been considered through a Becker-Grossman style model of health production (Becker, 1965; Grossman, 1972). At the root of this framework is the idea that health (H) is a function of environmental conditions (W, representing weather) and behavioral responses to environmental conditions ($B(W)$):

$$H = H(W, B(W)) \tag{1}$$

Totally differentiating Equation (1) equation with respect to weather yields a familiar result:

$$\frac{dH}{dW} = \frac{\partial H}{\partial W} + \frac{\partial H}{\partial B}\frac{\partial B}{\partial W} \tag{2}$$

The first component of this equation ($\frac{\partial H}{\partial W}$) represents the biological relationship between

[6]Studies of disease pathology indicate that cold and dry conditions stimulate influenza transmission (Lowen et al., 2007; Lofgren et al., 2007). A leading behavioral explanation is that individuals crowd indoors during to cold weather which increases the opportunity for the spread of infectious disease; there does exist evidence supporting this mechanism (Burström et al., 1999; Baker et al., 2000; Oxford et al., 2002; Souza et al., 2003; Adda, 2015; Stoecker et al., 2015).

weather exposure and health. For the purposes of this framework, simply assume that exposure to more extreme weather (high W) is biologically associated with worse health (low H), or that $\frac{\partial H}{\partial W}$ is negative.

The second component of Equation (2), $\frac{\partial H}{\partial B}\frac{\partial B}{\partial W}$, represents the role that behavioral responses play in modifying the overall relationship. Most of the previous literature on behavioral responses to environmental conditions has implicitly assumed that any behavioral response is directly aimed at mitigating the health consequences that would be realized through exposure to those conditions. As such, $\frac{\partial H}{\partial B}\frac{\partial B}{\partial W}$ is "protective" in nature and is assumed to be positive and at least as small in absolute value relative to $\frac{\partial H}{\partial W}$ (i.e., it attenuates the negative biological relationship between weather and health).

The assumption that these behavioral responses are protective holds only if individuals do not change their behavior in response to weather for considerations unrelated to expected health outcomes. It is probable, however, that individuals not only consider their expected health, but many other factors as well in choosing the types of activities in which they engage in response to different weather conditions. In the absence of the strong assumption that behavioral adjustments are solely protective, the relationship between weather and health becomes far more complex. In the remainder of this section, I do not attempt to derive any concrete predictions as to how weather will affect hospital usage; instead, I attempt to highlight the complexity of this relationship and the varied ways in which behavioral responses to weather can affect health. Consider the following modified framework.

Suppose there exists a set of n behaviors in which an individual can engage: $b_1, b_2, b_3, ..., b_n$. Furthermore, suppose there are two potential weather realizations: W_1 and W_2. Depending on the weather realization, individuals choose how to allocate a total of T minutes to each behavior. Denote the time allocation associated with each weather realization as B_1 and B_2:

$$\begin{gathered} t_1^1,\ t_2^1,\ t_3^1, ...,\ t_n^1 \in B_1 \\ t_1^2,\ t_2^2,\ t_3^2, ...,\ t_n^2 \in B_2 \\ \sum t_i^1 = \sum t_i^2 = T \end{gathered} \tag{3}$$

Where t_3^1, for example, is the number of minutes allocated to behavior b_3 under weather realization W_1. These time allocations are assumed to be different depending on the weather realization. There also exists a set of health outcomes H, where each element represents a specific disease. For example: suppose $h_1 \in H$ represents heart attack and $h_2 \in H$ represents influenza. Time spent engaging in each behavior is associated with a risk of experiencing each health outcome in H. Predicting how weather will affect health requires understanding exactly how individuals change their behavior in response to weather *and* how each of these behaviors maps into an extensive list of health outcomes.

Consider a simple example where W_1 is an ideal day in terms of temperature and W_2 is an abnormally cold day. Suppose an individual changes their behavior on only one dimension: they substitute a physically strenuous outdoor activity (e.g., yard work) for leisure time indoors (e.g., watching TV). Substituting away from a physically strenuous activity could translate into lower risk of injury or cardiovascular illness, while substituting to time spent indoors could translate into increased risk of contracting an infectious disease. The point on infectious diseases is important: it is not necessarily just the behavioral response of the individual in question that has the potential to affect their health, but the externalities associated with the responses of others as well (i.e., multiple individuals crowding indoors, facilitating the spread of illness). Even in this simple case, making any concrete predictions about even the direction of the overall impact that weather has on a general measure of health (e.g., the number of hospital visits) requires a number of strong assumptions.

Further complicating matters is the fact that the dynamics of the relationship between weather and hospital usage are potentially important. There are several ways in which hospital usage can be affected by weather not only on the day of the weather event, but in the days and weeks that follow. This is apparent in earlier studies of temperature and mortality, such as Deschênes and Moretti (2009), who find evidence of a "harvesting" phenomenon wherein high temperatures are found to lead to a forward displacement in mortality (i.e., exposure to hot temperatures leads to the early onset of conditions that would have been experienced in the absence of the weather shock). It is possible that a similar harvesting effect exists in the relationship between temperature and morbidity, though other possibilities exist as well. Another possibility is that symptoms develop or worsen in the days following a weather event. If, for example, exposure to cold weather (or the behavioral response to it) leads to an increased risk of contracting an infectious disease, the increase in hospital visits would not be expected to occur until several days or weeks after the event.

Furthermore, unique to this setting is that hospital usage represents a decision to seek treatment rather than an illness itself. This is important to keep in mind for a number of reasons. If extreme weather affects the cost of seeking treatment (i.e., people may have an aversion to travelling to seek medical treatment in times of harsh weather), then hospital usage may be affected even if health is not. This has potential implications for the dynamics of the relationship: treatment for a period t health shock may not be sought until period $t + 1$ or later.

This discussion highlights the notion that developing clear predictions of the effects of temperature on morbidity and hospital usage would be a difficult task requiring an extensive list of assumptions, and the empirical analysis proceeds with this consideration in mind. Indeed, the first goal of the empirical analysis and a primary contribution of this paper will be to document the reduced-form relationship between temperature and hospital usage. This includes documenting both the total net effect of temperature on hospital usage as well as the dynamics. Given the observed relationship, I then attempt to shed light on the factors that drive the overall relationship.

3 Data

This study merges data on hospital visits with meteorological conditions based on the patient's zip-code of residence and the exact date of the visit. The data represent the period 2005-2014.[7] Given its large size (in both geography and population, representing more than 10% of the total U.S. population) and the wide variety of climate zones across the state, California is an ideal microcosm for this analysis given that more widespread data is not available. On the other hand, California is a relatively warm state, and its colder regions are more sparsely inhabited than the rest of the state. The results generated for the cold end of the weather distribution, however, are not necessarily due to extreme cold weather, but deviations from ideals. This assertion is supported by the fact that I find similar results for all cold-weather temperature ranges; that is, the patterns I estimate for temperatures in the ranges 40-45, 45-50, and 50-55 mirror the patterns estimated for temperatures under 40°F, though with smaller magnitudes (similarly, the estimates for less extreme hot temperatures mirror the patterns estimated for the highest temperature bin). As such, the results for cold temperatures do not rely on a small number of particularly cold regions.[8]

3.1 Weather

Data on temperature and precipitation are drawn from the Global Historical Climatology Network (GHCN).[9] This dataset, according to the National Climatic Data Center (NCDC), contains the most comprehensive set of land-based weather stations available anywhere. Due to a large number of missing values in weather data, I only use station-years that have complete daily records. I conduct this procedure separately for stations that measure tem-

[7] 2005 was the first year in which data on outpatient ED visits are available.

[8] Further, all regressions are weighted by zip-code population so that all estimates can be interpreted as representative of the population.

[9] The GHCN does not provide measures of humidity. In the Online Appendix, I use an alternative set of weather stations to show that the inclusion of humidity does not alter the results.

perature and stations that measure precipitation.[10] Using this rule, the number of stations used in the analysis ranges from 174 (in 2006) to 328 (in 2014) for temperature, and from 167 (in 2006) to 324 (in 2014) for precipitation. These daily station level data are aggregated to the zip-code level by taking a weighted average of all stations within a 30km radius of each zip-code's centroid. The weight applied to each station is the inverse of the squared distance between the zip-code centroid and each station (using exact latitude and longitude coordinates) such that stations that are closer to the centroid of each zip-code are given more weight.

The primary variables of interest in this analysis are ten five-degree (Fahrenheit) mean temperature bins ranging from under 40°F to over 80°F. These variables indicate whether temperature measured at a given station falls in the specified temperature range; note that these variables are constructed prior to aggregation to the zip-code level to preserve the variation in temperature that exists within each zip-code (Deschênes and Greenstone, 2011; Dell et al., 2014). The 60-65°F bin is omitted in all regressions such that all estimates are interpreted as the impact of a day in the given temperature range relative to a day in the 60-65°F range.

3.2 Health

Data on ED visits and hospitalizations were obtained through California's Office of Statewide Health Planning and Development (OSHPD). This consists of two restricted data files for the period 2005-2014. The first file contains the universe of outpatient visits through the emergency department; the second contains the universe of inpatient visits, whether through the emergency department or not. It is important to note that the years included in each dataset are based on the patient's discharge date rather than admission date; this is of little consequence for outpatient visits as the admission date and discharge date are at maximum one day apart. This is problematic for inpatient admissions, however, as patients discharged after 2014 are not observed in the data: there exists a severe under-counting problem for inpatient admissions at the end of the sample period. For this reason, I drop admissions in December of 2014 from the analysis. It is still the case that patients admitted before December 2014 and released in 2015 or later are not counted, but since only 1.1% of patients in the sample stay for more than 31 days, it is unlikely to affect the analysis in any meaningful way.

For most of the analysis, I focus on visits that took place through the emergency depart-

[10] The reason for this is that there are a subset of stations that exclusively measure temperature and another subset that exclusively measure precipitation. Using stations that have complete records for both temperature and precipitation results in lower spatial coverage.

ment, as these are more likely to represent actual health shocks. The same argument cannot be made for all inpatient visits, of which many are scheduled (e.g., surgery), or in some way inevitable (e.g., childbirth). An emergency department visit that results in an inpatient stay can be interpreted as a more severe outcome.[11] For much of the analysis, outpatient and inpatient visits (through the ED) are combined into a single measure.

The patient's zip-code of residence and the exact date of the visit are used to merge health and weather data. In addition, the data contains information on the patient's age, gender, the patient's expected source of payment, and the principal – and up to 24 secondary – diagnoses. The diagnosis codes correspond to the ninth revision of the International Classification of Diseases (ICD-9-CM). In many cases, these codes (of which there are approximately 14,000) will be converted using Clinical Classifications Software (CCS). CCS codes were developed by the Healthcare Cost and Utilization Project (HCUP) for the purpose of collapsing the large number of ICD codes into clinically meaningful categories for use in data analysis. The CCS coding system offers four levels of aggregation, the highest of which aggregates all ICD codes into 18 categories. The final two CCS categories ("ill-defined conditions" and "residual codes") are omitted from the analysis of disease categories and thus only 16 categories are presented. While the CCS coding system will prove useful for categorizing diseases into highly aggregated groups, other aggregation schemes based on ICD codes will be used throughout portions of the analysis and are described in detail as they are used.

The primary outcome of interest is the ED visit rate per 100,000 population. This rate is calculated by dividing the total number of ED visits (outpatient visits plus visits that resulted in an inpatient stay) on a given day, within a given zip-code, by the total population in that zip-code.[12] Because calculating this rate for zip-codes with very small populations can result in very large visit rates, I drop the smallest zip-codes in the sample representing 1% of the total population of California.[13] Age-specific population counts are available, and these are used in calculating age-specific visit rates.

Summary statistics for these data are presented in Table 1. This includes information on temperature variables as well as the outcomes of interest (total visits, five age groups, and sixteen disease categories). For the data on ED visits, I include the total number of visits represented by each category; the primary outcome of interest (all visits) represents more

[11] It cannot be determined from the data whether an inpatient visit took place through a different hospital's emergency department. There is also information on whether the inpatient stay was scheduled, though it is redundant to condition on *unscheduled* when also conditioning on *emergency*, since emergency department visits cannot be scheduled.

[12] Zip-code population data are taken from the U.S. Census Bureau's population estimates for 2010. Note that zip-code level population data is only available on a decennial basis.

[13] These are zip-codes with populations less than 2,300 individuals. Ultimately, this makes very little impact on the final estimates as all models are weighted by zip-code population.

than 94 million visits to the emergency department. Means for the temperature variables should be interpreted as the percentage of days that fall into each temperature range.

4 Econometric Strategy

The goal of the econometric model is to estimate the causal effect of a temperature realization on day t on hospital visits not only on the same day, but for a number of days in the future as well. Conceptually, the goal is to estimate the following causal effects: $\frac{\partial \text{Hosp}_t}{\partial \text{Temp}_t}, \ldots, \frac{\partial \text{Hosp}_{t+30}}{\partial \text{Temp}_t}$. In practice, I employ a distributed lag model wherein the hospital visit rate is regressed on contemporaneous temperature as well as at least 30 lags of temperature. Such a model directly estimates the following partial effects: $\frac{\partial \text{Hosp}_t}{\partial \text{Temp}_t}, \ldots, \frac{\partial \text{Hosp}_t}{\partial \text{Temp}_{t-30}}$. The difference between the directly estimated partial effects and the conceptual goal lies only in the interpretation. To be clear, $\frac{\partial \text{Hosp}_t}{\partial \text{Temp}_{t-h}}$ simply represents the change in hospitalizations on day t, given a change in temperature h periods in the past; $\frac{\partial \text{Hosp}_{t+h}}{\partial \text{Temp}_t}$ has an equivalent interpretation. The fully-specified model used to estimate these partial effects is described below.

$$\begin{aligned}\text{Hosp}_{z,t} = \alpha + \sum_{j=1}^{9}\sum_{h=0}^{30}\beta_{j,t-h}\text{Temp}_{j,z,t-h} + \sum_{h=0}^{30} f\left(\text{Precip}\right)_{z,t-h} \\ + \delta_{zip-week} + \theta_{county-year} + \eta_{day-of-week} + \kappa_{holidays} + \varepsilon_{z,t}\end{aligned} \tag{4}$$

In this model, $\text{Hosp}_{z,t}$ represents the hospital visit rate for individuals living in zip-code z, on day t. The coefficients of interest are the 279 coefficients on the temperature variables, $\beta_{j,t-h}$. These variables represent nine 5-degree temperature bins as described in Section 3.1. The bin representing temperatures in the 60-65°F range is omitted; as such, all estimates are measured relative to this temperature range. Controls for precipitation are included for each day in which temperature is measured; in the main specification, these controls enter in the form of a third-order polynomial.

The inclusion of a suitable set of fixed effects is extremely important to identify the causal effect of random variations in temperature on hospital visits. Most importantly, it is necessary to account for the fact that both hospital visits (and health in general) and weather vary seasonally. In the baseline specification, I include zip-by-week-of-year fixed effects. This allows seasonality to be measured at a relatively fine scale (weekly) and allows seasonality effects to vary by zip-code. The use of these highly flexible week-of-year fixed effects in place of a more aggregate measure of seasonality (e.g., months or seasons) is potentially important if changes in health are driven by changes in behavior. For example, the last week of December, a relatively cold week on average, is very different from a typical week in a behavioral sense (which could be related to health). Also included are county-by-year fixed effects, which control for annual factors (such as health policy or inter-annual weather events like El Niño) that are allowed to vary by county. Additionally, all specifications include both day-of-week effects and indicators for all national holidays as both hospital usage (and behavior in general) can vary substantially by day-of-week or on holidays. It should be noted, however, that since weather is orthogonal to both the day of week and holidays conditional on seasonal controls, these effects are not necessarily included for the purpose of causal identification, but to increase the precision of the estimates.

While this panel fixed-effects methodology does allow for the inclusion of a wide variety of controls in the form of highly flexible fixed effects, it should be noted that this analysis relies on variation in weather within a single state. As such, an important caveat is the fact that within a given time period, weather is correlated (albeit not perfectly) across all cross-sectional units (i.e., zip-codes). To the extent that weather is imperfectly measured, correlation in weather across zip-codes will induce correlation in the error term across zip-codes within a given time period. To account for this, I employ a two-way clustering strategy and cluster on both the county (i.e., a cross-sectional unit that is wider than the unit of analysis) and year-month. This strategy allows for arbitrary correlation in the error term across all zip-codes and days within a given county *and* across all zip-codes and days within a given year-month.[14]

It is important to note that the estimated temperature effects represent the effect of temperature in time period t on hospital visits in time period $t + h$, controlling for meteorological conditions on every other day in the 31-day period. While the specific features of the dynamic relationship are discussed, I focus on two measures derived from this empirical model. The first is the contemporaneous effect: $\beta_{j,t}$, which measures the impact of temperature on hospital visits on the day of the event. The second is the cumulative effect, which measures the 31-day cumulative effect of the one-day temperature shock. This "dynamic cumulative effect" is constructed by summing all coefficients within a given temperature bin: $\sum_{h=0}^{30} \beta_{j,t-h}$.[15] Note that because the cumulative effect is a simple linear combination of coefficients, the calculation of the standard errors for these estimates is straightforward and requires no additional assumptions. The cumulative effect is meant to capture the total

[14] In a previous draft of the paper, I compared multiple strategies for estimating the standard errors including clustering at only the county level, allowing for spatial correlation (Conley, 1999), and the two-way strategy used here. I found that the two-way strategy resulted in the most conservative estimates (i.e., the largest standard errors). In this previous draft, all models were estimated at the county level (with very similar results).

[15] For more on dynamic causal effects and dynamic cumulative effects, see Stock and Watson (2003).

effect of temperature on hospital visits, and corresponds to estimates calculated at more highly aggregated time-scales (i.e., monthly and annual) that are common elsewhere in the literature. In the following section, I present additional estimates from a model estimated at the monthly level with results that are very similar to the cumulative effects estimated at the daily level.

5 Results

5.1 Baseline Analysis

The main results are summarized in Table 2. Each set of results (i.e., each regression) is displayed in two columns. The first column for each set of results represents the contemporaneous effect ($\beta_{j,t}$) and the second column represents the cumulative effect ($\sum_h \beta_{j,t-h}$). Though the models are estimated in levels throughout the analysis, these results are presented in percent changes (calculated by dividing the corresponding coefficient or sum of coefficients by the mean daily visit rate). This is done simply for ease of interpretation, which will be especially important in the sub-group analyses to come. Furthermore, for the sake of brevity, I report estimates for only the four most extreme temperature bins (two hot and two cold). The full set of results (i.e., all temperature bins) are available in the Online Appendix. The results labelled "All Emergency" in Table 2 are the baseline estimates, which represent all inpatient and outpatient hospital visits through the emergency department. The interpretation of the contemporaneous effect estimate for the <40°F temperature bin is as follows: a day under 40°F is associated with approximately 6.1% fewer ED visits on the day of the event, relative to a 60-65°F day. The level effect can be recovered by multiplying the percentage effect by the mean daily visit rate (77.1 visits per 100,000 individuals); this implies approximately 4.7 fewer visits per 100,000 residents.

Taken on its own, the contemporaneous effect would suggest that population health benefits from cold weather. When visits in the days following a cold temperature shock are considered, however, the relationship changes considerably. The cumulative effect estimate for the <40°F bin implies that a one-day cold temperature shock leads to a total net *increase* in ED visits equal to approximately 11.0% of the mean daily visit rate. In levels, this implies approximately 8.5 additional visits per 100,000 residents.

Hot temperatures are associated with both contemporaneous and cumulative increases in the ED visit rate. Specifically, a day over 80°F is associated with a 3.5% contemporaneous increase, which grows to an 5.1% cumulative increase over the course of the 31-day period.

These results are summarized in Figure 1, which mimics a figure used in many recent analyses of the effects of temperature (Deschênes and Greenstone, 2011; Barreca, 2012; Barreca et al., 2016). The cumulative effects display a U-shaped relationship that is comparable to estimates generated in much of the literature on temperature and mortality. The estimates in such studies are typically conducted at more highly aggregated time scales (monthly, bimonthly or annual) and are meant to capture the total effect of a temperature shock on mortality. This figure also shows that for both cold and hot temperatures, the less extreme bins display comparable patterns to their more extreme counterparts, but with lower magnitudes. This emphasizes the fact that these results are not driven by extreme temperatures, but deviations from the ideal range.

In addition to the main results, Table 2 presents results for different types of visits. The results labelled "Outpatient Emergency" and "Inpatient Emergency" decompose the main effect into these two categories; these results show that outpatient visits (i.e., less severe visits) are more responsive to temperature than inpatient visits. The "Non-Emergency" results, which represent visits that did not occur through the ED (and are not included in the baseline estimates) show that inpatient non-emergency visits are not responsive to temperature; this is unsurprising given that many such visits are planned.

The choice of a 31-day window was motivated by the prior literature on temperature and mortality, which has shown that the lagged effects of temperature tend to die out within approximately one month (Basu and Samet, 2002; Armstrong, 2006; Deschênes and Moretti, 2009). The same pattern need not hold for morbidity. The first four columns of Table 3 display the cumulative effects of temperature using up to 60 lags in temperature. For cold temperatures, an extended lag window meaningfully impacts the estimates, which climb from 11.0% using a 31-day window to 17.0% using a 61-day window (though the impacts seem to flatten after 50 days). That being said, there is a substantial loss in precision in estimating models over such a long time horizon; this loss in precision is amplified in the sub-group analyses to follow. As such, I consider the 31-day cumulative impacts to be the preferred specification throughout the analysis, though it should be noted that these estimates are likely a lower bound on the total net effect of cold temperatures on hospital usage.

The final column of Table 3 represents estimates of the cumulative effects derived from a model estimated at the monthly level that is typical of the literature on temperature and mortality (Barreca, 2012; Barreca et al., 2016). These estimates are presented as a check on the baseline model to ensure that model specification is not driving the results as estimated in the main analysis. The regressors of interest in this model are the total number of days in a month within each temperature bin and are meant to capture the effect of an additional day within each bin.[16] Reassuringly, the estimates generated from the model estimated at the monthly level are quite similar to the estimates using daily data.

For much of the analysis, I summarize the dynamics of the temperature-health relationship by reporting just the contemporaneous and cumulative effects. It is useful, however, to understand exactly how this relationship evolves. Figures 2 and 3 are intended to illustrate how hospital usage is affected following a temperature shock. These figures also present a falsification test by estimating models with leads in temperature instead of lags. In each figure, points to the left of zero represent leads (i.e., placebo estimates) and points to the right of zero represent lags (i.e., the treatment effects). In all cases, the falsification tests suggest that future temperature shocks do not meaningfully affect past ED visit rates. Each figure displays two plots, one describing the dynamics associated with a day under 40°F, and another describing a day over 80°F.

Figure 2 plots the raw coefficient estimates. The contemporaneous effect can be identified on these plots as the point where days relative to the shock equals zero; the point corresponding to 30 days after to the shock represents the effect of temperature on the ED visit rate thirty days after the temperature shock.

Figure 3 plots the sum of all coefficients up to and including the relevant lag (or lead). For example, the point corresponding to two days after to the shock represents the sum of coefficients on contemporaneous temperature as well as the first and second temperature lags ($\beta_{j,t}+\beta_{j,t-1}+\beta_{j,t-2}$). In other words, this point reflects the 3-day cumulative effect. Similarly, the point corresponding to 30 days after to the shock represents the 31-day cumulative effect ($\sum_{h=0}^{30} \beta_{j,t-h}$). These plots are intended to display how the total number of visits accumulate following a given temperature shock. In each plot, the impacts are reported as level effects rather than percent effects; this is for a subsequent comparison with Figure 4 which decomposes the total effect (reported here) into disease categories.[17]

For cold temperatures, the contemporaneous drop is plain to see in either plot (indeed, the estimates at $t = 0$ are equivalent across both figures). Figure 2 reveals that the coefficients

[16] The monthly model is estimated using zip-by-month and county-by-year fixed effects and includes a one-month lag in temperature to allow for longer-run impacts. It is important to note that a monthly model that uses only the current month is *not* directly analogous to the 31-day cumulative effect estimated with daily data. This is because it is possible that the temperature shock of interest could have occurred at any day within the month; if the shock occurred on the last day of the month, only the contemporaneous daily effect of temperature would be captured. Under the assumption that the temperature shock occurs on any day of the month with equal probability, a monthly model with only current month (i.e., no lagged month) estimates $\sum_{h=0}^{30} \frac{31-h}{31}\beta_{t-h}$ rather than $\sum_{h=0}^{30} \beta_{t-h}$. Similarly, a monthly model that does include a one-month lag is not directly analogous to the 61-day cumulative effect estimated using daily data, though longer-run impacts are more fully captured.

[17] Note that the level effects can easily be converted to percentage terms (as reported in the tables) by dividing by the mean daily visit rate (approximately 77.1).

are near zero for roughly a week after the shock, after which they hover above zero for the remainder of the 31-day period. Figure 3 shows how this bears out in terms of total visits: the sum of positive coefficients observed several days out from the initial shock is eventually enough to make up for the contemporaneous drop. The total net effect of a cold day on ED visits not only climbs back to zero, but becomes positive (and statistically significant) by the end of the period.

For hot temperatures, Figure 2 reveals a large contemporaneous increase in ED visit rates followed by smaller statistically significant increases for each of the next four days. Afterwards, the coefficients hover near zero for the remainder of the 31-day period. Again, Figure 3 shows how these dynamics play out in terms of total visits: there is a large increase in visits on the day of the shock, then the cumulative number of visits increases for approximately one week before leveling out. Note that the majority of the coefficients beyond 10 days from the shock are slightly negative, resulting in a 31-day cumulative effect that it slightly smaller in magnitude than (though not statistically different from) the cumulative effect measured 10 days beyond the shock.

Interestingly, this pattern for hot temperatures is quite distinct from the dynamics observed in studies of temperature and mortality. In particular, the increases in ED visits in the days that follow the shock are distinct. Furthermore, the literature on mortality has found evidence of a "harvesting" phenomenon, under which increases in contemporaneous mortality would be curbed by decreases in mortality in the following days, resulting in a net effect close to zero (Deschênes and Moretti, 2009; Karlsson and Ziebarth, 2016). Harvesting is generally interpreted as a forward displacement in mortality. There is some evidence of a harvesting effect here as well in that the total net effect (after 30 days) is smaller than the peak effect (after 10 days), but the total net effect is larger – rather than smaller – than the contemporaneous effect.

The goal to this point has been to document the overall relationship between temperature and hospital usage. In the sections that follow, the goal is to uncover heterogeneity in this relationship and to provide evidence on the mechanisms at work.

5.2 Age Heterogeneity

Separating hospital visits into five age groups (under 5, 5-14, 15-24, 25-64, and over 64) reveals significant impacts of temperature across all age groups. These estimates are reported in Table 4. Though the signs of the estimates are consistent across ages, there are substantial differences in magnitudes. What is immediately clear and of first-order importance is that the youngest age groups are most responsive to both cold and hot temperatures. This is significant because temperature-driven mortality is typically concentrated among the elderly. The implication of this is that studies of temperature and health that focus on mortality do not capture the potentially large impacts experienced by groups whose health is sensitive to temperature fluctuations even if they are unlikely to die from such exposure.

This point is illustrated in the cumulative impacts: the impacts of cold weather for the most affected age group (children under 5), at 27.7% above the mean daily visit rate, is four times as large as that of the least affected group (non-elderly adults aged 25-64) at 6.9%. Similarly, the cumulative effects of hot weather for the youngest two categories, at 8.6% and 9.5% for the under 5 and 5-14 groups, are more than twice as large as the impacts for any other group (the largest of which is 4.0%). It should be noted that interpretation of the cumulative impacts should be taken with some caution as the estimates lack precision, though they are often statistically different from zero (especially the largest impacts).

More subtle differences across age groups are also worth noting: consider the differences across groups in the difference between contemporaneous and cumulative effects for hot weather. The most striking comparison is between the 5-14 and the Over 64 age groups. The 5-14 group experiences a 2.1% contemporaneous increase that climbs to a 9.5% cumulative effect. The Over 64 group experiences a 3.7% contemporaneous increase that *falls* to a 2.7% cumulative effect. This likely reflects differences across groups in the types of diseases that are responsive to temperature (or the behavioral response to temperature), but this also provides some intuition regarding the evidence of harvesting in the all-age analysis. The harvesting effect (a contemporaneous increase that is offset by later decreases) appears to be driven by the oldest age category; this result suggests that the harvesting phenomenon observed in the elderly in other studies does not necessarily extend to other populations.

5.3 Disease Heterogeneity

Breaking the analysis down by disease category provides an illuminating view of the broad relationships described above. Not only can this inform us as to which specific disease types drive the overall temperature-hospital relationship, but these estimates provide insight into how both biological and behavioral factors influence this relationship. The variable that makes this sub-sampling possible is each patient's principal diagnosis, reported as an ICD-9-CM diagnosis code. I aggregate these ICD codes using Clinical Classification Software (CCS), which is used to convert approximately 14,000 ICD codes into 16 clinically meaningful categories.

I separately estimate Equation (4) for each category and summarize the results in Figure 4. This figure plots – for both hot and cold temperatures – the cumulative impacts (similar to Figure 3) for all 16 disease categories. The estimates are reported in levels such that the sum of the estimates across all categories approximately equals the total impact described in Figure 3. The purpose of Figure 4 is to highlight whether any specific disease types drive the overall estimates. Because of the large number of categories represented on each plot, confidence intervals for each category are not displayed. Furthermore, I highlight the six largest disease categories using colored markers and display the other ten categories as grey lines. Note that the number of visits for each category is presented in Table 1. While small categories typically do not generate large impacts in absolute terms, they may generate large relative impacts and this is not well captured in Figure 4; individual plots (with confidence intervals) for each of the 16 categories are presented in Figures A1 and A2.

For cold temperatures, the contemporaneous decrease in ED visits is common across all categories and the largest categories contribute most to the total impact. The cumulative increase in visits following the shock, however, is driven by the Respiratory category far more than any other disease category. Indeed, the cumulative increase for all disease types implies an additional 8.5 visits (per 100,000) and the Respiratory category alone accounts for approximately 6.0 additional visits (per 100,000). This is a large increase in visits whether considered in absolute terms or relative terms: 6.0 additional visits represents an increase of approximately 60% above the mean daily visit rate for the Respiratory category. This is not the only disease category for which there is a large cumulative effect in relative terms: the cumulative effect for the Infectious/Parasitic category is an increase equivalent to approximately 35% of the mean daily visit rate for this category, though the contribution of the Infectious/Parasitic category to the overall impact is small due to the relatively small number of visits for these diseases.

An important common factor between these two disease categories (Respiratory and Infectious/Parasitic) is that a high proportion of the specific diseases within each category are communicable. More than half of the diseases in the Respiratory category fall under the "Respiratory Infections" sub-category, of which specific diseases such as influenza and pneumonia are a part.[18] While this evidence does not necessarily speak to a specific mechanism (e.g., indoor crowding), the timing of the relationship between cold weather and ED visits for these categories is consistent with the hypothesis that cold weather leads to increased exposure or susceptibility to communicable disease. The relationship is relatively flat in the days immediately following the shock followed by consistent increases throughout the remainder of the period; the flat period is consistent with an incubation period following exposure to a communicable illness (Adda, 2015).

[18] For reference, Tables A1 and A2 display the total number of ED visits in the sample by CCS sub-category; there are 136 such sub-categories.

It should also be noted that while the Infectious/Parasitic and Respiratory disease categories display the most pronounced cumulative increase, they are not the only categories characterized by these dynamics. Indeed, diseases in the Digestive, Nervous System, Endocrine and Circulatory categories all show roughly similar patterns, though to a lesser extent. Though each of these categories may contain some specific disease codes corresponding to communicable diseases, there are other possible explanations for such a phenomenon.[19] One possibility is an increase in the exposure to communicable diseases – even diseases that are relatively benign in most people (e.g., the common cold) – leads to the exacerbation or triggering of symptoms for individuals who were already ill or at risk. Thus, it is not necessarily the communicable disease for which a person requires medical care, but another disease for which the communicable disease acts as a trigger. Another possibility is a deferral of treatment: given the large decrease in contemporaneous ED visits, it could be that some or all of the "missing" ED visits are displaced to a later date. That being said, there are several categories for which the dynamics are negative and flat for the entire period, though the cumulative effects that are not always significantly different from zero (e.g., Injuries, Mental Illness, and Skin Conditions).

For hot temperatures, it is clear from Figure 4 that there are several disease categories that contribute to the overall relationship; this includes several of the largest disease categories (e.g., Injuries) as well as several of the smaller categories. While the effect on Injuries is the largest in absolute terms and thus these visits contribute most to the overall relationship, this is not the case in relative terms where the largest cumulative increases are exhibited by Endocrine Disorders and Diseases of the Perinatal Period. It is not surprising that diseases in either of these categories would be quite sensitive to high temperatures. The most common sub-category of the Endocrine Disorders category is "Fluid and Electrolyte Imbalances"; it is likely that increases in these types of diseases are driven by dehydration. Given the finding in Section 5.2 indicating that young children are quite sensitive to temperature in general, it is unsurprising that diseases of the perinatal period are strongly affected.

It is also worth mentioning the relatively large cumulative decrease for Respiratory diseases in response to hot weather. It may be the case that hot weather (or the behavioral response to hot weather) suppresses the spread of communicable disease in the same way that cold weather encourages it. That being said, the cumulative effect for the Respiratory category is not statistically different from zero at conventional levels and as such any interpretation should be taken with caution.

[19] Examples of specific communicable diseases in these categories include types of gastroenteritis (Norovirus or Rotavirus (Glass et al., 2009; Dennehy, 2011)) in the Digestive category and Bacterial Meningitis in Nervous System category.

It is beyond the scope of this paper to conduct an in-depth analysis of how visits for each specific disease type responds to temperature; the focus instead is on explaining the overall relationship. Consider the following question: if exposure to heat leads to an increase in ED visits, why do they not all take place on the day of the shock? It could be that there is a delay in the manifestation of symptoms that require medical treatment, but this may also indicate that there is a delay between when symptoms are experienced and when treatment is sought. The latter scenario is potentially problematic if earlier treatment is associated with better outcomes. Analyzing these patterns by disease "deferrability" in the following section will provide some insight into these possibilities.

5.4 Mechanisms

The purpose of this section is to shed light on and discuss the mechanisms that underlie the overall relationships between both cold and hot weather and ED visits as summarized in Figures 2 and 3.

First, what factors drive the contemporaneous decrease in ED visits during cold weather? This is perhaps the most puzzling result as it is unlikely that exposure to cold weather is health-improving. Behavioral factors are far more plausible. One possibility is a decreased willingness to seek treatment during unfavorable weather; it is well known that many ED visits are for ailments that do not require emergency care (Taubman et al., 2014). It is then plausible that any factor that increases the cost of seeking treatment (e.g., bad weather) will decrease the rate at which individuals seek treatment. Under this hypothesis, there would not necessarily be any immediate change in health as a result of cold weather, but there would be a change in hospital usage. Alternatively, it is possible that there is an actual improvement in health (or suppression of symptoms) during cold weather. Again, while exposure to cold weather is unlikely to improve health, it is possible that the behavioral response to cold weather is health-improving. This would be the case if individuals engage in a different set of activities depending on the weather and cold weather activities are less conducive to negative health outcomes.[20]

The goal of the following exercise is to determine whether individuals are simply less likely to go to the hospital during cold weather or if there is an actual decrease in the incidence of disease. The strategy relies on classifying hospital visits into categories representing diseases that are more or less deferrable in nature. I use a classification developed by Billings

[20]Another possibility is that more individuals die during these cold weather events before going to the hospital. This mechanism is implausible, however, due to the numbers involved here: a 6.1% decrease in ED visits implies 4.7 fewer deaths per 100,000. This is more than double the daily mortality rate as reported in Deschênes and Moretti (2009). Furthermore, estimates from Deschênes and Moretti (2009) also suggest a small but statistically insignificant decrease in mortality during cold weather.

et al. (2000) that assigns each ICD diagnosis code a probability of belonging to one of four categories: Non-Emergent, Primary Care Treatable, ED Care Required and Preventable, or ED Care Required and Non-Preventable. These probabilities were developed by a team of ED physicians who classified thousands of individual cases into these categories; ICD codes are assigned probabilities rather than mutually exclusive categories because some of these cases had the same principal diagnosis yet were classified into different categories. More information on this procedure can be found in the original article or in subsequent articles that have made use of these classifications (e.g., Taubman et al. (2014)). I merge these probabilities to the hospital data and construct zip-by-day counts for each category by summing the probabilities associated with each visit. Because I am only interested in the relative deferrability of categories, I compare admissions for the least emergent category with those in the most emergent. Specifically, I consider visits in the "Non-Emergent" category to be relatively more deferrable (e.g., joint pain and warts) and visits in the "ED Care Required and Non Preventable" category to be relatively less deferrable (e.g., heart attack and stroke).[21]

I estimate the effects of temperature on visit rates for diseases in these two categories and present the results in Table 5. For now, let us only consider cold temperatures: the estimates of the contemporaneous effects for the "more deferrable" and "less deferrable" categories are -6.5% and -5.8%, respectively. While the estimate for "less deferrable" diseases is slightly smaller in magnitude, we would expect that this estimate would be near zero if individuals were simply choosing not to seek treatment. The implication is that there appears to be a decrease in the actual incidence of disease on the day of a cold weather event. While it is difficult to make any more specific claims about the mechanisms at work, this is an interesting finding in that it suggests that cold weather (or the behavioral response individuals take in response to it) has the potential to *improve* health on net, at least for some types of disease. For most disease types, the total effect does rebound to at least zero in subsequent weeks, but there are some disease categories for which the total net effect is negative: Injuries, Skin Conditions and Mental Illness.

Let us now turn to the increase in ED visits in the month that follows a cold weather event. One possible mechanism is related to the contemporaneous decrease in visits: the "missing" visits on the day of the shock may simply be deferred to a later date; similarly, activities that are unappealing during cold weather and carry some risk of acute illness or injury may have been deferred to a later date (e.g., outdoor activities that require exertion such as yard

[21] Note that in the Online Appendix, I test the robustness of this classification method by using an alternate method to classify diagnoses as more or less deferrable. This alternate method is based on the method used in Dobkin (2003); Card et al. (2009); Doyle et al. (2015) and generates very similar results to those presented here.

work). It is certainly possible that this is responsible for a portion of the increased visits that follow a cold weather shock, but if this were the only factor it would not be expected that the cumulative increase in admissions would cross zero and become positive. The increase in contagious illnesses uncovered in the analysis of disease heterogeneity is responsible for the majority of this increase in ED visits.

Prior research has show that individuals are more likely to stay indoors during cold temperatures (Graff-Zivin and Neidell, 2014). It is possible that indoor crowding (the simultaneous behavioral response of many individuals) increases the likelihood that contagious illnesses are spread, and there exists empirical evidence on this point (Burström et al., 1999; Baker et al., 2000; Oxford et al., 2002; Souza et al., 2003; Adda, 2015; Stoecker et al., 2015). It is not possible in the context of this study, however, to concretely discern whether the increase in contagious illnesses following a cold weather event is due to biological or behavioral factors. The finding remains important, however, as it shows that contagious illnesses are highly responsive to cold weather itself rather than other factors that vary at the seasonal level.

Let us now consider the mechanisms that drive the relationship between hot weather and ED visits. It is unlikely that the effect of heat on ED visits observed in this study operates entirely through a behavioral channel as opposed to a biological mechanism resulting from heat exposure. This notion is supported by the finding that disease categories representing heat-related illness are most strongly affected. For example, the contemporaneous effect for the Endocrine category (which includes heat-related illnesses such as Dehydration) represents an approximate 14% increase above the mean daily visit rate for this category, approximately twice as large as the second most affected category in relative terms.

The final point to be made regarding mechanisms addresses the increase in ED visits that follow a hot weather event. Unlike the increased visits that follow cold weather, these tend to occur in the few days that follow the shock. A simple explanation returns to the notion that individuals may defer treatment for an illness: if individuals experience a health shock on the day of the hot weather event, they may not seek treatment for that ailment until a later date. The estimates presented in Table 5 support this claim. For illnesses that are more deferrable, the cumulative effect is more than twice the size of the contemporaneous effect; for illnesses that are less deferrable, the contemporaneous and cumulative effects are of very similar magnitude. The implication is that when individuals have the ability to defer treatment, they will often do so. This has potentially important welfare implications that extend beyond the scope of temperature-induced health shocks if early treatment is associated with better outcomes.

6 Implications

6.1 Cost Estimates

In this section, I abstract from the specific mechanisms at work and any heterogeneous treatment effects to consider the total costs associated with a given temperature shock and provide implications for the effects of climate change in California on both the number of ED visits and health care costs. It must be noted that I do not estimate the willingness to pay to avoid a negative health outcome, only the cost of providing hospital services; as such, all estimates here represent a lower-bound on the total costs of temperature-induced morbidity.

A simple way to estimate these costs would be to multiply the estimated effect of temperature on ED visits by an estimate of the average cost of an ED visit. As I have shown, however, it is not only the number of ED visits that are responsive to temperature, but the composition. To account for this, I take a different approach and re-estimate Equation (4) using total hospital costs (per 100,000 population) in a given zip-code and day as the dependent variable. To accomplish this, I assign each ED visit a cost estimate.

The process of assigning costs to visits depends on whether the visit resulted in an inpatient stay. The process is slightly simpler for inpatient visits because the inpatient files from OSHPD report the total billed charges for each visit whereas the outpatient ED files do not.[22] These reported charges still do not represent the cost of providing services, however. The Healthcare Cost and Utilization Project (HCUP) provides data on the ratio of costs to charges for all hospitals included in their National Inpatient Sample (NIS). I construct a measure of the cost of providing services for each inpatient visit by multiplying the total billed charge for each visit by the California-specific mean cost-to-charge ratio (0.288) obtained from the 2010 NIS.

To calculate outpatient ED costs, I first need to estimate the charges associated with each visit. Though information on charges is not available in Calfornia, HCUP's Nationwide Emergency Department Sample (NEDS) does contain information on ED charges for a subset of observations (this subset does not include observations in California, though the NEDS does not provide state-identifying information in any case). To assign charges to each visit in the California data, I calculate the median charge for each ICD-by-age group cell in the

[22]Note that there are two potential issues with the data on inpatient charges. First, charges are missing from approximately 12.0% of observations in the inpatient data; for these observations, I assign charges to be equal to the median charge in each ICD-by-age-group cell. Second, the distribution of charges is highly skewed, with numerous observations in the millions of dollars, including 69 observations top-coded at $9,999,999. So that total charges in each zip-by-day observation are not driven by these extreme values, I further top-code the data on charges at the 99th percentile of the distribution ($404,117).

2010 NEDS data, and assign these charges to the relevant observations in the California data. I then multiply by the national average cost-to-charge ratio (0.507) derived from the 2010 NIS to obtain an estimate of ED costs for each observation.[23]

The object of interest in these regressions is the (31-day) cumulative effect, and thus only these impacts are reported here. The first column of Table 6 displays the effects on total cost (per 100,000 population) for all visits that took place through the ED. The second two columns decompose this effect by inpatient and outpatient visits. While there are substantially more outpatient than inpatient visits, the impacts on total costs are heavily influenced by inpatient visits because the cost of providing these services is much greater. This can be seen by looking at the mean value of the dependent variable in each case: the mean daily cost of providing inpatient hospital services is approximately \$184,916 per 100,000 population, compared with \$45,583 for outpatient services.

The estimates in the first column suggest that a day over 80°F is associated with a statistically significant \$7,994 increase in hospital costs per 100,000 population. The estimate for the <40°F bin indicates an increase in hospital costs of \$12,176, though these estimates are not statistically different from zero. These estimates can easily be scaled up to represent a particular population. Suppose for instance that the entire population of California (38.8 million) were subjected to one-day mean temperature shocks of either <40°or >80°F: these estimates indicate the total cost of providing hospital services associated with these shocks to be approximately \$4.7 and \$3.1 million dollars, respectively.

6.2 Climate Change Impacts

In this final section, I calculate the California-specific impacts of predicted climate change on hospital usage and costs. In order to make these calculations, I require information on the predicted change that an average person will experience in the number of days per year that fall into each temperature bin. The exact process of calculating predicted temperature changes is described in the Online Appendix, but these predictions are based on the RCP8.5 greenhouse gas concentration scenario, and are made using the Hadley GEM2-ES climate model, both of which are used in the Intergovernmental Panel on Climate Change's (IPCC) Fifth Assessment Report. The historic and predicted number of days per year in each temperature bin, as well as the predicted changes, are summarized in Figure A3. With this information in hand, I multiply the predicted change in the number of days in a given

[23] The rationale for using the national rather than the California-specific cost-to-charge ratio can be seen by looking at the inpatient data: average charges are higher and cost-to-charge ratios are lower in California, so using the California-specific cost-to-charge ratio with estimates of charges based on national data would underestimate the total cost of each visit.

temperature bin per year by the corresponding cumulative effect estimate, then sum across bins.[24]

I estimate the effects of climate change on hospital visits and costs in California to be negligible. The point estimate for the effect of climate change on annual emergency-related hospital visits is small and not statistically different from zero; the 95% confidence interval for the percent change in the annual number of visits is (-1.20%, 1.40%). Similarly, the point estimate for the effect of climate change on annual emergency-related hospital costs is statistically indistinguishable from zero; the 95% confidence interval for the annual percent change in costs is (-1.57%, 0.61%). These negligible climate impacts are not surprising given the result that both cold and hot temperatures are estimated to have negative health effects.

7 Conclusion

This paper uses daily data on the near-universe of ED visits in California between 2005 and 2014 to provide new insight into the link between weather, human behavior and health. I provide some of the first estimates of the temperature-morbidity relationship for a large geographically diverse region using panel fixed-effects methods.

The empirical results indicate that a day under 40°F, relative to a 60-65°F day, is associated with a 6.1% decrease in ED visits on the day of the event. When accounting for visits that occur in the 30 days that follow, however, my estimates indicate a total net increase in ED visits of 11.0%. The contemporaneous decrease is similar for disease types that are classified as deferrable or non-deferrable, suggesting an improvement in actual health outcomes as opposed to a decrease in willingness to seek treatment at a hospital. The increase in ED visits following a cold temperature shock is driven by diseases that are communicable in nature, suggesting that cold temperatures lead to increased susceptibility to communicable disease. It is important to note that while on average there is a total net increase in ED visits due to a cold temperature shock, there is significant heterogeneity in this process such that are some disease types for which cold temperatures actually decreases the incidence of disease (e.g., Injuries, Skin Conditions, and Mental Illness).

Additionally, the results indicate that a day over 80°F leads to a 3.5% increase in ED visits on the day of the event and a total net increase of 5.1% after 30 days, where the additional visits occur in the few days immediately following the shock. The pattern of increased visits in the days that follow is much stronger for disease types that are classified as deferrable, suggesting that individuals do not necessarily seek treatment for an illness at

[24] In practice, the predicted climate impact is calculated as follows: $\sum_J \left(\sum_{h=0}^{30} \beta_{j,t-h}\right) * \Delta\text{Days}_j$. Since this is still a linear combination of coefficients, calculation of the standard errors is straightforward.

the earliest chance.

This study departs from much of the previous literature on the relationship between environmental conditions and health in how behavioral responses are considered. Importantly, I argue that behavioral changes in response to weather can affect health outcomes even if the behavior was not motivated by health concerns. This simple assumption makes a variety of outcomes possible: one is that behavioral responses to environmental conditions need not be protective in nature (i.e., it is possible that the behavioral response to weather can be health-damaging rather than health-improving). Another is that behavioral responses to environmental conditions may go beyond mitigating health damages associated with exposure, and actually decrease the probability of the realization of a negative health outcome relative to the counterfactual in which the weather shock was never experienced (e.g., the disease categories that experienced a net decrease in visits following a cold weather event).

In this paper, I have remained fairly agnostic as to the specific behaviors that drive the effects of temperature on ED visits. While the behaviors discussed here were undertaken in reaction to changes in temperature, the behaviors themselves likely influence health independent of weather conditions. Future research could seek to identify the specific behaviors responsible for the results presented here and determine how such behaviors influence health risks in general.

Figure 1: Contemporaneous and Cumulative Effects

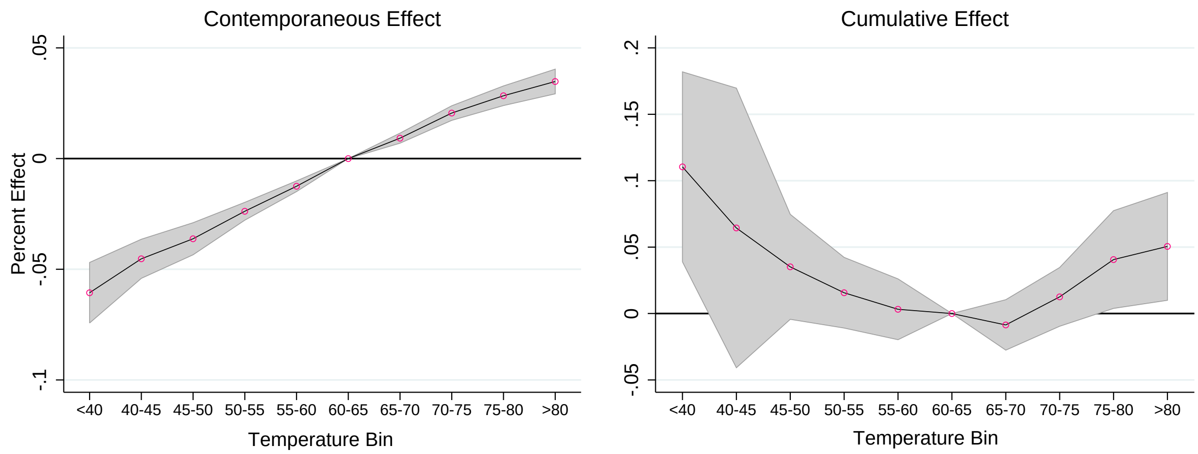

Note – The "Contemporaneous Effect" corresponds to $\beta_{j,t}$ in Equation (4), divided by the mean daily visit rate such that it can be interpreted as a percentage change. The "Cumulative Effect" corresponds to $\sum_{h=0}^{30} \beta_{j,t-h}$, divided by the mean daily visit rate. Shaded regions represent 95% confidence intervals.

Figure 2: Daily Impacts (Coefficient Estimates)

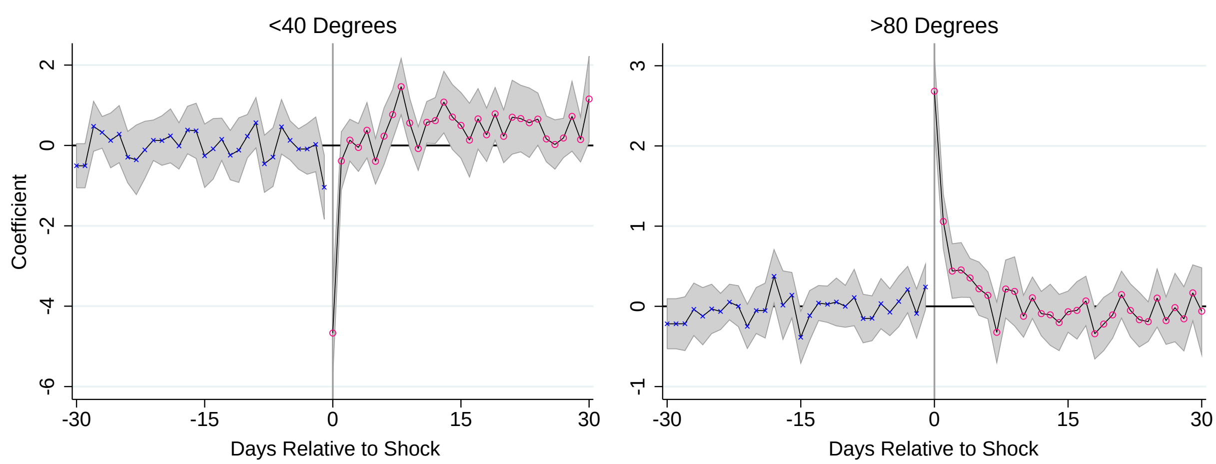

Note – Estimates to the left of zero represent leads (i.e., placebo estimates) and estimates to the right of zero represent lags (i.e., treatment effects). Points on these plots represent coefficient estimates, such that each point on the plot may be interpreted as the impact of a weather event in period $t = 0$ on ED visit rates h periods removed. Estimates are reported in levels for comparison with Figures 3 and 4. Percent effects can be calculated by dividing these level effects by the mean daily visit rate (77.1). Shaded regions represent 95% confidence intervals.

Figure 3: Cumulative Impacts

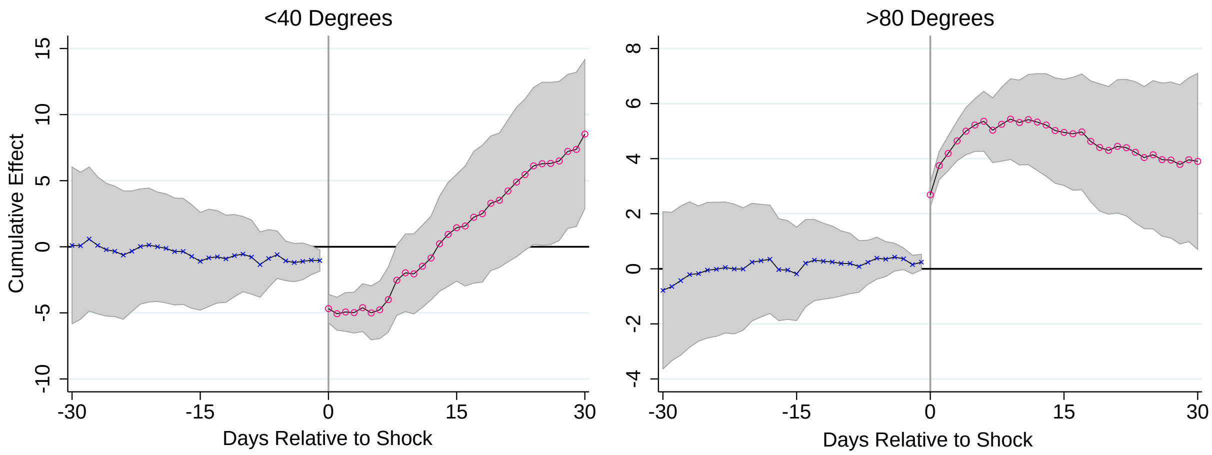

Note – Estimates to the left of zero represent leads (i.e., placebo estimates) and estimates to the right of zero represent lags (i.e., treatment effects). Points on these plots represent the sum of all coefficients for a given temperature bin up to and including the corresponding lag (or lead). For example, the point at which days relative to shock equals 0 represents the contemporaneous effect ($\beta_{j,t}$), and the point at which days relative to shock equals 30 is the 31-day cumulative effect ($\sum_{h=0}^{30} \beta_{j,t-h}$). The cumulative treatment effect estimates include the contemporaneous effect and the cumulative placebo estimates do not. Estimates are reported in levels for comparison with Figures 2 and 4. Percent effects can be calculated by dividing these level effects by the mean daily visit rate (77.1). Shaded regions represent 95% confidence intervals.

Figure 4: Cumulative Impacts by Disease Category

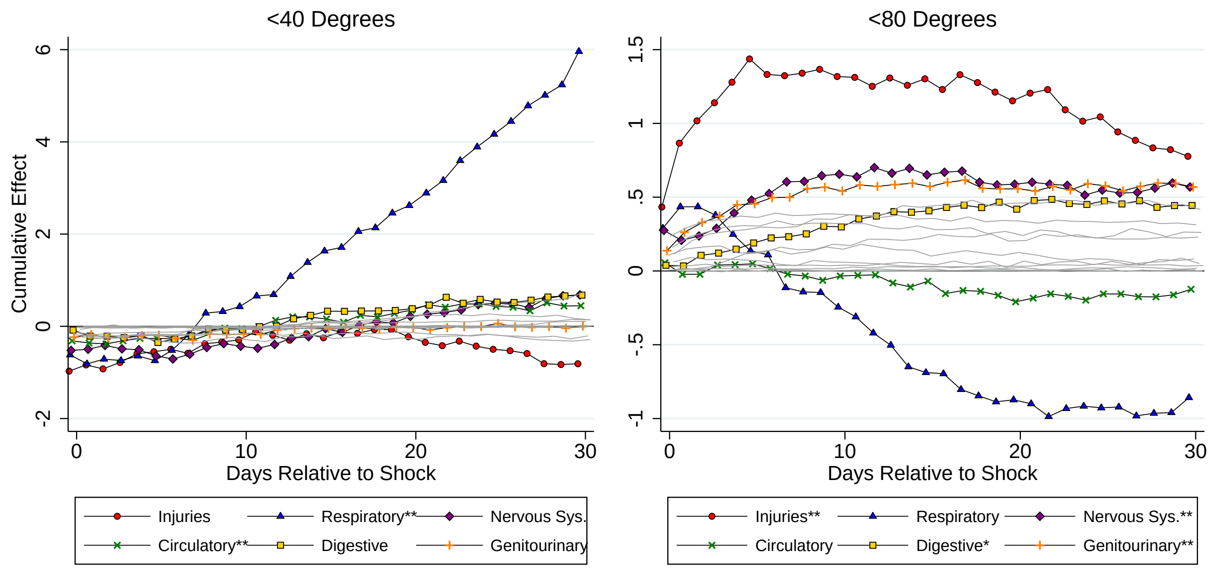

Note – These plots represent cumulative impacts over days relative to the temperature shock for 16 disease categories. The six largest disease categories are highlighted with colored markers as these tend to drive the overall estimates; the estimates for the ten remaining categories are included as grey lines. Estimates are reported in levels such that the sum of level effects across categories roughly equals the total impact (reported in Figure 3). Standard errors are not represented on this figure; stars on each of the highlighted categories (in the legend) denote the level of significance for the 31-day cumulative effects. ** significant at the 1% level. * significant at the 5% level.

Table 1: Summary Statistics

	Number of Visits	Mean	SD
All Visits	94,225,592	77.09	33.96
Age Groups (rates are age-specific):			
Age Under 5	10,322,464	118.38	103.33
Age 5-14	8,105,756	48.12	63.94
Age 15-24	12,911,893	71.31	52.94
Age 25-64	45,855,312	71.34	38.70
Age Over 64	17,030,166	126.38	117.15
CCS Disease Categories:			
Infectious/Parasitic	2,620,942	2.14	2.91
Neoplasms	549,877	0.45	1.21
Endocrine	2,032,303	1.66	2.40
Blood	472,829	0.39	1.11
Mental Illness	3,985,253	3.26	3.76
Nervous System	8,331,257	6.82	5.67
Circulatory	8,031,574	6.57	5.15
Respiratory	12,078,529	9.88	8.39
Digestive	7,381,386	6.04	5.03
Genitourinary	5,690,685	4.66	4.28
Pregnancy	2,540,114	2.08	2.78
Skin	3,208,853	2.63	3.30
Musculoskeletal	5,045,618	4.13	4.20
Congenital	43,224	0.04	0.33
Perinatal	225,341	0.18	0.74
Injuries	19,357,264	15.84	9.06
Temperature Bins:			
<40	-	0.012	0.012
40-45	-	0.032	0.032
45-50	-	0.069	0.069
50-55	-	0.132	0.132
55-60	-	0.176	0.176
60-65	-	0.170	0.170
65-70	-	0.161	0.161
70-75	-	0.110	0.110
75-80	-	0.079	0.079
>80	-	0.060	0.060

Note – The data on hospital visits represents all visits that took place through the emergency department (whether outpatient or inpatient); the mean represents the mean daily visit rate, and is calculated as the number of daily visits per 100,000 residents. For age categories, this is calculated as the number of daily visits per 100,000 age-specific residents. Temperature variables are indicators, and the mean value of these variables should be interpreted as the percentage of days that falls into each category.

Table 2: Baseline Results

	All Emergency		Outpatient Emergency		Inpatient Emergency		Non-Emergency	
	Contemp.	Cumul.	Contemp.	Cumul.	Contemp.	Cumul.	Contemp.	Cumul.
<40	-0.061**	0.110**	-0.064**	0.121**	-0.040**	0.053*	-0.001	-0.008
	(0.007)	(0.036)	(0.008)	(0.042)	(0.005)	(0.025)	(0.010)	(0.036)
40-45	-0.045**	0.064	-0.048**	0.070	-0.029**	0.036	-0.002	0.009
	(0.005)	(0.054)	(0.005)	(0.059)	(0.004)	(0.032)	(0.006)	(0.017)
75-80	0.028**	0.041*	0.029**	0.047**	0.026**	0.004	0.004	0.001
	(0.002)	(0.019)	(0.003)	(0.020)	(0.002)	(0.014)	(0.003)	(0.011)
>80	0.035**	0.051*	0.035**	0.054*	0.034**	0.033	0.007	0.015
	(0.003)	(0.021)	(0.003)	(0.022)	(0.004)	(0.017)	(0.005)	(0.015)
Mean Dep. Var.	77.1		65.0		12.1		14.7	
# Admissions	94,225,592		79,491,760		14,733,830		18,020,504	
N	3,905,239		3,905,239		3,905,239		3,905,239	
Zip-by-Week	X		X		X		X	
County-by-Year	X		X		X		X	
Day-of-Week	X		X		X		X	
Holidays	X		X		X		X	

Note – The results labelled "All Emergency" represent the baseline estimates. The "Outpatient Emergency" and "Inpatient Emergency" are components of the baseline estimates; "Non-Emergency" visits represent inpatient visits that did not occur through the ED, and are not a component of the baseline estimates. The set of controls (and fixed effects) listed here are the same set of controls used throughout the remainder of the analysis; regression results in the Online Appendix probe the robustness of these estimates to alternative sets of controls. Standard errors are two-way clustered at the county and year-month levels. All regressions are weighted by total zip-code population. Regressions are estimated in levels, but reported in percent changes (the corresponding coefficients or sum of coefficients divided by the reported mean dependent variable). ** significant at the 1% level. * significant at the 5% level.

Table 3: Cumulative Effect Specification

	30 Days (Baseline)	40 Days	50 Days	60 Days	Monthly Model
<40	0.110**	0.150**	0.167**	0.170**	0.153**
	(0.036)	(0.042)	(0.049)	(0.056)	(0.058)
40-45	0.064	0.083	0.100	0.089	0.088
	(0.054)	(0.058)	(0.057)	(0.056)	(0.065)
75-80	0.041*	0.045*	0.044	0.041	0.047
	(0.019)	(0.022)	(0.023)	(0.026)	(0.028)
>80	0.051*	0.049*	0.048	0.042	0.046
	(0.021)	(0.023)	(0.026)	(0.029)	(0.033)
Mean Dep. Var.	77.1	77.1	77.1	77.1	2,345.6
# Admissions	94,009,920	94,009,920	94,009,920	94,009,920	94,330,160
N	3,891,199	3,891,199	3,891,199	3,891,199	128,575

Note – Estimates in the first four columns represent cumulative effects from models estimated with various lag lengths. Sample sizes are slightly smaller as an additional 30 days had to be dropped from the beginning of the sample to accommodate 60 lags in temperature. Column five represents estimates at the monthly level with zip-by-month and county-by-year fixed effects. Standard errors are two-way clustered at the county and year-month levels. All regressions are weighted by total zip-code population. Regressions are estimated in levels, but reported in percent changes. Estimates for regressions at the daily level are the level effects divided by the mean daily visit rate; estimates for regressions at the monthly level are the level effects divided by the mean monthly visit rate and multiplied by 30.5 such that they are comparable with the estimates using daily data. ** significant at the 1% level. * significant at the 5% level.

Table 4: Age Heterogeneity

	Under 5		5-14		15-24		25-64		Over 64	
	Contemp.	Cumul.	Contemp.	Cumul.	Contemp.	Cumul.	Contemp.	Cumul.	Contemp.	Cumul.
<40	-0.070**	0.277**	-0.117**	0.222*	-0.061**	0.102*	-0.049**	0.069	-0.056**	0.115**
	(0.014)	(0.109)	(0.017)	(0.108)	(0.009)	(0.049)	(0.007)	(0.036)	(0.007)	(0.043)
40-45	-0.064**	0.137	-0.083**	0.128	-0.037**	0.043	-0.037**	0.034	-0.049**	0.100**
	(0.010)	(0.108)	(0.011)	(0.121)	(0.005)	(0.060)	(0.005)	(0.043)	(0.006)	(0.036)
75-80	0.048**	0.074	0.024**	0.098*	0.029**	0.037	0.022**	0.035	0.034**	0.007
	(0.008)	(0.051)	(0.007)	(0.042)	(0.003)	(0.019)	(0.003)	(0.020)	(0.003)	(0.012)
>80	0.064**	0.086	0.021**	0.095*	0.039**	0.034	0.027**	0.040	0.037**	0.027
	(0.012)	(0.053)	(0.008)	(0.047)	(0.004)	(0.024)	(0.005)	(0.024)	(0.004)	(0.018)
Mean Dep. Var.	118.4		48.1		71.3		71.3		126.4	
# Admissions	10,322,464		8,105,756		12,911,893		45,855,312		17,030,166	
N	3,905,239		3,905,239		3,905,239		3,905,239		3,905,239	

Note – The visit rate in these regressions is calculated by dividing the total number of visits by the age-specific population in each zip code. Standard errors are two-way clustered at the county and year-month levels. All regressions are weighted by total zip-code population. Regressions are estimated in levels, but reported in percent changes (the level effect divided by the reported mean dependent variable). ** significant at the 1% level. * significant at the 5% level.

Table 5: Disease Deferrability

	More Deferrable		Less Deferrable	
	Contemp.	Cumul.	Contemp.	Cumul.
<40	-0.065**	0.170**	-0.058**	0.098**
	(0.011)	(0.041)	(0.006)	(0.021)
40-45	-0.044**	0.115	-0.042**	0.050
	(0.007)	(0.073)	(0.004)	(0.037)
75-80	0.026**	0.064	0.030**	0.031*
	(0.003)	(0.026)	(0.003)	(0.015)
>80	0.030**	0.065*	0.036**	0.039*
	(0.003)	(0.027)	(0.003)	(0.018)
Mean Dep. Var.	13.9		10.6	
# Admissions	16,939,112		12,926,923	
N	3,905,239		3,905,239	

Note – The classification of ICD codes into more and less deferrable categories is based on the classification developed by Billings et al. (2000), which was constructed by having a team of ED physicians classify thousands of individual cases based on whether the specific case required emergency care. Standard errors are two-way clustered at the county and year-month levels. All regressions are weighted by total zip-code population. Regressions are estimated in levels, but reported in percent changes (the level effect divided by the reported mean dependent variable). ** significant at the 1% level. * significant at the 5% level.

Table 6: Healthcare Cost Impacts

	All Emergency-Related	Emergency Inpatient	Emergency Outpatient
<40	12,176	8,542	3,406*
	(10,877)	(10,202)	(1,422)
40-45	11,364	8,543	2,573
	(9,998)	(8,705)	(1,884)
75-80	1,639	-499	2,086**
	(3,732)	(3,270)	(803)
>80	7,994*	5,498	2,301**
	(3,773)	(3,310)	(882)
Mean Dep. Var.	$230,846	$184,916	$45,583
# Visits	94,225,592	14,733,830	78,745,280
N	3,905,239	3,905,239	3,905,239

Note – The dependent variable in these regressions is dollars per day, per 100,000 population. A mean of $45,583 in the third column implies that the average daily cost of providing ED outpatient services is $45,583 per 100,000 population. Likewise, the average daily cost of providing all hospital services is $230,846 per 100,000 population. Standard errors are two-way clustered at the county and year-month levels. All regressions are weighted by total zip-code population. Regression estimates are reported in levels (rather than percent changes) such that the interpretation is as follows: the total change in hospital costs attributable to a given temperature shock per 100,000 individuals in the population. ** significant at the 1% level. * significant at the 5% level.

References

Adda, J. (2015). Economic activity and the spread of viral disease: Evidence from high-frequency data. *Quarterly Journal of Economics*, forthcoming.

Armstrong, B. (2006). Models for the relationship between ambient temperature and daily mortality. *Epidemiology*, 17(6):624–631.

Baker, M., McNicholas, A., Garrett, N., Jones, N., Stewart, J., Koberstein, V., and Lennon, D. (2000). Household crowding a major risk factor for epidemic meningococcal disease in auckland children. *The Pediatric Infectious Disease Journal*, 19(10):983–990.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159.

Barreca, A. I. (2012). Climate change, humidity, and mortality in the united states. *Journal of Environmental Economics and Management*, 63(1):19–34.

Barreca, A. I. and Shimshack, J. P. (2012). Absolute humidity, temperature, and influenza mortality: 30 years of county-level evidence from the united states. *American journal of epidemiology*, 176(suppl 7):S114–S122.

Basu, R. and Samet, J. M. (2002). Relation between elevated ambient temperature and mortality: A review of the epidemiologic evidence. *Epidemiologic Reviews*, 24(2):190–202.

Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal*, pages 493–517.

Billings, J., Parikh, N., and Mijanovich, T. (2000). Emergency department use: The new york story.

Burström, B., Diderichsen, F., and Smedman, L. (1999). Child mortality in stockholm during 1885–1910: the impact of household size and number of children in the family on the risk of death from measles. *American Journal of Epidemiology*, 149(12):1134–1141.

Card, D., Dobkin, C., and Maestas, N. (2009). Does medicare save lives? *The Quarterly Journal of Economics*, 124(2):597–636.

Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of econometrics*, 92(1):1–45.

de Pablo, F., Tomás, C., Rivas, L., and Diego, L. (2009). Winter circulation weather types and hospital admissions for cardiovascular, respiratory and digestive diseases in salamanca, spain. *International Journal of Climatology*, 29:1692–1703.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? the new climate–economy literature. *Journal of Economic Literature*, 52(3):740–798.

Dennehy, P. H. (2011). Viral gastroenteritis in children. *The Pediatric Infectious Disease Journal*, 30(1):63–64.

Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, pages 152–185.

Deschênes, O. and Moretti, E. (2009). Extreme weather events, mortality, and migration. *The Review of Economics and Statistics*, 91(4):659–681.

Dobkin, C. (2003). Hospital staffing and inpatient mortality. *Unpublished Working Paper*.

Dowell, S. F. and Ho, M. S. (2004). Seasonality of infectious diseases and severe acute respiratory syndrome–what we don't know can hurt us. *The Lancet Infectious Diseases*, 4(11):704–708.

Doyle, J., Graves, J., Gruber, J., and Kleiner, S. (2015). Measuring returns to hospital care: Evidence from ambulance referral patterns. *The Journal of Political Economy*, 123(1):170.

Glass, R. I., Parashar, U. D., and Estes, M. K. (2009). Norovirus gastroenteritis. *New England Journal of Medicine*, 361(18):1776–1785.

Graff-Zivin, J. and Neidell, M. (2009). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management*, 58(2):119–128.

Graff-Zivin, J. and Neidell, M. (2013). Environment, health, and human capital. *Journal of Economic Literature*, 51(3):689–730.

Graff-Zivin, J. and Neidell, M. J. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 31(1).

Grossman, M. (1972). On the concept of health capital and the demand for health. *The Journal of Political Economy*, pages 223–255.

Houser, T., Hsiang, S., Kopp, R., Larsen, K., Delgado, M., Jina, A., Mastrandrea, M., Mohan, S., Muir-Wood, R., Rasmussen, D., et al. (2014). American climate prospectus.

Isen, A., Rossin-Slater, M., and Walker, R. (2015). Heat and long-run human capital formation. *Unpublished Manuscript*.

Karlsson, M. and Ziebarth, N. (2016). Population health effects and health-related cost of heat events: Comprehensive evidence from germany. *Working Paper*.

Knowlton, K., Rotkin-Ellman, M., King, G., Margolis, H. G., Smith, D., Solomon, G., Trent, R., English, P., et al. (2009). The 2006 california heat wave: impacts on hospitalizations and emergency department visits. *Environmental Health Perspectives*, 117(1):61–67.

Lofgren, E., Fefferman, N. H., Naumov, Y. N., Gorski, J., and Naumova, E. N. (2007). Influenza seasonality: underlying causes and modeling theories. *Journal of Virology*, 81(11):5429–5436.

Lowen, A. C., Mubareka, S., Steel, J., and Palese, P. (2007). Influenza virus transmission is dependent on relative humidity and temperature. *PLoS Pathog*, 3(10):e151.

Moretti, E. and Neidell, M. (2011). Pollution, health, and avoidance behavior evidence from the ports of los angeles. *Journal of Human Resources*, 46(1):154–175.

Neidell, M. (2009). Information, avoidance behavior, and health the effect of ozone on asthma hospitalizations. *Journal of Human resources*, 44(2):450–478.

Oxford, J., Sefton, A., Jackson, R., Innes, W., Daniels, R., and Johnson, N. (2002). World war i may have allowed the emergence of spanish influenza. *The Lancet Infectious Diseases*, 2(2):111–114.

Portier, C. J., Tart, K. T., Carter, S. R., Dilworth, C. H., Grambsch, A. E., Gohlke, J., Hess, J., Howard, S., Luber, G., Lutz, J., et al. (2013). A human health perspective on climate change: a report outlining the research needs on the human health effects of climate change. *Journal of Current Issues in Globalization*, 6(4):621.

Souza, L. S. d. F., Ramos, E. A. G., Carvalho, F. M., Guedes, V. M. C. R., Souza, L. S., Rocha, C. M., Soares, A. B., Velloso, L. d. F., Macedo, I. S., Moura, F. E. A., et al. (2003). Viral respiratory infections in young children attending day care in urban northeast brazil. *Pediatric Pulmonology*, 35(3):184–191.

Stock, J. H. and Watson, M. W. (2003). *Introduction to econometrics*, volume 104. Addison Wesley Boston.

Stoecker, C., Sanders, N. J., Barreca, A., et al. (2015). Success is something to sneeze at: Influenza mortality in regions that send teams to the super bowl. *American Journal of Health Economics*, forthcoming.

Taubman, S. L., Allen, H. L., Wright, B. J., Baicker, K., and Finkelstein, A. N. (2014). Medicaid increases emergency-department use: evidence from oregon's health insurance experiment. *Science*, 343(6168):263–268.

Ye, X., Wolff, R., Yu, W., Vaneckova, P., Pan, X., and Tong, S. (2012). Ambient temperature and morbidity: a review of epidemiological evidence. *Environmental Health Perspectives*, 120(1):19–28.

Figure A1: Cumulating Effects by Disease Category: <40 °F

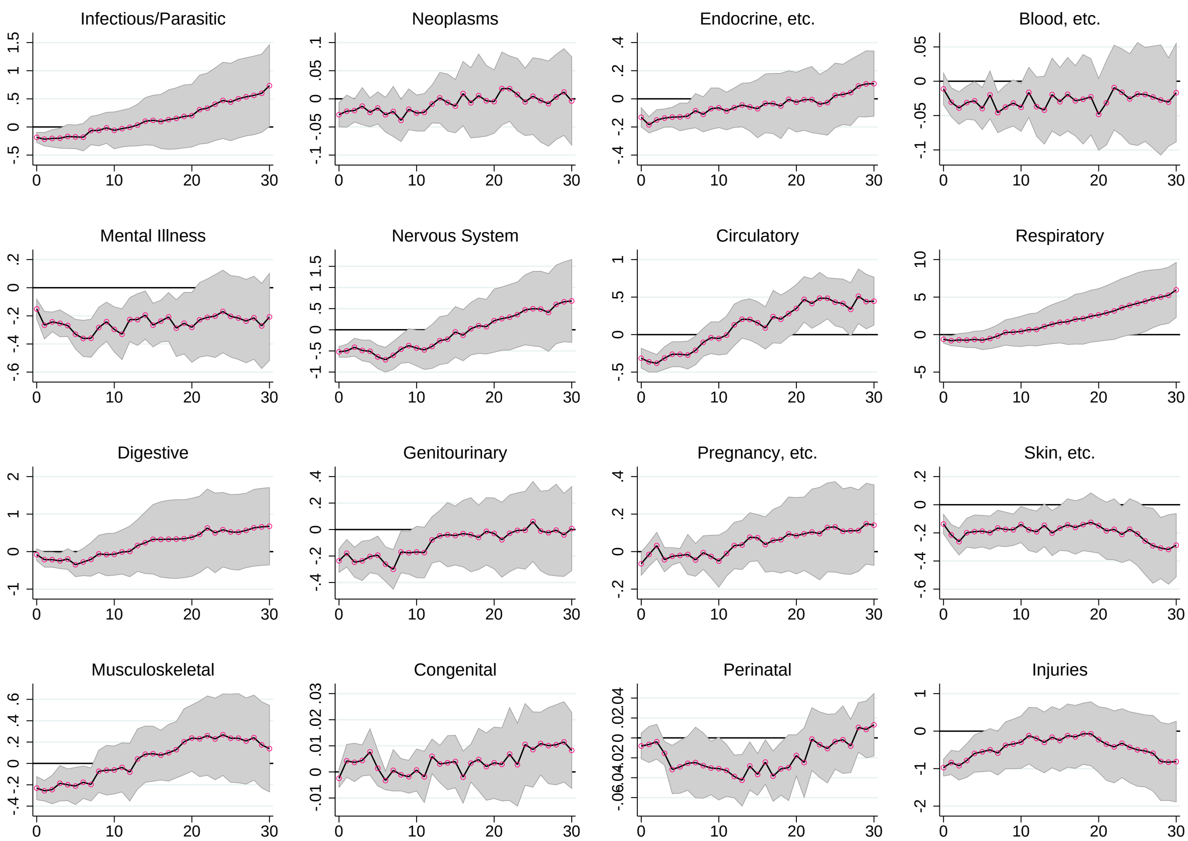

Note – These are cumulating effects graphs as in Figure 3, for the coldest temperature bin only. Each graph represents one of sixteen different CCS categories, at their highest level of aggregation. Note that each plot is on a different scale and that estimates are reported in levels rather than percentage terms (i.e., estimates are not divided by the mean visit rate). Shaded regions represent 95% confidence intervals.

Figure A2: Cumulating Effects by Disease Category: >80 °F

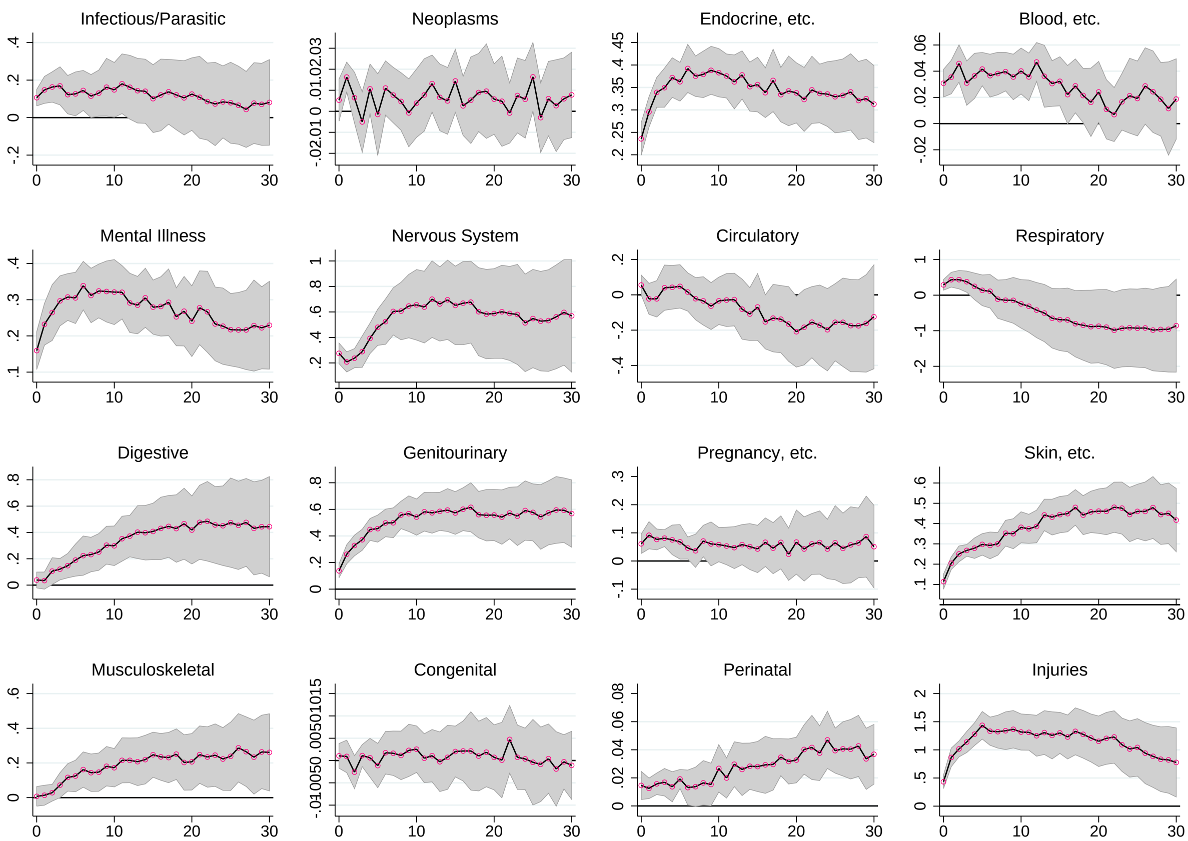

Note – These are cumulating effects graphs as in Figure 3, for the hottest temperature bin only. Each graph represents one of sixteen different CCS categories, at their highest level of aggregation. Note that each plot is on a different scale and that estimates are reported in levels rather than percentage terms (i.e., estimates are not divided by the mean visit rate). Shaded regions represent 95% confidence intervals.

Figure A3: Climate Predictions

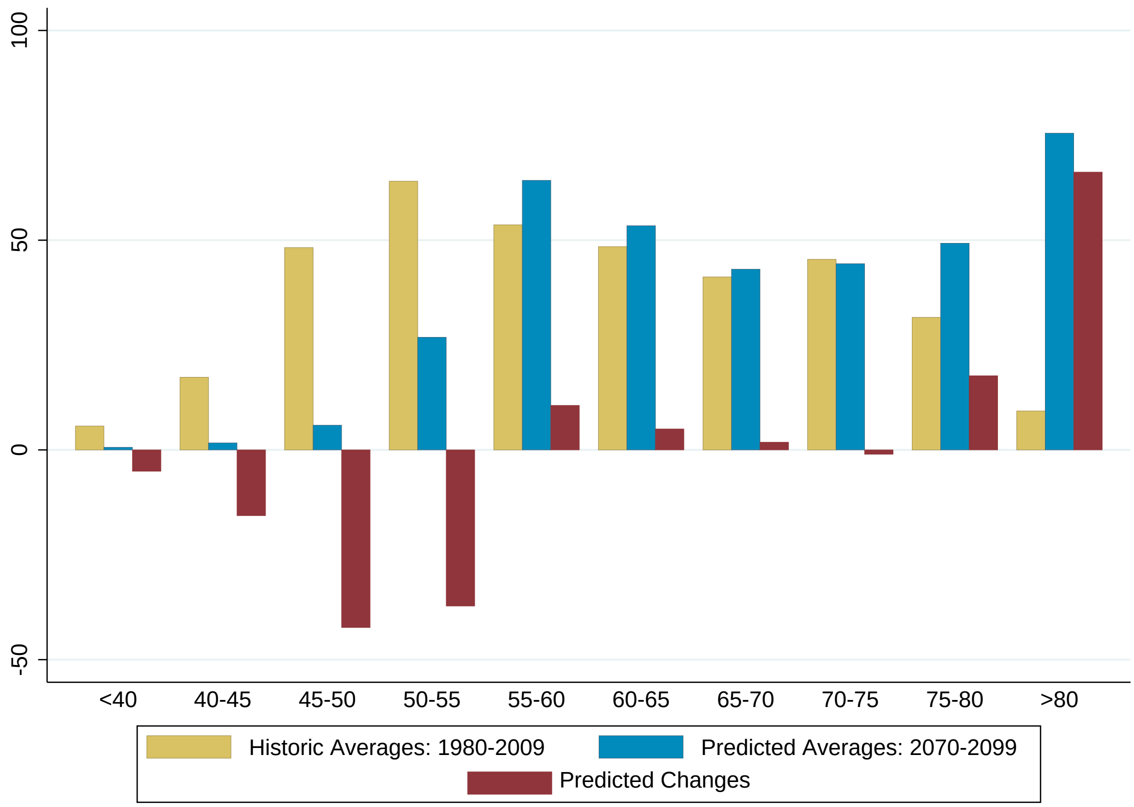

Note – Bars represent the number of days per year in each temperature bin. Predictions are based on the RCP8.5 greenhouse gas concentration scenario adopted by the IPCC's Fifth Assessment Report and are made using the Hadley GEM2-ES climate model. All calculations are made for zip-codes in California and are population-weighted.

Table A1: Specific Disease Category Counts (Part 1)

Code	Count	Label
1		**Infectious Diseases**
1.1	901,078	Bacterial infection
1.2	201,056	Mycoses
1.3	1,209,989	Viral infection
1.4	103,208	Other infections; including parasitic
1.5	75,251	Immunizations and screening for infectious disease
2		**Neoplasms**
2.1	40,986	Colorectal cancer
2.2	58,737	Other gastrointestinal cancer
2.3	55,396	Cancer of bronchus; lung
2.4	3,615	Cancer of skin
2.5	10,827	Cancer of breast
2.6	8,579	Cancer of uterus and cervix
2.7	9,070	Cancer of ovary and other female genital organs
2.8	8,975	Cancer of male genital organs
2.9	14,051	Cancer of urinary organs
2.1	55,331	Cancer of lymphatic and hematopoietic tissue
2.11	33,807	Cancer; other primary
2.12	109,199	Secondary malignancies
2.13	5,909	Malignant neoplasm without specification of site
2.14	28,907	Neoplasms of unspecified nature or uncertain behavior
2.15	3,634	Maintenance chemotherapy; radiotherapy
2.16	88,871	Benign neoplasms
3		**Endocrine, etc.**
3.1	37,676	Thyroid disorders
3.2	250,246	Diabetes mellitus without complication
3.3	664,435	Diabetes mellitus with complications
3.4	73,577	Other endocrine disorders
3.5	6,598	Nutritional deficiencies
3.6	4,700	Disorders of lipid metabolism
3.7	104,749	Gout and other crystal arthropathies
3.8	789,642	Fluid and electrolyte disorders
3.9	2,143	Cystic fibrosis
3.1	1,050	Immunity disorders
3.11	62,810	Other nutritional; endocrine; and metabolic disorders
4		**Blood, etc.**
4.1	343,085	Anemia
4.2	57,244	Coagulation and hemorrhagic disorders
4.3	47,963	Diseases of white blood cells
4.4	4,195	Other hematologic conditions
5		**Mental Illness**
5.1	47,991	Adjustment disorders
5.2	825,655	Anxiety disorders
5.3	27,742	Attention deficit conduct and disruptive behavior disorders
5.4	137,583	Delirium dementia and amnestic and other cognitive disorders
5.5	26,065	Developmental disorders
5.6	6,551	Disorders usually diagnosed in infancy childhood or adolescence
5.7	4,479	Impulse control disorders not elsewhere classified
5.8	649,198	Mood disorders
5.9	8,110	Personality disorders
5.1	514,477	Schizophrenia and other psychotic disorders
5.11	900,617	Alcohol-related disorders
5.12	360,743	Substance-related disorders
5.13	82,198	Suicide and intentional self-inflicted injury
5.14	74,619	Screening and history of mental health and substance abuse codes
5.15	147,360	Miscellaneous mental disorders
6		**Nervous System**
6.1	55,157	Central nervous system infection
6.2	85,766	Hereditary and degenerative nervous system conditions
6.3	10,801	Paralysis
6.4	909,929	Epilepsy; convulsions
6.5	2,100,734	Headache; including migraine
6.6	167,981	Coma; stupor; and brain damage
6.7	943,847	Eye disorders
6.8	2,822,769	Ear conditions
6.9	1,127,338	Other nervous system disorders
7		**Circulatory System**
7.1	665,669	Hypertension
7.2	5,860,823	Diseases of the heart
7.3	718,609	Cerebrovascular disease
7.4	273,046	Diseases of arteries; arterioles; and capillaries
7.5	374,254	Diseases of veins and lymphatics

Table A2: Specific Disease Category Counts (Part 2)

Code	Count	Label
8		**Respiratory System**
8.1	6,776,940	Respiratory infections
8.2	1,219,839	Chronic obstructive pulmonary disease, etc.
8.3	1,476,767	Asthma
8.4	137,406	Aspiration pneumonitis; food/vomitus
8.5	145,454	Pleurisy; pneumothorax; pulmonary collapse
8.6	239,300	Respiratory failure; insufficiency; arrest (adult)
8.7	11,265	Lung disease due to external agents
8.8	1,454,696	Other lower respiratory disease
8.9	629,809	Other upper respiratory disease
9		**Digestive System**
9.1	316,558	Intestinal infection
9.2	729,476	Disorders of teeth and jaw
9.3	167,275	Diseases of mouth; excluding dental
9.4	939,676	Upper gastrointestinal disorders
9.5	214,809	Abdominal hernia
9.6	1,155,086	Lower gastrointestinal disorders
9.7	709,491	Biliary tract disease
9.8	192,433	Liver disease
9.9	305,828	Pancreatic disorders (not diabetes)
9.1	492,413	Gastrointestinal hemorrhage
9.11	961,048	Noninfectious gastroenteritis
9.12	1,100,548	Other gastrointestinal disorders
10		**Genitourinary System**
10.1	4,003,703	Diseases of the urinary system
10.2	356,935	Diseases of male genital organs
10.3	1,191,579	Diseases of female genital organs
11		**Pregnancy, etc.**
11.1	6,065	Contraceptive and procreative management
11.2	241,534	Abortion-related disorders
11.3	1,978,024	Complications mainly related to pregnancy
11.4	27,611	Indications for care in pregnancy; labor; and delivery
11.5	21,723	Complications during labor
11.6	167,671	Other complications of birth;
11.7	52,434	Normal pregnancy and/or delivery

Code	Count	Label
12		**Skin, etc.**
12.1	2,333,937	Skin and subcutaneous tissue infections
12.2	107,930	Other inflammatory condition of skin
12.3	74,757	Chronic ulcer of skin
12.4	646,404	Other skin disorders
13		**Musculoskeletal System**
13.1	47,175	Infective arthritis and osteomyelitis (no TB or STD)
13.2	1,106,749	Non-traumatic joint disorders
13.3	2,219,077	Spondylosis; intervertebral disc disorders; etc.
13.4	1,058	Osteoporosis
13.5	41,176	Pathological fracture
13.6	17,678	Acquired deformities
13.7	26,547	Systemic lupus erythematosus, etc.
13.8	1,304,775	Other connective tissue disease
13.9	136,988	Other bone disease
14		**Congenital Anomalies**
14.1	6,919	Cardiac and circulatory congenital anomalies
14.2	10,295	Digestive congenital anomalies
14.3	7,995	Genitourinary congenital anomalies
14.4	1,880	Nervous system congenital anomalies
14.5	14,955	Other congenital anomalies
15		**Diseases Related to the Perinatal Period**
15.1	1,266	Liveborn
15.2	347	Short gestation; low birth weight; etc.
15.3	269	Intrauterine hypoxia and birth asphyxia
15.4	254	Respiratory distress syndrome
15.5	43,430	Hemolytic jaundice and perinatal jaundice
15.6	1,196	Birth trauma
15.7	175,089	Other perinatal conditions
16		**Injuries**
16.1	415,247	Joint disorders and dislocations; trauma-related
16.2	3,068,594	Fractures
16.3	10,285	Spinal cord injury
16.4	481,004	Intracranial injury
16.5	166,835	Crushing injury or internal injury
16.6	4,230,354	Open wounds
16.7	3,488,448	Sprains and strains
16.8	3,585,265	Superficial injury; contusion
16.9	264,802	Burns
16.1	931,078	Complications
16.11	675,666	Poisoning
16.12	1,987,362	Other injuries and conditions due to external causes