Ambient Temperature and Occupational Accidents

by

Lucy Page

Stephen C. Sheppard and Peter L. Pedroni, Advisors

Abstract

I use plausibly exogenous variation in temperature within counties to estimate the causal impacts of temperature on the rate of accidents across the United States from 1990 to 2010. I focus on a set of "temperature-sensitive" industries, where workers are likely to be exposed to outdoor weather conditions. I find that accident rates rise significantly at high temperatures and fall significantly at low temperatures. Based on my estimates for the short-run impacts of temperature, I project that climate change may cause between about 33,600 and 102,400 additional accidents in temperature-sensitive industries per year between 2070 and 2099. These additional accidents could pose costs of between $750 million and $2.30 billion per year between 2070 and 2099. These estimates may overstate the true long-run impacts of temperature on accident rates if individuals can adapt to these impacts, or they may understate these long-run impacts if more prolonged climatic change intensifies these risks by more than do the transient weather fluctuations in my data.

Department of Economics

Williams College

Williamstown, MA

May 11, 2016

Contents

1 Introduction

Climate change is expected to have a wide range of impacts around the world, from changing the incidence and geographic range of vector- and water-borne diseases in Africa to affecting the tourism industry by propelling tourists to higher altitudes and latitudes (IPCC WGII, 2014). To maximize economic efficiency, policies aimed at reducing greenhouse gas emissions should equate the marginal costs and benefits of these reductions; to do so, they should rely on a comprehensive understanding of climate change's extensive and varied impacts. Unfortunately, a complete treatment of these impacts is impossible; not only can we not pinpoint all of the multitude of possible effects of a changing climate, but projections of possible impacts rely on climate models that remain uncertain. While we may never achieve the perfect knowledge required for ideal climate policy, however, we can do our best to inform policy design with as much information on the impacts of climate change as is possible.

Currently, an active economic literature is working to estimate the impacts of climate change on a range of economic outcomes. So far, this work has estimated the effect of climate change on health, through deaths from air quality, deaths from extreme temperature, loss of labor productivity, and damages from water quality; on infrastructure, like bridges, roads, urban drainage, and coastal property; on the demand and supply of electricity; on water resources, through inland flooding, drought, and water shortages; on agriculture and forestry; and on ecosystem services, through coral reefs, shellfish, freshwater fish, wildfire, and carbon storage (US EPA, 2015). However, large gaps in our knowledge remain. Climate change will affect all sectors of the global economy, as well as ecosystems and other non-market assets that are difficult to quantify, and many of these sectors have been studied only partially, if at all. In this analysis, I add to this literature by studying the potential impacts of climate change on the incidence of workplace accidents in temperature-sensitive industries in the United States. To date, no studies have estimated the impacts of climate change on worker safety, though there are several public health studies that have evaluated the relationship between temperature and accident incidence in Quebec (Adam-Poupart et al, 2015) and Adelaide, Australia (Xiang et al., 2014b, 2014c).

The occupational health effects of increased accident risk could be a significant impact of climate change in the United States and around the world. As temperatures rise around the country under a changing climate, one might expect that more frequent extreme temperature highs would pose a risk to laborers working outside in industries like construction, utility services, forestry, and agriculture. At the same time, milder winters might reduce these risks. Workplace injuries impose a significant cost on the United States economy, so even a modest change in the incidence of these accidents could be economically significant. Indeed, a 2011 study published by the National Institute for Occupational Safety and Health (NIOSH) reported that fatal occupational injuries in the United States were associated with a total cost of over $53 billion between 1992 and 2002. This figure corresponds to an average of about 6,125 fatal occupational accidents each year over that period. These costs are both direct and indirect, including medical expenses, lost production, and lost future wages. Leigh (2011) finds that fatal and nonfatal occupational injuries and illnesses in the US inflicted total costs of about $250 billion in 2007, with about $192 billion of this cost associated with accidents. Thus, any changes in the risk of accidents could have important economic impacts.

In this research, I investigate the impact of temperature on the incidence of workplace accidents in a panel of counties across the United States from 1990 through 2010, focusing on industries where workers are plausibly exposed to outdoor temperatures, like construction, agriculture, forestry, and utilities servicing. This analysis draws both from the recent body of work on the impacts of climate change and from this work in public health. Like much of the recent research on climate impacts, I use a suite of geographic and temporal fixed effects to identify the impacts of short-term, plausibly exogenous local variation in weather (Dell et al., 2014). This identification strategy allows robust causal inference of the impacts of temperature. Like other research on climate impacts, I estimate this relationship across the contiguous United States and using a series of ten-degree temperature bins, improving on the city-specific analysis and strict functional form restrictions of previous public health research in this area. However, I also draw from the public health literature by looking closely at the daily impact of temperature on accident incidence and by using Poisson regression, which is ideal for modeling the arrival of random events, like workplace accidents.

Thus, I present the first analysis of accident incidence and temperature to make use of a broad geographic spatial scale and to flexibly model the impacts of temperature, while retaining a framework of daily-level Poisson analysis that is well-suited to analysis of accident counts. I conduct this analysis using a panel of weather, accident, and employment data for a panel of daily-level data for about 2,500 counties across the contiguous United States from 1990 to 2010. This dataset includes records for 71,218 occupational injuries and fatalities from the Occupational Safety and Health Administration (OSHA), daily-level temperature and precipitation data from the National Weather Service cooperative station network and from the North America Land Data Assimilation System, and monthly-level employment data from the Bureau of Labor Statistics Quarterly Census of Employment and Wages.

I find a statistically significant relationship between daily maximum temperature and accident rate, where accident rate rises at high temperature extremes and falls at low temperature extremes. At high temperatures, I find that accident rate increases by 8.2% on days with maximum temperature between 90 and 100°F and by 30.0% on days with maximum temperature over 100°F, both relative to a day with maximum temperature between 60 and 70°F. At low temperatures, accident rate falls by 21.0% on days with maximum temperature between 10 and 20°F and by 30.1% on days with maximum temperature between 0 and 10°F, both relative to a day with maximum temperature between 60 and 70°F. These impacts are heterogeneous across the United States, but this variation does not appear to be systematically related to historical climate conditions. Next, it appears that both extremely high temperatures on the preceding day and temperature spikes increase accident rate today. It is important to note that my estimates for the impacts of temperature incorporate both the impact of temperature on the risk of accidents for a given duration and intensity of work and behavioral changes in response to temperature, or adaptation. In future research, I hope to better disentangle the relative contributions of these two forces.

I take preliminary steps to project these results into estimates of the impacts of climate change associated with changes in the incidence of workplace accidents. Based on my estimates for the short-run impacts of temperature, I project that climate change may cause between about 33,600 and 102,400 additional accidents in temperature-sensitive industries per year between 2070 and 2099. These additional accidents could pose costs of between $750 million and $2.30 billion per year between 2070 and 2099. It is important to interpret these results with caution. My analysis identifies the impacts of weather fluctuations on accident incidence, where we can think of weather as a particular draw from a climate distribution. The impacts of these short-run fluctuations in weather are not necessarily analogous to the impacts of climate change on worker safety, where climate change occurs gradually over a long timescale. In particular, my estimates may overstate the impacts of climate change if adaptation occurs in the long-run, perhaps through the development of new worker safety technologies or regulations. On the other hand, my estimates may understate the impacts of climate change if these impacts are characterized by intensification, where prolonged periods of heat have more drastic impacts on accident rate than are revealed in the impacts of transient weather fluctuations (Dell et al, 2014).

The rest of the paper is organized as follows. Section 2 reviews relevant literature on the relationship between temperature and occupational safety. Section 3 describes my data sources and reports summary statistics, and Section 4 outlines my econometric strategy. Section 5 then presents my primary results, and Section 6 gives extensions of this primary analysis. Section 7 outlines preliminary analysis of the role of adaptation in my results, Section 8 gives directions for future research, and Section 9 concludes the paper.

2 Literature Review

2.1 Climate Change and Industry

A large body of research has worked to estimate the effects of temperature on aggregate industrial output. To date, much of this research has focused on developing countries. For example, Hsiang (2010) estimates the impacts of temperature and cyclones on output in Caribbean countries between 1970 and 2006. He uses annual data for 28 countries, simultaneously controlling for precipitation, countries' average production and temperature, country-industry time trends in production, and year-level shocks in regional-industry production. Thus, Hsiang identifies the impact of temperature on output using plausibly exogenous local year-to-year fluctuations in weather. He finds that unusually hot periods significantly reduce output, with losses of about 2.4% of nonagricultural output per 1°C increase; these impacts are driven by productivity losses in mining and utilities, wholesale, retail, restaurants, and hotels, and other service-related industries.

Similarly, Dell, Jones, and Olken (2012a) and Jones and Olken (2010) estimate the impacts of temperature on aggregate output in a global sample. First, Dell, Jones, and Olken (2012a) study industrial value-added output across 125 countries between 1950 and 2003, and they find that increases in average temperature reduce output by 2% per increase of 1°C in poor countries. Then, Jones and Olken (2010) use trade data from wealthy countries to estimate the impacts of temperature on output in a range of narrow sectors. They find that warming of 1°C reduces exports from poor countries by about 2.4%. Sector-level estimates show that temperature significantly reduces output in twenty of the sixty-six two-digit export categories considered, including manufactured goods ranging from wood, metal, and rubber manufactures to footwear. Like Hsiang's (2010) work on weather and output, each of these two papers uses a suite of geographic and temporal fixed effects to identify the impacts of plausibly exogenous year-to-year fluctuations in temperature and precipitation.

While these studies find that temperature significantly impacts productivity in developing countries, Dell, Jones, and Olken (2012a) do not find evidence that this relationship holds in wealthy countries. Burke, Hsiang, and Miguel (2015) refute this apparent discrepancy by accounting for non-linearities in the impacts of temperature on aggregate macroeconomic productivity. Using country and time fixed effects and country-specific quadratic trends in growth rates, the authors examine the potentially non-linear impacts of country-specific annual deviations from temperature and precipitation trends on country-specific deviations from growth trends. Burke et al. find that country-level aggregate production is smooth, non-linear, and concave with respect to temperature for all countries, with a maximum at 13°C. This relationship appears to be consistent since 1960 and across agricultural and non-agricultural production in both rich and poor countries. Thus, these results suggest that economic activity depends on climate in all countries, regardless of level of development. Deryugina and Hsiang (2014) explicitly investigate the possibility that wealth protects countries from the adverse of impacts of temperature by estimating the impact of daily temperature on annual income in United States counties between 1969 and 2011. Like the other studies I summarize here, they use county and year fixed effects to isolate the impact of county-specific deviations in temperature on income. Deryugina and Hsiang (2014) find that the average productivity of individual days falls roughly linearly by about 1.7% for each 1°C in daily average temperature above 15°C, with this impact remaining consistent across their study period.

Cachon, Gallino, and Olivares (2012) supplement these sector-level estimates with a detailed study of the impacts of weather on plant-level output in the automobile industry in the United States. The authors use weekly production data from 64 automobile plants from 1994 through 2005 estimate the impact of extreme rain, snow, heat, and wind on the production of automobiles. They use a broad collection of controls, including plant fixed effects, a number of controls for seasonality, national time effects, and variables for the introduction and ramp-down of products. As in the other papers I have summarized here, these controls allow Cachon, Gallino, and Olivares to plausibly identify the causal impact of local weather fluctuations.

While they find that production does not fall significantly in weeks with one day with maximum temperature over 90°F or with between two and five days with maximum temperature over 90°F, production falls by 8.75% on weeks with six or seven days with maximum temperature over 90°F. Cachon et al. do not find evidence that these production losses are recouped in the week following the temperature shock, though their analysis leaves open the possibility that they are recouped over a larger time window. This result is particularly interesting in that it identifies negative impacts of outdoor temperature in an industry where work is indoors and likely to be air-conditioned.

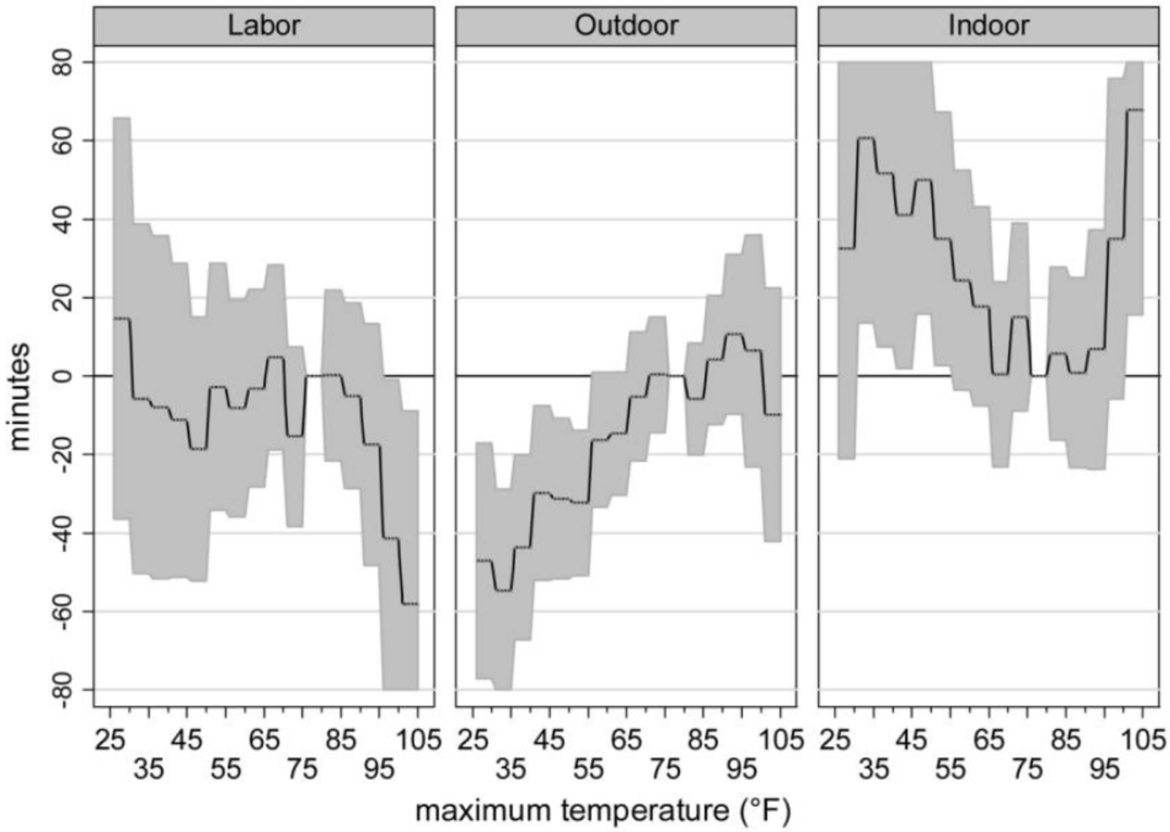

Figure 1: Relationship between temperature and time allocation for temperature-sensitive industries, where the panel labeled "Labor" gives the impact of temperature on time spent working, the panel labeled "Outdoor" gives the impact of temperature on time spent on outdoor leisure, and the panel labeled "Indoor" gives the impact of temperature on time spent on indoor leisure (Graff Zivin and Neidell, 2014).

These impacts might arise because indoor air-conditioning is insufficient when it is extremely hot outside, because exposure to temperature outside of work reduces employees' at-work productivity, or because of other disruptions to production outside the plant interior. In any case, Cachon et al.'s results suggest that the impacts of temperature may extend beyond those industries that are obviously sensitive to weather conditions.

Besides these studies of the impact of temperature on industrial output, a recent study by Graff Zivin and Neidell (2014) looks at the impact of temperature on the allocation of time to labor. The authors use a panel of county-level data from 2003 through 2006 from the American Time Use Survey to estimate how temperature impacts the amount of time spent working in industries where workers are plausibly exposed to weather, like agriculture, forestry, mining, construction, and utilities. They also categorize manufacturing as temperature-sensitive, noting the questionable use of air-conditioning in manufacturing plants. Using temporal and geographic fixed effects, the authors identify the impacts of temperature on time allocation using local and plausibly exogenous variation in weather. Overall, they find that time spent working in these temperature-sensitive industries falls significantly at high temperature extremes, with labor supply falling by about an hour per day on days with maximum temperatures over 100°F relative to a day with maximum temperature between 76 and 80°F. The authors do not find that temperature significantly reduces time spent working in indoor industries other than manufacturing.

2.2 Background on the Relationship Between Temperature and Accidents

Despite this growing body of work on the impacts of climate change on industry, no studies have as yet estimated the potential risks that climate change may pose to worker safety. Any change in the incidence of occupational accidents induced by climate change represents a shock to worker health; to the extent that accidents divert resources from production processes, we can also think of these impacts as shocks to labor productivity. Then, the impacts of temperature on accident incidence could be one mechanism through which temperature reduces aggregate economic output. Workers across the United States are subjected to a wide range of environmental conditions, from construction workers exposed to outdoor weather to iron smelters exposed to process-generated heat. Working under these conditions may pose significant risks; if so, the shifts in weather patterns anticipated under climate change might have sizable implications for worker safety.

2.2.1 Physical Impacts of Temperature on Workers

These impacts may partially arise through the physiological impacts of temperature on the human body. A large body of academic and policy-oriented work has sought to examine the health risks associated with working in different weather conditions; to date, this work has focused on the impacts of high temperatures. Working under high temperatures may put workers at risk of heat stress, which is defined as exposure to excess heat from the combination of one's own metabolism and environmental sources. Since outdoor workers likely generate significant metabolic heat as they work, they are particularly susceptible to heat stress. The body's responses to heat stress include elevated heart rate, vasodilation, greater circulation to the skin, and sweating. If heat stress is extended, the loss of plasma and electrolytes from continuous sweating and changes in blood circulation can overwhelm the body's thermoregulatory systems, allowing core body temperature to rise and compromising the cardiovascular and central nervous systems. The symptoms of heat stress can range from discomfort to death, categorized as heat illnesses including heat rash, heat syncope, heat cramps, heat exhaustion, or heat stroke (Jackson and Rosenberg, 2010).

While there are no government-mandated standards for workplace thermal conditions, the National Institute for Occupational Safety and Health (NIOSH) periodically publishes recommended standards for temperature exposure. NIOSH produces standards for heat stress as a combination of environmental and metabolic heat, differentiating between tolerable heat stress limits for different work durations and for unacclimatized versus acclimatized workers. Figure 2 presents NIOSH's (2013) recommended heat stress limits for the "standard" acclimatized worker, where this worker weighs 154 lbs and has 1.8 m^2 of body surface. This figure gives both a "ceiling" limit, denoted by "C," which gives a heat load that should never be exceeded without adequate heat-protective clothing and equipment, and a series of "Recommended Alert Limits," denoted by "RAL," below which most healthy workers would be protected from adverse health impacts. For example, based on Figure 2, NIOSH recommends that workers producing 233 watts of metabolic heat be exposed to 90°F for only 15 minutes per hour. For context, Campbell (2012) estimates that working at a desk produces 95 watts of metabolic heat, level walking at 5.5km/hr or moderately hard work produce 250 watts of metabolic heat, and level walking at 5.5km/hr with a 20-kg pack or sustained hard work produce 350 watts of metabolic heat.

Recent research has shown that heat exposure has historically posed a significant threat to worker safety. The Center for Disease Control (2008) states that 423 workers died from the direct cause of exposure to environmental heat between 1992 and 2006, where 24% of these laborers worked in agriculture, forestry, fishing, or hunting. Then, in 2010, there were 4,190 heat-related injuries and illnesses resulting in one or more days of lost work (NIOSH). In the same year, 40 workers died from heat exposure, with 18 working in construction, 6 working in agriculture and mining, 6 in professional and business services, likely in waste management remediation services, and 3 in manufacturing. The Occupational Safety and Hazard Administration (OSHA) writes that

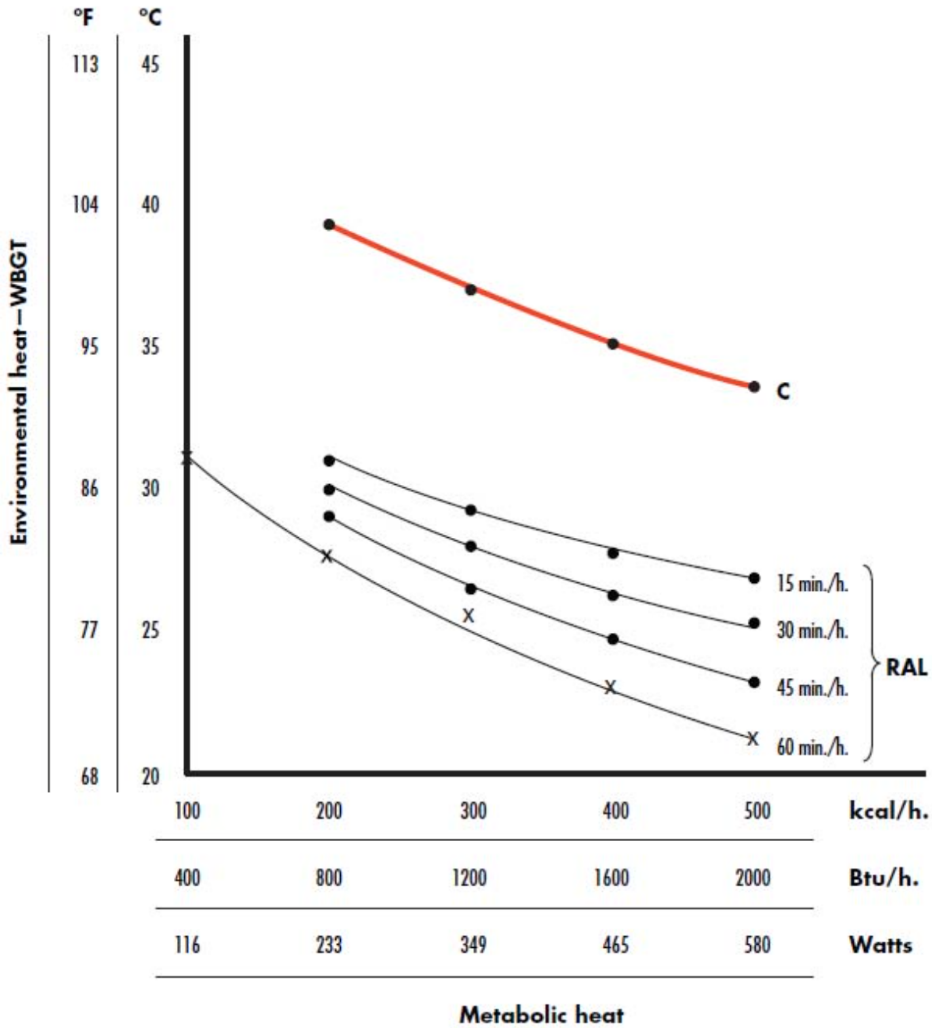

Figure 2: NIOSH's (2013) recommended heat stress alert limits for a "standard" unacclimatized worker. This standard worker weighs 154 lbs and has 1.8 m^2 of body surface. "C" denotes a Ceiling limit and "RAL" denotes Recommended Alert Limit.

the industries most affected by heat-related illness include construction, transportation, utilities, agriculture, building and grounds maintenance, landscaping services, and support activities for oil and gas operations (OSHA, 2015), and studies have evaluated the prevalence of heat stress in these and other industries (Ayyappan et al., 2009, CDC, 2008, Wasterlund, 1997, Donoghue, 2004). For example, a 2004 paper on heat illness in the U.S. mining industry states that 583 cases of heat illness were reported to the Mine Safety and Health Administration between 1983 and 2001, with the highest incidence in stone mills, metal mills, and underground metal mines.

While most research on work in extreme temperature has focused on the risks of working in extremely hot conditions, workers may also be at risk in very cold environments. In particular, workers exposed to extreme cold may be at risk of cold stress. In cold environments, the body shifts blood flow from its extremities to its core in order to maintain a sufficiently high core temperature, exposing the skin and extremities to rapid cooling and increasing the risk of frostbite. If cold stress is prolonged, the body may begin to lose heat faster than it can be produced, causing core body temperature to fall and ultimately leading to hypothermia (OSHA). As hypothermia becomes more severe, its symptoms can progress from goose bumps and shivering to cardiac and respiratory failure, and ultimately to death (Stocks et al., 2004).

To date, most studies have examined the influence of temperature on explicitly *heat-related* illnesses and injuries among workers, like heat illness and heat stress; however, there is reason to believe that extreme temperatures might impact the rate of occupational accidents more broadly. Indeed, while heat and cold stress can be the primary cause of workplace mortality or morbidity, their symptoms also likely make workers more susceptible to accidents. Even with heat illness of low severity, workers are likely to grow dizzy or faint. As heat illness progresses, victims may exhibit fatigue, nausea, lack of coordination, and confusion, all of which would impair their ability to work safely (Jackson and Rosenberg, 2010). Similarly, workers suffering from cold stress may become sluggish, confused, and unable to use their hands properly, again increasing accident risk (Occupational Health Clinics for Ontario Workers, Inc., 2005).

Furthermore, research in physiology suggests that, even before serious symptoms of heat or cold stress appear, both high and low extreme temperatures might make workers more prone to accidents by impairing cognitive function. A series of studies in the mid- to late-1900s evaluated the associations between temperature exposure and measures of cognitive function that link directly to accident risk, like coordination, vigilance, reaction time, and mental performance. Many of these studies find negative associations, particularly at high temperature extremes (Mackworth, 1947, Fraser, 1957, Pepler, 1960, Bell et al. 1964, Azer et al. 1972, Fine and Kobrick, 1978).

For example, Epstein et al. (1980) report that when study participants were asked to shoot at a square target on a video screen at 70°F, 86°F, and 95°F, their proportion of error increased from 7.9% to 15.9%, and then to 16.6% between those temperatures. Similarly, Wyon et al. (1996)

find that the proportion of signals missed by drivers increases by 50% and reaction time increases by 22% between 70 and 80°F. While other studies have produced conflicting estimates (Chiles, 1958, Pepler, 1959, Colquhoun, 1969, Grether et al., 1971, Ramsey, 1975, Ramsey and Pai, 1975, Lewis et al., 1983), several meta-analyses have found consistent negative associations between heat exposure and task performance after accounting for studies' measures of heat exposure and task type (Hancock and Vasmatzidis, 2003). In a review of 150 studies, Ramsey and Kwon (1992) find that performance of tasks requiring more demanding perceptual motor skills, like vigilance and coordination, is impaired at high temperatures, and Pilcher et al. find in a 2002 meta-analysis of task performance studies that exposure to hot and cold temperatures impairs performance of attention- and perception-based, mathematical, and reaction time-based tasks, with the worst performance above 90°F and below 50°F. Similarly, Seppanen, Fisk, and Faulkner (2003) find that productivity in various cognitive tasks falls by about 2% for each increase of 1°C at temperatures above 25°C.

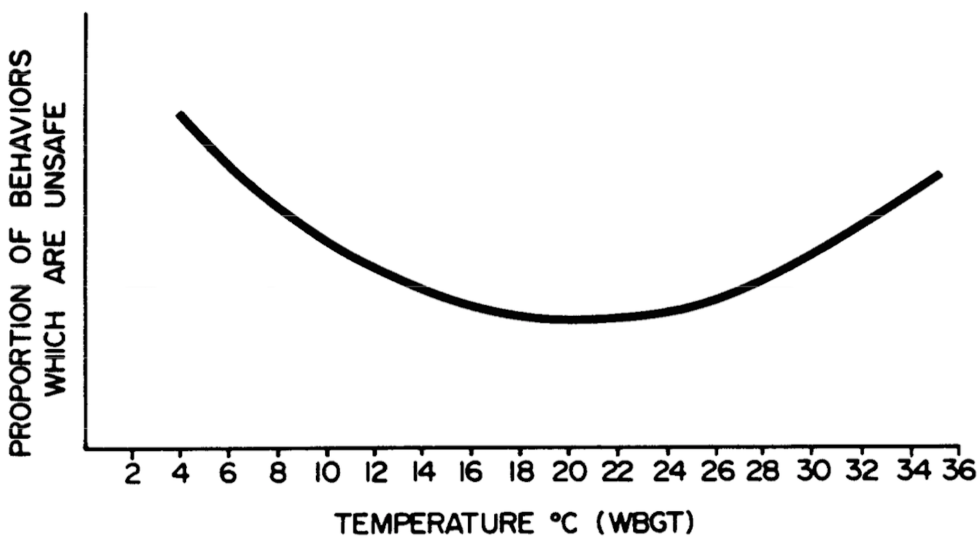

Figure 3: Ramsey et al.'s (1983) observed relationship between ambient temperature and the Unsafe Behavior Index (UBI) in two industrial plants, where the (UBI) is a function of the proportion of behaviors that are unsafe.

In turn, these physiological effects may translate into increasingly risky worker behavior. In a 1983 study of the effect of workplace heat exposure on at-work safety behaviors, Ramsey et al. observe worker behavior at a metal products manufacturing plant and a foundry, and develop an outcome variable for risky behavior called the Unsafe Behavior Index (UBI). The authors match observations of the UBI with measure of wet bulb temperature, which they split into five bins ranging from temperature below 60°F (15°C) to temperature above 90°F (30°C). Using a second order quadratic model and controlling for workers' metabolic workload and job risk group and the observation's time of day and day of the week, Ramsey et al. find a statistically significant U-shaped relationship between thermal exposure and unsafe work behavior. In particular, they find that unsafe behavior is lowest between about 63°F and 73°F, but increases roughly proportionally both at high and low temperatures. This U-shaped curve is pictured in Figure 3.

2.2.2 Adaptation in the Temperature-Accident Relationship

Based on these mechanisms, we might hypothesize that there exists a U-shaped curve between temperature and accident risk, with the risk of workplace accidents rising both at extremely high and extremely low temperatures. Even if the risk of accidents for a given amount of work follows this U-shape, however, the observed relationship between temperature and the incidence of accidents could be altered by adaptation. In this context, we can think of adaptation as processes or decisions that help reduce the negative impacts of short- or long-run shifts in temperature. This adaptation could be physiological or economic: in this context, physiological adaptation describes the process by which human bodies physically acclimatize to adverse weather conditions, and economic adaptation describes actions and behavior changes that economic actors take to reduce the negative impacts of temperature shocks on accident incidence. I will describe each of these forms of adaptation in more detail.

Research in physiology suggests that workers can adjust fairly quickly to adverse temperature environments, creating a mechanism for physiological adaptation that could significantly reduce the impacts of temperatures on the risk of workplace accidents. Most research has focused on acclimatization at high temperatures, where it advances primarily through increased sweating efficiency, with earlier onset, increased volume, and lower electrolyte concentration of sweat (NIOSH, 2013). In turn, enhanced sweating efficiency allows workers to function with a lower core temperature, reduced heart rate, and overall reduced thermoregulatory strain in hot conditions. Workers can typically achieve acclimatization with daily heat exposure over one to two weeks, though the speed of this process depends on worker age and health. This process of acclimatization, which has been well-documented both in laboratory experiments and field studies (Lind and Bass, 1963, WHO, 1969), would likely reduce the risk of both heat illness and accidents among workers that are consistently exposed to high ambient temperatures. While physiological adaptation to cold temperatures is less rapid and has been the subject of less extensive study, research suggests that workers would likely show a similar capacity to adapt to consistently cold temperatures (Kaciuba-Uscilko et al, 1989). Then, these physiological adaptive mechanisms might reduce the impact of temperature on accident incidence.

While physiological adaptation is a biological process that does not involve decisions by economic actors, economic adaptation to the impacts of temperature is any action taken by workers or employers to reduce the negative impact of potential shifts in temperature on accident incidence. These adaptive behaviors can occur at a wide range of timescales. In the short-term, we might expect that workers and employers would reduce the duration or intensity of work in response to uncomfortably hot or cold working conditions. For example, a draft publication by the National Institute for Occupational Safety and Health on heat stress among workers makes a series of recommendations for short-term protective behaviors against hot weather, including advising that workers take frequent breaks for rest, food, and water and that employers provide areas for rest and recovery, provide cooling protective clothing, reschedule particularly strenuous jobs to cooler times whenever possible, and schedule safety and health training (NIOSH 2013).

Similarly, a 2010 paper on prevention of heat illness among agricultural workers outlines a range of protective behaviors for both employees and employers, including modifying work assignments in response to weather, providing shaded rest stations, creating heat acclimation plans, and taking frequent breaks of adequate duration (Jackson and Rosenberg, 2010). Recently, studies have found some evidence that workers adjust time spent working in response to temperature. For example, Graff Zivin and Neidell's 2014 paper on the impact of temperature on time allocation found that workers in certain outdoor industries dedicate about an hour less to labor on days with maximum temperature over 100°F relative to days with maximum temperature between 76°F and 80°F.

At longer timescales, other forms of economic adaptation to extreme temperatures would become possible. If working outdoors becomes more dangerous as climate change progresses, for example, we might expect to see the gradual development of new worker safety technologies or increased access to air conditioning for workers that are exposed to outdoor conditions. We might also see the development of new worker safety regulations or shifting norms related to working in adverse conditions. These adaptive responses would help to avert some of the costs of extreme temperature associated with changing the risk of occupational accidents.

Then, even if accident *risk* is U-shaped curve with respect to temperature, these physiological and economic adaptive behaviors might alter the shape of the relationship between temperature and the *incidence* of accidents, both in the short- and long-term. We would expect to see that short-term behavioral adaptations, like reductions in workload at extreme temperatures, would reduce the impacts of temperature on accident risk; more precisely, they would lessen the increase in accident incidence at temperature extremes implied by the increase in accident risk per unit of work at those temperatures. Even in the short term, then, we might expect to see a dampened U-shaped relationship between temperature and the incidence of workplace accidents, or even a reversed U-shaped relationship (Xiang et al., 2014c).

Then, long-run adaptation, like the development of new policies and technologies, would bend down the impact of temperature on accident incidence in the long-run. That is, we would expect to see that the impact of temperature on occupational accidents would become increasingly damped as time progresses and adaptive mechanisms develop. Physiological acclimation might work in a similar way. Assuming that high temperatures increase accident risk and become more frequent under climate change, physiological acclimatization would reduce the impacts of temperature on accident incidence as temperatures grow consistently higher and workers become better acclimated to those conditions. While these long-term adaptation possibilities may not show up in my estimates of the short-term impacts of temperature, extrapolating these short-term estimates to the long-term changes in accident incidence that we would expect to see under climate change would require considering the extent to which adaptation will reduce these impacts.

2.3 Previous Research on Temperature and Accident Incidence

A series of studies have sought to estimate the association between temperature and the incidence of occupational accidents. Overall, these studies have identified a spread of nonlinear relationships between temperature and accidents, ranging from reversed U-shaped curves, where accidents become less frequent at extreme temperatures, to U-shaped curves, where accidents become more frequent at extreme temperatures. It seems that this divergence may be largely explained by studies' inclusion or exclusion of "denominator data," where the denominator is a measure of work volume or output that can be used to calculate a rate of accidents per unit of work. In other words, this denominator data partially captures the extent of short-term adaptation to temperature, as I described in Section 1.3.2. Xiang et al. (2014a) summarize the role of this denominator data, writing that while adaptive behaviors may "result in the unexpected decline in the number of work-related injuries," the association between temperature and accidents tends to be U-shaped when "denominator information is available for calculating work-related injury rates".

The first quantitative study of the relationship between temperature exposure and occupational accidents was published in 1922, when Osborne et al. reported to Britain's Industrial Fatigue Research Board on the effect of temperature on accident frequency. The authors examine the frequency of accidents in three British munitions factories. While Osborne et al. do not collect work output denominator data, they seek to address broad differences in work volume or duration by controlling for shift type and verifying that employment and output remain relatively constant over the year. The authors compare accident incidence across shift types and temperature categories, and they find that the frequency of accidents is lowest between 65 and 69°F. The rate of accidents increases at temperatures below 65°F, and increases sequentially between 65 to 69°F and 70 to 74°F, and then again between 70 to 74°F and temperatures over 75°F. The only exception to this steady U-shaped curve is a dip in accident frequency below 49°F, which the authors attribute to adaptive behavior, suggesting that work may slow or cease altogether at those temperatures. I present the authors' depiction of this nonlinear relationship between temperature and accident frequency in Figure 4.

Powell et al. follow this work in 1971 with a study of the causes of 2,367 accidents in four

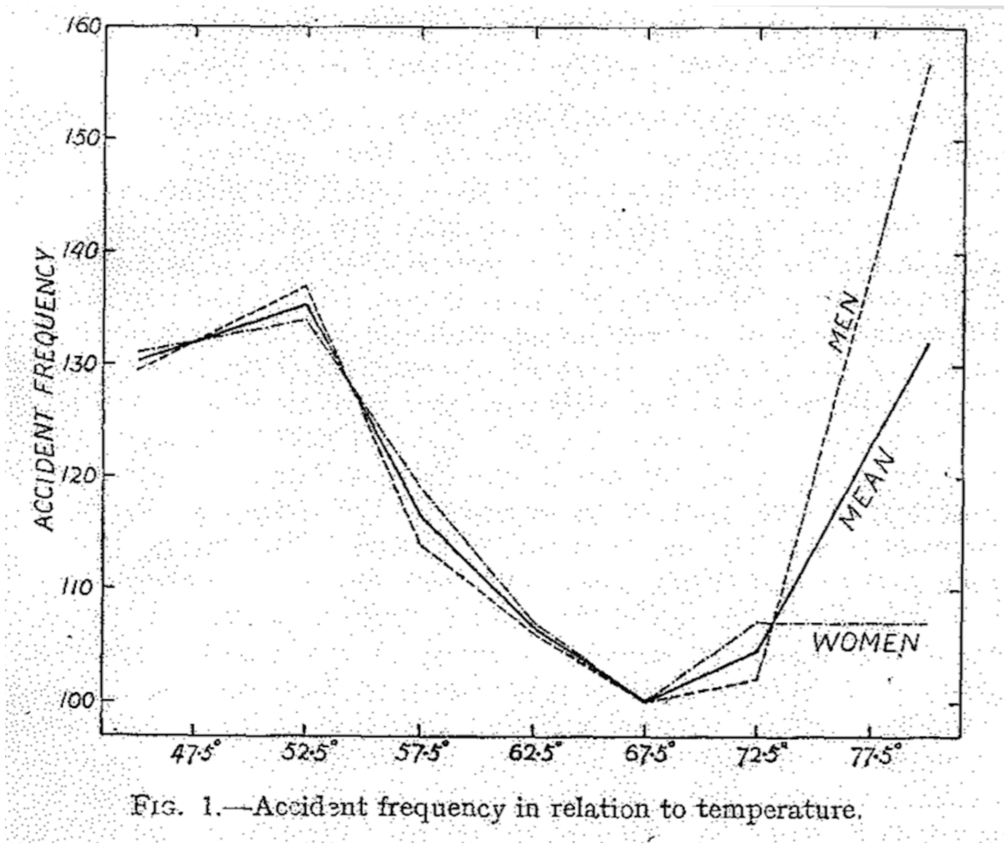

Figure 4: Osborne et al.'s (1922) observed relationship between accident frequency and temperature in two munitions factories.

different industrial workshops between 1966 and 1969. In each of the four facilities, the authors calculate the frequency of accidents that occur over a spread of temperature ranges and then use simple t-tests to determine whether the mean temperature at which accidents occur in each facility is significantly different from the mean temperature over the study period. They find significant relationships in two of the four shops, with an increase in the frequency of accidents at temperatures below about 68°F. They find no increase of accidents at high temperatures, which they attribute to "the noticeable slackening of work done in hot weather" (20, Powell et al. 1971). Like Osborne et al. (1922), Powell et al. do not collect data on work volume at these facilities, so their estimates of the impact of temperature on accident frequency are shaped by the impact of temperature on work duration and intensity.

In a study of safety records at a Midwestern aluminum smelting plant, Fogleman et al. (2005) adjust the methodology of these early studies by incorporating detailed data on person-hours worked

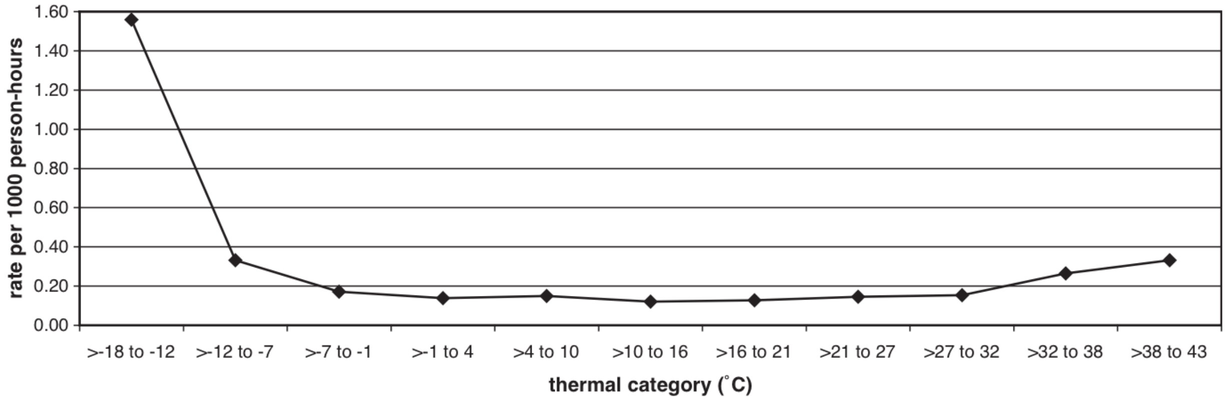

Figure 5: Fogleman et al.'s (2005) observed acute injury rates by outdoor thermal category in a Midwestern aluminum smelter factory.

at this facility; this denominator data allows them to evaluate the impact of temperature on the number of accidents for a given amount of work, more cleanly measuring changes in accident risk. The authors use an outdoor heat index based on temperature and humidity as a surrogate for heat conditions inside the open plant, splitting this index into 10°F bins ranging from between 0 and 10°F to between 100 and 110°F. They generate data on the ratio of acute injuries to person-time for a series of temperature categories and work locations, and then use Poisson regression to evaluate the effect of categories of this heat index on these ratios of incidents to person-time after controlling for age and work location. Fogleman et al. find significantly elevated odds ratios for accidents for an outdoor heat index both at temperatures below 20°F (-7°C) and at temperatures above 90°F (32°C), though this increase is much larger in magnitude at low temperatures. See Figure 5 for the curve between injury rate and outdoor thermal category that Fogleman et al. trace out.

Since Fogleman et al. use a measure of accident frequency adjusted for work volume, their estimates are unique in the literature on temperature and accidents in that they are able to estimate the impact of temperature on the risk of accidents for a given amount of work, rather than simply estimating the impact of temperature on the number of accidents that take place. So far, Fogleman et al's work is the only study to make use of detailed denominator data.

More recently, several studies have aimed to evaluate the effect of temperature on occupational safety at the level of cities and regions. While these studies may be more widely relevant than the facility-specific design of earlier work, this broader spatial scale also makes it increasingly difficult to identify meaningful denominator data. In the first such paper, Morabito et al. (2006) explore the effect of high daily apparent temperature on summer hospitalizations for work accidents at six hospitals in Tuscany, Italy between 1998 and 2003. The authors use non-parametric tests, including the Mann-Whitney U test and the Kruskal-Wallis H test, to look for differences in work accident rates between quartiles of daily maximum, 24-hour mean, and daytime mean apparent temperature (AT). Overall, the authors find a reversed U-shape relationship between AT and hospitalizations for workplace accidents, with the highest accident rates occurring between about 76°F and 82°F, representing high, but not extreme, apparent temperature. Like other authors, Morabito et al. attribute this reduction in accident incidence at extreme temperatures to behavioral changes that actually reduce accident incidence below that at moderately high temperatures, despite the plausible increase in the risk of accidents for a given volume of work at very hot temperatures.

While Morabito take a non-parametric approach, several recent studies in public health have sought to expand this work by more precisely quantifying the impact of temperature on the incidence of accidents. In general, these public health studies are based on city- or region-level daily accident counts, use Poisson or negative binomial regression, and allow for some level of geographic and temporal heterogeneity.

Xiang et al. (2014b, 2014c) investigate the association between high temperature and workplace accidents in Adelaide, Australia using workers' compensation data. In the first of these, the authors estimate the association between daily maximum temperature and injury counts on weekdays from 2001 to 2010 using a generalized estimating equation model with a negative binomial distribution accounting for overdispersion and a first-order autocorrelation structure. The authors allow for nonlinearity in the relationship between temperature and work-related injury by constructing a piecewise linear spline function with one knot at around 98°F, allowing the impact of heat on injuries to differ above and below this threshold. Xiang et al. restrict the impact of seasonality by limiting their analysis to Australia's warm season between October and March and including month fixed effects; they also include year and day-of-week fixed effects.

The authors stratify analysis by "indoor" and "outdoor" industries, where outdoor industries in-

Figure 6: Estimated exposure-response relationships between daily maximum temperature and daily injury worker compensation claims for (A) total effects, (B) outdoor industries, and (C) agriculture, forestry, and fishing in Adelaide, Australia (Xiang et al., 2014c)

clude agriculture, forestry, fishing, construction, electricity, gas, and water services. Overall, Xiang et al. find that an increase of 1.8°F (1°C) is associated with a 0.2% increase in the number of daily injury claims at temperatures below 100°F, but that the number of daily injury claims decreases by 1.4% with an increase in daily maximum temperature of 1.8°F (1°C) at temperatures above 100°F (37.7°C). Thus, the relationship between daily maximum temperature and worker compensation claims takes a reversed U-shape. Figure 6 shows the shape of this curve among all workers, among workers in outdoor industries, and among workers in agriculture, forestry, and fishing. While the y-axis scales differ between these figures, these industries showed very similar percentage changes in accident rate in response to temperature. As in other papers on accident and temperature, Xiang et al. attribute the reduction in injuries at extremely high temperatures to protective measures. In particular, they point out that they observe no such decrease in injuries within workers in "electricity, gas, and water" industries at extreme temperatures, which they suggest may be because workers in those industries must ensure the continuous supply of electricity, gas, and water through extreme conditions, and are therefore unable to make behavioral adjustments.

In a second paper, Xiang et al. (2014c) investigate the impact of heat waves on workers' compensation claims in Adelaide, Australia between 2001 and 2010. Here, they define a heat wave as a period of at least three consecutive days with daily maximum temperature over 95°F. As in their first 2014 study, the authors use generalized estimating models with a negative binomial distribution, and they attempt to eliminate the influence of seasonality by limiting analysis to weekdays between October and March and by including month fixed effects; they also include year and day-of-week fixed effects. Again, they classify industries as either "outdoor" or "indoor," where outdoor industries include agriculture, forestry, fishing, construction, electricity, gas, and water services. Xiang et al. find that while heatwave conditions do not significantly increase worker safety claims among all workers, claims within outdoor industries increase by 6.2% during heatwaves.

Adam-Poupart et al. (2015) follow a similar methodology in their analysis of the association between heat-related occupational injury compensations and exposure to summer maximum daily temperatures in each of the 16 heath regions of Quebec for the years 1998 through 2010. Like Xiang et al. (2014b, 2014c), the authors estimate the association between daily compensation counts and daily hourly maximum temperature using a generalized linear model with negative binomial regression. They also include a similar suite of controls, including day-of-week and month with fixed effects, a linear year term, and controls for public holidays and humidity. However, while Xiang et al. use a spline function in temperature to allow for a nonlinear impact of temperature on accidents, Adam-Poupart et al. model a linear relationship between daily maximum temperature and heat-related compensation counts. In analyzing each of the health regions of Quebec, this analysis is the first to examine the impact of temperature on accidents across different working populations; Adam-Poupart et al. include the log of monthly working population in each health region as an offset for the negative binomial regression to adjust for the sizes of these worker pools. After estimating these risk functions for each health region in Quebec, Adam-Poupart et al. pool these effects using a random-effects model. Overall, the authors find that an increase of 1°C in daily maximum temperature is associated with a 42% increase in the count of daily heat-related occupational injury compensations in Quebec.

To date, all substantive work on the impacts of temperature on worker safety has focused on North America, Europe, and Australia. However, climate change may pose particular risks to workers in low- and middle-income countries (Kjellstrom et al, 2009). In particular, the impacts of temperature on accident incidence could be particularly severe due to limited access to air-conditioning, the prevalence of output-based payment systems that discourage protective behavior in favor of high productivity, and industrial structures centered around outdoor industries like agriculture, forestry, or mining. To date, Tawatsupa et al. (2013) have found that heat stress is significantly associated with occupational injuries in Thailand, though the study is limited in that it relies on qualitative measures of heat exposure and occupational injuries reported by participants in the 2005 Thai Cohort Study.

3 Data Sources

3.1 Weather Data

I use daily-level weather data generated from raw data drawn from the National Climatic Data Center (NCDC) Summary of the Day Data files (TD-3200). These files contain daily weather observations from stations in the National Weather Service (NWS) cooperative station network across the United States. While most of these stations are operated by private volunteers, they include NWS principal climatological stations and stations operated by employees of other federal agencies, like the Federal Aviation Administration and the National Park Service. About 8,000 stations in this cooperative network are currently active, but about 23,000 stations have been active since the network's formation in the mid 1900s. My primary data elements of interest are daily maximum temperature, minimum temperature, and precipitation.

To ensure the accuracy of these weather readings, I use weather data from stations selected using a rule modeled after that used by Deschênes and Greenstone (2011), who use the same weather data. First, I ensure the quality of weather data by only including data for a given year from stations that provide measurements for maximum temperature and precipitation for each day of that year. Over 10,000 weather stations meet this selection criterion for at least one year between 1984 and 2010, and an average of 4,407 stations meet this selection rule in a given year in that range. Next, while Deschênes and Greenstone only use weather data from stations at elevations below 7000 feet to reflect the weather conditions experienced by most residents of a county, I seek to better accomplish this goal by only using stations located in census tracts that had population density of over 100 people per square km in 2004. Indeed, a rule based on station elevation only loosely identifies the areas where people live, and such a rule potentially places far more restrictions on stations in mountainous areas of the country, like Wyoming and Montana, than on stations in areas like the Midwest.

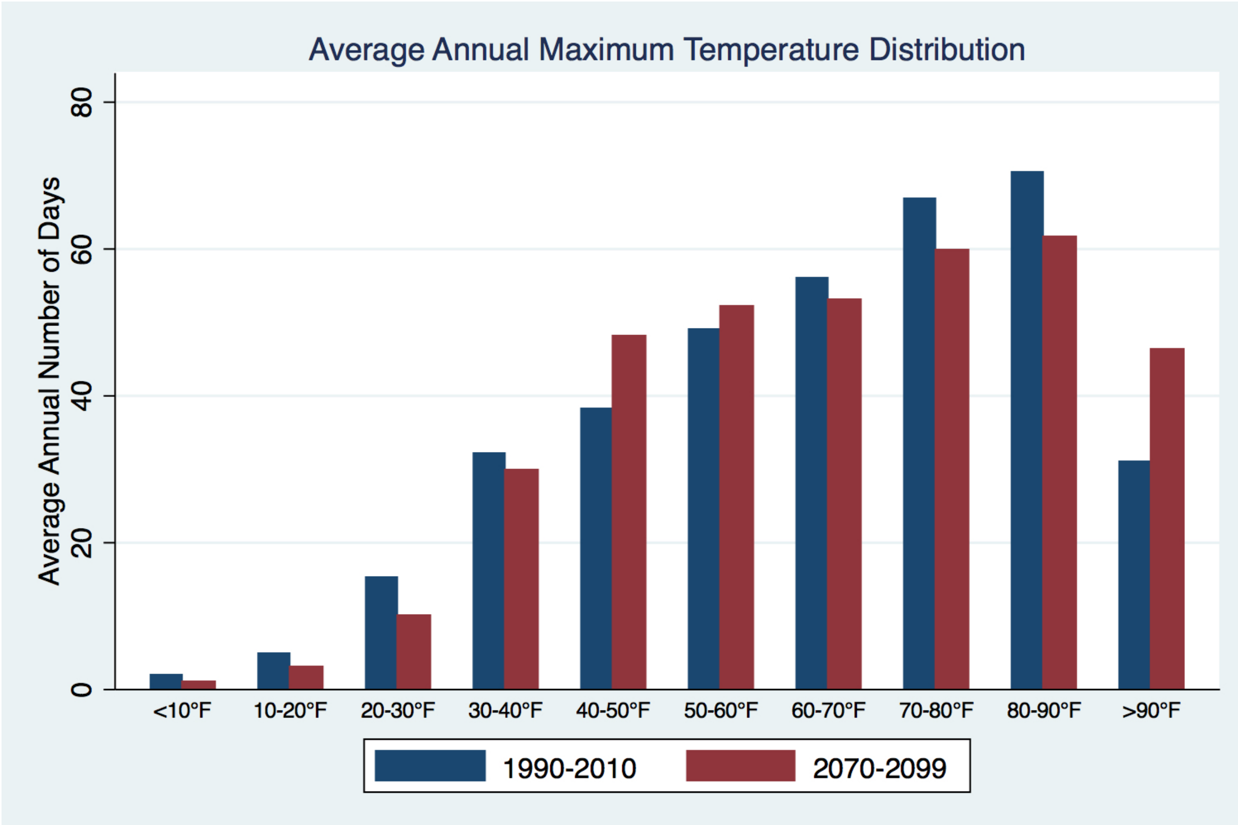

Figure 7: Historical and forecast distribution of daily maximum temperature across ten temperature bins. Bars for 1989-2010 represent the average number of days in each temperature bin per year for a county in my sample. Bars for 2070-2099 represent the average number of days in each temperature bin per year in a county under the Hadley 3 climate model under the highest warming scenario (A1). Forecast data is from replication files for Deschênes and Greenstone. (2011).

A total of 3,322 stations meet our final selection rule for at least one year between 1984 and 2010, and an average of 1,632 stations meet this rule for a given year in that range. I then produce county-level weather data by taking an inverse-distance weighted average of the data from all remaining stations within 200 km of the centroid of each county. In some cases, my data for a county's weather conditions is drawn entirely from data from stations outside of that county. Then, inverse-distancing weighting allows us to approximate weather in counties without stations while also privileging the readings from within-county stations for those counties with functioning weather stations within their borders. I construct a measure for maximum temperature in a county using data from weather stations in a different county. Inverse-distance weighting

Next, I merge this weather data with additional daily data on temperature from the North

Table 1: Temperature Summary Statistics

Census Division	Avg Daily Max Temp (°F)	Avg # Days with Max Temp < 10°F	Avg # Days with Max Temp < 20°F	Avg # Days with Max Temp > 90°F	Avg # Days with Max Temp > 100°F
All	64.53	2.005	7.045	30.477	3.637
New England	54.172	2.156	9.296	3.389	0.305
Middle Atlantic	56.476	1.04	7.705	2.506	0.006
E.N. Central	58.196	2.292	10.936	8.714	0.328
W.N. Central	60.042	6.818	19.336	30.866	4.067
South Atlantic	69.425	0.045	0.526	25.759	0.553
E.S. Central	69.023	0.114	0.663	25.023	0.988
W.S. Central	75.953	0.034	0.252	83.137	14.623
Mountain	57.875	1.867	7.745	20.053	1.969
Pacific	61.774	0.578	3.132	22.27	4.721

Note: All figures *per annum* and by county. Averages calculated first by county and then averaged across counties in my sample. Census divisions are defined as follows: New England (CT, ME, MA, NH, RI, VT), Middle Atlantic (NJ, NY, PA), East North Central (IN, IL, MI, OH, WI), West North Central (IA, KS, MN, MO, NE, ND, SD), South Atlantic (DE, DC, FL, GA, MD, NC, SC, VA, WV), East South Central (AL, KY, MS, TN), West South Central (AR, LA, OK, TX), Mountain (AZ, CO, ID, NM, MT, UT, NV, WY), and Pacific (CA, OR, WA).

America Land Data Assimilation System (NLDAS). NLDAS is a collaborative project between NOAA, NASA, Princeton University, and the University of Washington which has produced a record of county-level daily maximum temperature and heat index across North America from 1979 to 2011. Heat index is a measure of apparent temperature that incorporates humidity, and which therefore may be a better measure of temperature exposure than are dry temperature measurements. NLDAS calculates heat index for days with temperatures at or above 80°F according to the following formula, where HI is heat index, T is ambient dry bulb temperature, and RH is relative humidity:

$$HI = -42.379 + 2.049 * T + 10.14333127 * RH - 0.22475541 * T * RH - 0.00683783 * T^2 - 0.05481717 * RH^2 + 0.00000199 * T^2 * RH^2$$

I then merge this NLDAS maximum temperature and heat index data with data on maximum temperature and precipitation from the NCDC Summary of the Day datafiles. These data sources are in close agreement when both are available; NLDAS and NCDC data on maximum temperature have a correlation coefficient of 0.9561. In some cases, only NLDAS temperature data is available for a particular county-year combination. Then, since NLDAS data does not include a measure of precipitation, I assign daily precipitation as the simple average of available precipitation data from neighboring counties.

I restrict my sample to the contiguous United States, excluding Alaska, Hawaii, and any other U.S. territories. The contiguous United States contains 3,108 county equivalents; based on my two sources for weather data, all 3,108 counties have daily weather data for at least one year between 1989 and 2010, and 2,840 counties have daily weather data for every year between 1989 and 2010. About 92% of accidents reported to OSHA during my study period can be matched to this weather data.

3.2 Accident Data

My data for occupational accidents comes from the Occupational Safety and Health Administration's Enforcement Inspection and Accident Investigation Data. Since the Occupational Safety and Health Administration (OSHA) was formed with the passage of the Occupational Safety and Health Act of 1970, employers have been required to report certain workplace incidents to OSHA. The first accident reporting requirements were instituted by a bill passed in July 1971. This bill, 36 FR 12612, required that employers report all workplace fatalities and accidents resulting in the hospitalization of five or more employees to OSHA within 48 hours of the incident. In 1994, these regulations were expanded to require that employers report accidents resulting in the hospitalization of only three workers and that employers make these reports within eight hours of learning of the incident. These regulations were strengthened again in 2015 to require that employers report to OSHA all work-related fatalities within 8 hours of the event, and report all work-related inpatient hospitalizations of one or more employees, all work-related amputations, and all work-related losses of an eye within 24 hours of the event.

A total of 93,566 accidents were reported to OSHA from the contiguous United States between 1970 and 2013. There is significant variation in the number of accidents reported each year over that period, but it is not clearly tied to these changes in reporting regulations. Accidents were first reported in large volume in the mid-1980s.[1] OSHA's records for occupational accidents also

[1] I have spoken to current OSHA employees as to this discrepancy, and none as yet have been able to explain it.

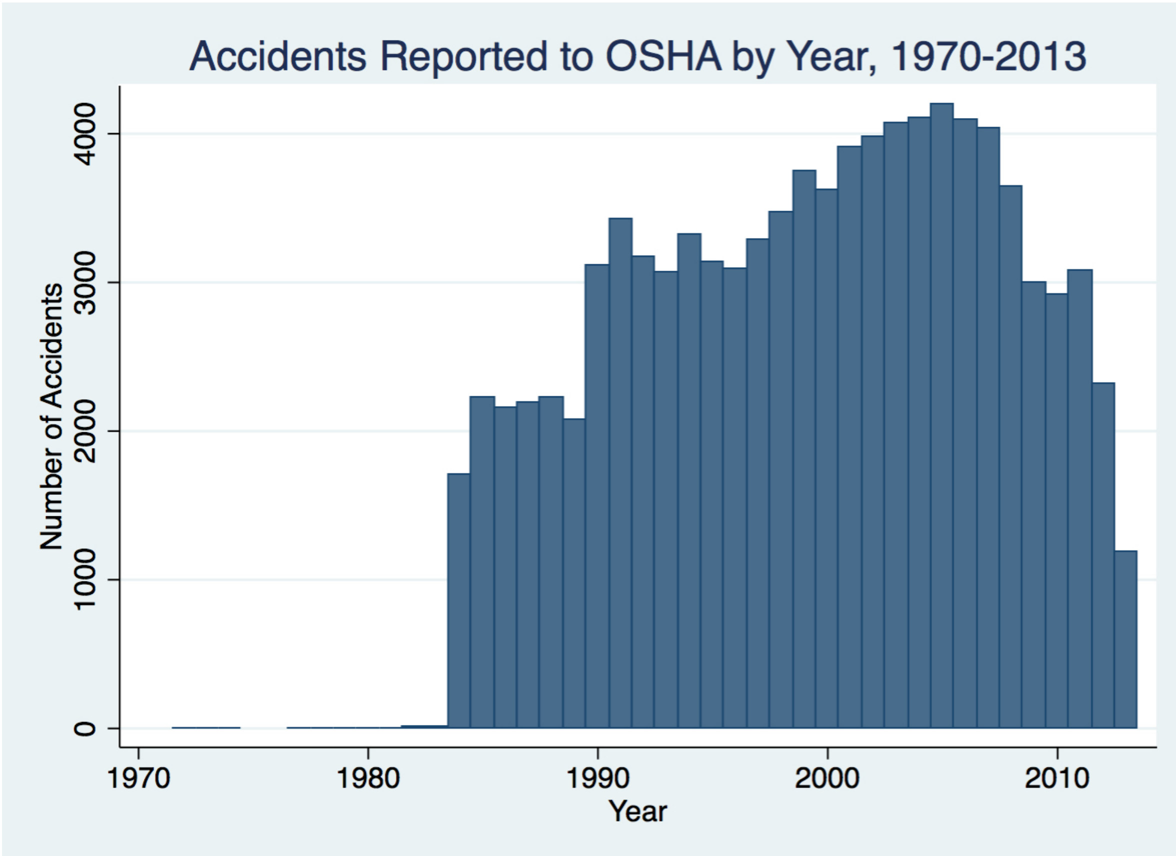

Figure 8: Total number of fatal and nonfatal accidents reported to the Occupational Safety and Health Administration (OSHA) by year. Figures are based on OSHA's Enforcement Inspection and Accident Investigation Data.

drop off in recent years due to a backlog in accident investigation and processing, since accident investigation cases must pass through a review process before being reported to the public. While OSHA states that this backlog of cases currently extends back to August 2013, the marked fall in the number of accidents following 2010 suggests that some cases from 2011 and 2012 may also currently be unreported (DOL Data Enforcement). We also see a sudden increase in the volume of accidents reported to OSHA in 1990, which does not correspond to a change in accident reporting requirements; this shift may be tied to a transfer of recordkeeping responsibilities from BLS to OSHA during that year.[2]

[2]OSHA's Office of Health Enforcement could not provide an explanation. I hypothesize that this increase in accident reporting in 1990 may be tied to a transfer of recordkeeping requirements from BLS to OSHA in that year. A Memorandum of Understanding dated July 11, 1990 delegated responsibility for administration of accident recordkeeping to OSHA, while leaving BLS responsible for conducting the Annual Survey Of Occupational Injuries And Illnesses.

In order to account for these somewhat ambiguous changes in accident reporting over time, I restrict my analysis to the years 1990 through 2010, a range in which there are no abrupt shifts in the volume of accidents reported each year. About 79.5% of accidents reported to OSHA occurred in this period. I further correct for changes in the volume of accidents reported to OSHA with year fixed effects.

It is important to note that OSHA's accident records are not a complete record of workplace injuries and deaths in the United States. Not only do these records include only particularly severe accidents, but they include only those accidents that occur in certain industries and among certain groups of workers. While OSHA has ultimately responsibility for the safety and health of all employees in the United States, its jurisdiction is limited in several respects. First, while OSHA covers almost all private sector employees in all 50 states, Washington D.C., and other U.S. jurisdictions, its coverage of public sector employees varies across the United States. Over 95% of the accidents reported to OSHA between 1970 and 2013 occurred in the private sector. Again, twenty-two states and territories have an OSHA-approved state program, which are required to be at least as stringent as the Federal OSHA program; these plans cover state and local government employees in addition to private sector employees. In addition to these twenty-two states, Connecticut, Illinois, Maine, New Jersey, New York, and the Virgin Islands have OSHA-approved plans that specifically cover public sector employees. In other states, public sector employees are not protected under OSHA, and accidents concerning them may not be reported to OSHA. In order to bypass uncertainty in coverage of public-sector employees by state, we will limit our analysis to accidents involving private-sector employees. Next, besides excluding public-sector employees in a number of states, OSHA does not cover self-employed workers, the immediate relatives of farm employers, or domestic workers.

In addition to this limited coverage of public sector employees, the self-employed, and family farm workers, OSHA's purview over certain industries is limited. While OSHA has ultimate responsibility for the safety and health of all employees, it can be preempted in a specific task or industry by another federal agency. Based on communication with OSHA and my best analysis of available resources on OSHA's jurisdiction, OSHA's accident records do not include accidents related to the defense industry, accidents related to mining, accidents related to aviation besides those involving ground crews or other airport-based labor, trucking accidents except those in off-highway loading and unloading, accidents involving Department of Energy contractors, fishing accidents except those at shipyards or near the coast, motor vehicle accidents on public streets or highways unless in a construction work zone, and accidents that occurred on a commercial or public transportation system.

Furthermore, employers likely do not report all accidents to OSHA. The Committee on Education and Labor of the House of Representatives (2008) notes that employers have strong incentives to underreport workplace accidents and illnesses, since businesses with fewer accident reports are more likely to receive government contracts, are less likely to receive OSHA inspections, and have lower workers' compensation insurance premiums. Thus, workers report that their employers actively discourage or intimidate them from reporting injuries and illnesses (Committee on Education and Labor, 2008).

Several studies have attempted to estimate the rate of underreporting to OSHA, focusing primarily on underreporting to the Survey of Occupational Injuries and Illnesses (SOII). The SOII is an annual program administered by OSHA and the Bureau of Labor Statistics that uses detailed accident and illness logs from a sample of establishments around the county to estimate nation-wide accident incidence, so it is distinct from the incident-specific reporting of severe accidents to OSHA that is required of all businesses. However, rates of underreporting of accidents to the SOII likely reflect the rate of compliance with these incident-specific reporting requirements. Several studies of the SOII suggest that the rate of underreporting may be quite high; Rosenman et al. (2006) analyze Michigan worker compensation and BLS data for 1999 through 2001, finding that between 60% and 69% of injuries and illnesses went unreported to the BLS during that period. Another paper finds that 90% of firms in Washington State either willfully or accidentally failed to fully comply with OSHA's record-keeping requirements in 2008 (Wuellner and Bonauto, 2014). While some failed to meet the required reporting timeframe or did not report all eligible incidents, 12% of the establishments surveyed failed to even keep an OSHA accident log. Finally, Dong et al. (2011) estimate that small construction establishments reported only 25% of severe injuries among Hispanic workers and 60% of severe injuries among white workers nationally between 1992 and 2006.

Thus, accidents reported to OSHA make up only a portion of all accidents that occur each year, both because they exclude those occurring among large groups of public sector employees and in a number of risky industries and due to likely noncompliance with reporting requirements. Ican estimate how the volume of accidents reported to OSHA compares to the total volume of workplace accidents that occur each year using data on workplace fatalities from the Census of Fatal Occupational Injuries (CFOI) at the Bureau of Labor Statistics. CFOI combines federal and state-level data to create reliable counts of occupational fatalities. These national fatality counts include all industries and both public and private-sector workers, including workers covered under OSHA, those covered under other federal or state agencies, and those workers outside the scope of regulatory coverage. Besides collecting data on standard paid employees, CFOI monitors fatalities among volunteers, self-employed or unpaid family workers, laborers on small farms, and undocumented workers performing the same duty as paid employees. Furthermore, the CFOI is less susceptible to concerns of underreporting. While the CFOI is based partially on employers' reports of fatalities to OSHA, CFOI also verifies the accuracy of workplace fatality statistics using death certificates, workers' compensation records, news reports, and reports to other Federal and State agencies (Committee on Education and Labor, 2008).

Thus, CFOI fatality counts should theoretically provide a complete measure of workplace fatalities in the United States, and we can get a sense for the completeness of OSHA's accident records by comparing its total fatality counts to those reported under CFOI. Across the years 1992 to 2010, OSHA's records include an average of 29.9% of the workplace fatalities recorded by CFOI. This reporting rate ranges from a low of 25.2% in 1993 to a high of 35.1% in 2005. Thus, we see that due to some combination of underreporting by employers and limitations in OSHA's coverage, OSHA's fatality records include slightly less than one third of the total occupational fatalities occurring in the US. Since fatalities are easily detectable, it is likely that they are reported at higher rates than are non-fatal injuries (Morantz, 2014). Thus, we might expect the ratio of non-fatal injuries reported to OSHA to the total volume of non-fatal workplace injuries to be even lower.

Despite these limitations, OSHA's accident records remain the preferable source of data on accident incidence because of their geographic and temporal specificity. Other sources of accident data, like CFOI, provide aggregate accident counts, but do not disclose data on specific accidents.

In particular, CFOI's must granular accident data give annual accident counts by Metropolitan Statistical Area. In contrast, OSHA's accident records provide specific data on each accident, including the street address of the accident site, the Standard Industrial Classification (SIC) code of the associated industry, and the date of the accident. This geographic and temporal specificity allows me to match particular accidents with a close measure of the weather conditions in which they occurred. While CFOI's data on occupational fatalities is more complete and likely to be somewhat more accurate than OSHA's accident records, using these annual accident totals would likely obscure information on the contemporaneous relationship between temperature and accident incidence that could be revealed in analysis at the daily level. Thus, I proceed with OSHA's records of occupational accidents.

Table 2: Accident Summary Statistics

Census Division	Average # Accidents per County	Average # Fatalities per County	Average # TS Accidents per County	Average # Accidents per Million	Average # TS Accidents per Million
All	1.071	0.531	0.890	12.830	10.944
New England	1.945	1.134	1.546	13.971	11.105
Middle Atlantic	1.423	1.200	1.196	5.934	5.066
E.N. Central	0.706	0.520	0.604	7.493	6.438
W.N. Central	0.244	0.182	0.195	11.412	8.753
South Atlantic	0.809	0.585	0.690	9.784	8.434
E. S. Central	0.361	0.329	0.314	8.271	7.431
W. S. Central	0.560	0.539	0.480	17.864	16.240
Mountain	0.674	0.421	0.560	15.210	13.162
Pacific	12.547	1.666	10.255	53.620	45.434

Note: All figures *per annum* and by county. Averages calculated first by county and then averaged across counties in my sample. Census divisions are defined as follows: New England (CT, ME, MA, NH, RI, VT), Middle Atlantic (NJ, NY, PA), East North Central (IN, IL, MI, OH, WI), West North Central (IA, KS, MN, MO, NE, ND, SD), South Atlantic (DE, DC, FL, GA, MD, NC, SC, VA, WV), East South Central (AL, KY, MS, TN), West South Central (AR, LA, OK, TX), Mountain (AZ, CO, ID, NM, MT, UT, NV, WY), and Pacific (CA, OR, WA).

Again, these records include a wide range of information on each accident and the resulting workplace inspection, including accident location, an SIC code for the associated industry, and the accident date. Using the SIC code associated with each accident, I classify each accident by whether it is associated with a "temperature-sensitive" industry and whether it is a fatality or a nonfatal hospitalization. Broadly speaking, I define temperature-sensitive industries as agriculture, forestry, and fishing; construction; manufacturing; transportation, communications, electric, gas, and sanitary services; oil and gas extraction, and other miscellaneous outdoor services, like amusement park operation and police and fire protection. In general, I model these classifications on Graff Zivin and Neidell's (2014) designations of industries at "high risk" of temperature exposure. My list of temperature-sensitive industries is a combination of broad SIC divisions, like Division A, Agriculture, and a set of 2- and 4-digit SIC coded industries. See the Appendix for a complete list of SIC codes associated with temperature-sensitive industries. Thus, my definition of "temperature-sensitive" accidents is not limited to accidents that are directly attributed to heat stress, but rather includes all accidents in temperature-sensitive industries. This classification reflects the body of research in physiology that suggests that extreme temperature may impair measures of cognitive function like coordination, reaction time, and vigilance. Then, extreme temperatures could affect the risk of all workplace accidents in outdoor industries, not just the incidence of heat stress or heat illness. See Section 1.3.1 for more details.

Ultimately, I compile records for 93,566 accidents in the contiguous United States reported to OSHA between 1970 and present. To bypass uncertainty concerning OSHA coverage of public-sector employees by state, we restrict analysis to occupational accidents in the private sector, and we limit our analysis to accidents occurring between 1990 and 2010 to exclude years with seemingly incomplete accident coverage. In total, then, our analysis is based on records for 71,218 occupational accidents between 1990 and 2010.

In Figures 9, 10, and 7.2.1, we explore the distribution of these accidents across industries, seasons, and regions. Figure 9 decomposes this collection of accidents by SIC division. We see that most accidents reported to OSHA occurred in construction (SIC Division C) and manufacturing (SIC Division D), though agriculture, forestry, and fishing (SIC Division A), transportation, communications, and utilities (SIC Division E), and services (SIC Division I) also report significant volumes of accidents. About 86.2% of these accidents occurred in temperature-sensitive industries. It is important to note that, in addition to reflecting variation in the risk of accidents across industries and the number of workers employed in each industry, this distribution of accidents across industries depends on OSHA's coverage of certain industries. For example, since OSHA does not cover accidents in most mining sectors, the volume of accidents reported in mining only includes

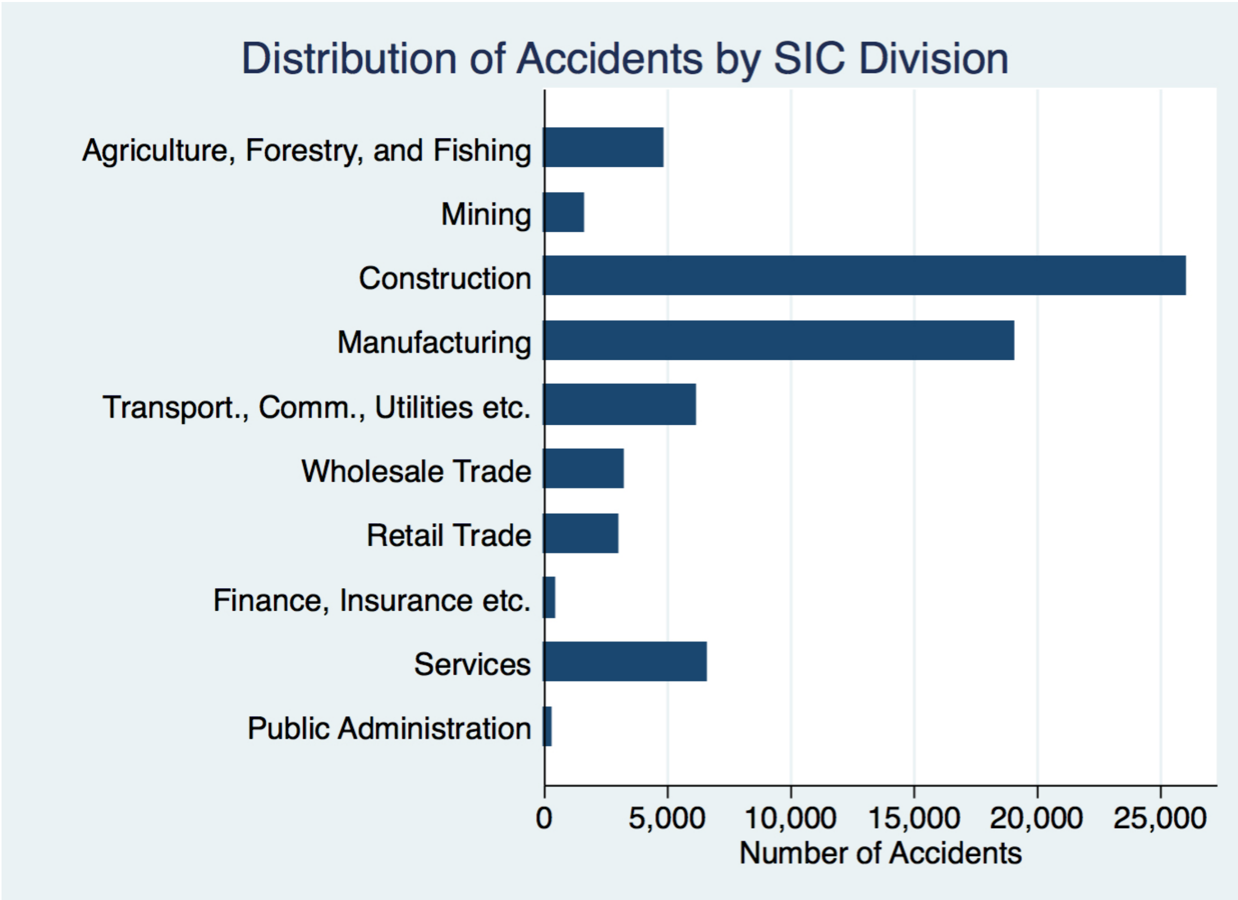

Figure 9: Number of accidents in my sample by associated industry division. My data includes a total of 71,218 private-sector accidents reported to OSHA between 1990 and 2010 in the contiguous United States. I classify accidents into divisions under the Standard Industrial Classification (SIC) system.

accidents in oil and gas extraction. Next, Figure 10 shows the distribution of accidents in my sample across months, first for accidents in non-temperature-sensitive industries and then for accidents in temperature-sensitive industries. We see that accident incidence follows a fairly constant seasonal cycle across accidents, with the number of accidents rising in the summer and falling in the winter. This cycle is fairly similar across temperature-sensitive and non-temperature-sensitive industries, though the number of accidents rises more consistently with warmer temperatures among temperature-sensitive industries.

Finally, the dark green columns in Figure 7.2.1 show the distribution of accidents in my sample across U.S. census divisions. We see that an overwhelming proportion of the accidents in my sample occurred in the Pacific U.S. census division. This divergence seems to be driven by Cali-

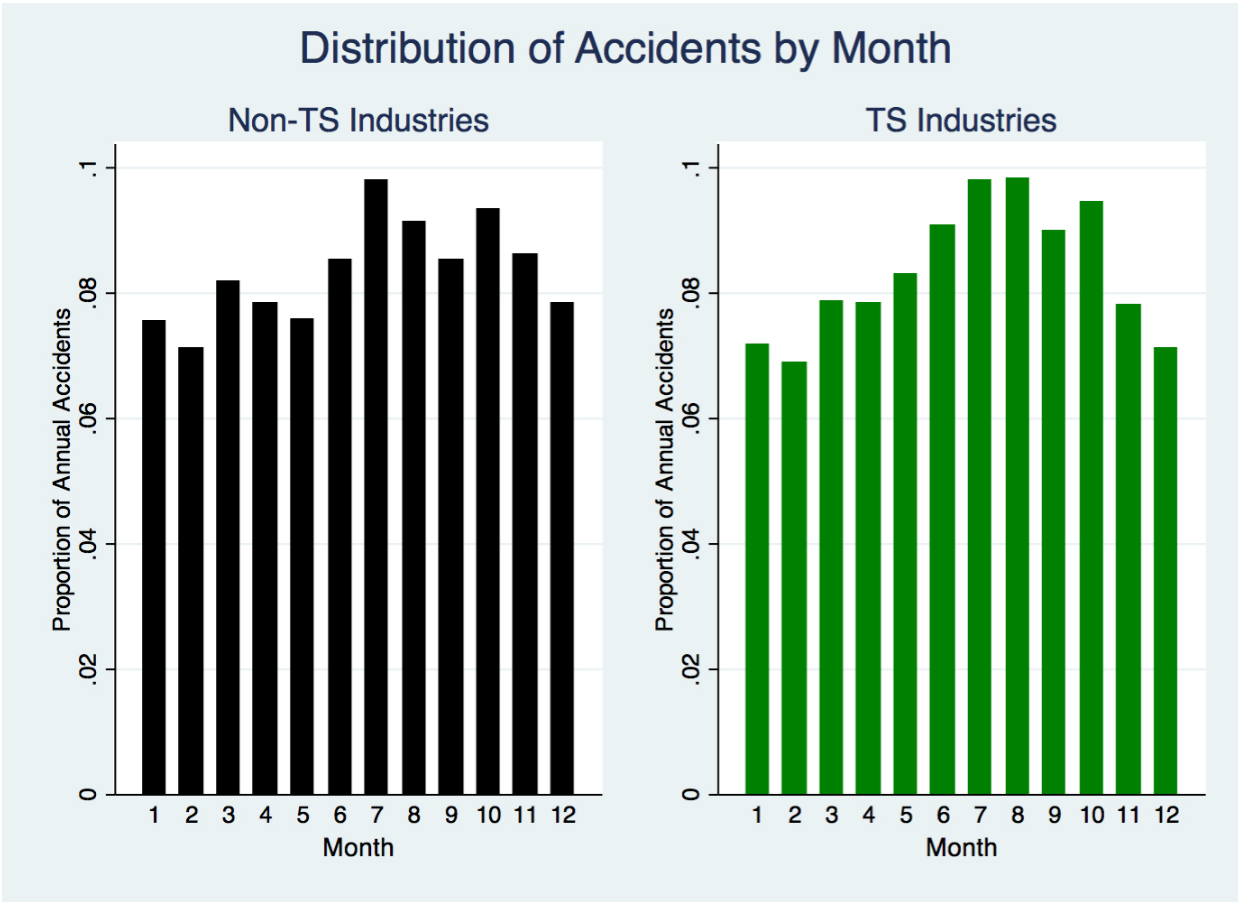

Figure 10: Distribution of accidents in my sample by month in which they occur. My data includes a total of 71,218 private-sector accidents reported to OSHA between 1990 and 2010 in the contiguous United States. I present seasonal distributions separately for non-temperature-sensitive (non-TS) industries and temperature-sensitive (TS) industries. See the Appendix for a complete list of temperature-sensitive industries.

fornia, which accounts for about 38.6% of the accidents in my sample. Geographic variation in the volume of accidents reported to OSHA likely stems from a number of sources, including industrial composition and overall labor force size. It may also relate to differences in regulations for accident reporting across regions. Twenty-two states and territories have OSHA-approved state programs, which are required to be at least as stringent as the Federal OSHA program.[3] In addition to differences in regulations on accident reporting, regions of the country likely vary in their norms for compliance with these regulations. According to an OSHA representative, the huge number of accidents reported from California is a product both of particularly stringent reporting requirements

[3]According to OSHA's Directorate of Cooperative and State Programs, there are no unified resources on differences across state OSHA regulations. Given more time, I will look through the text of state regulations to identify relevant differences in accident reporting requirements

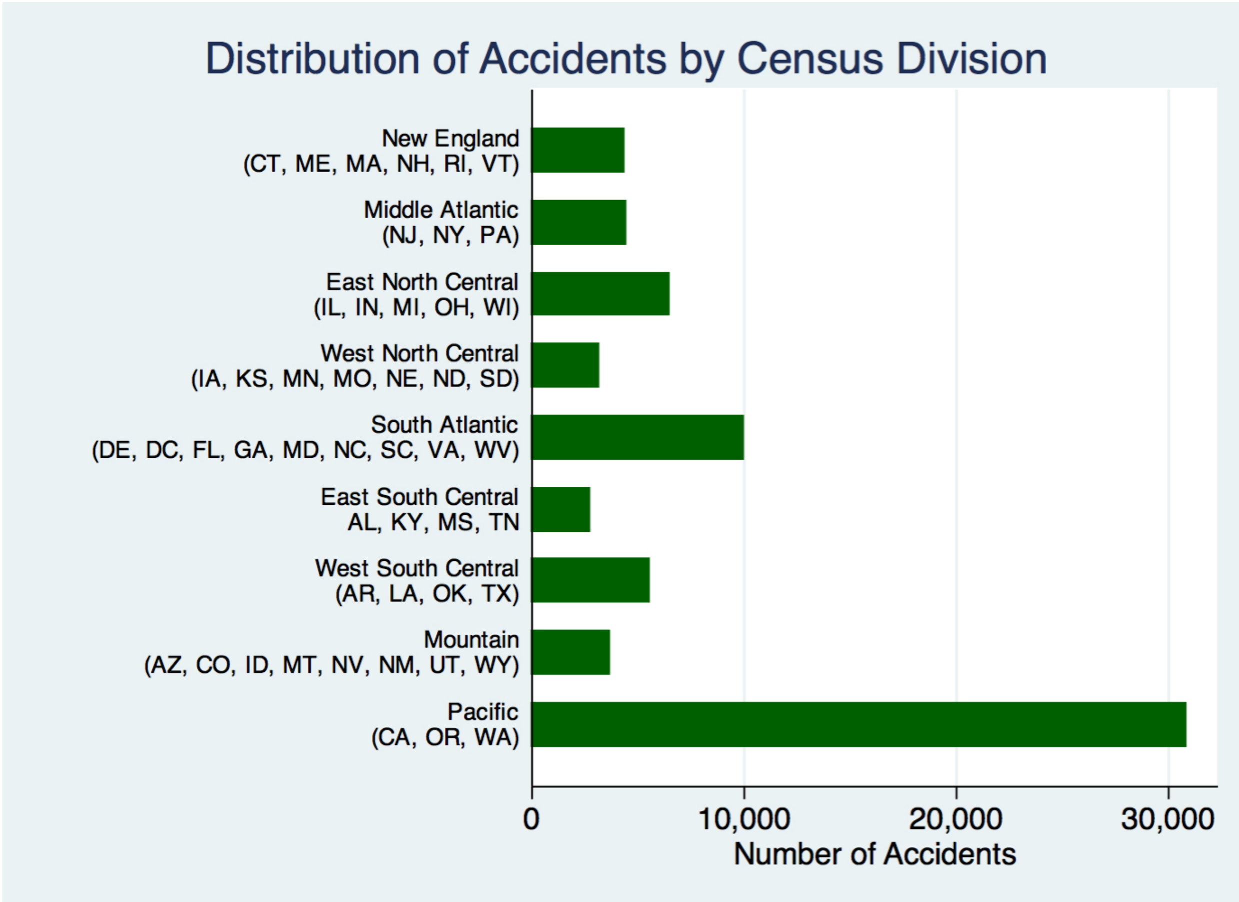

Figure 11: Distribution of accidents in my sample by Census Division in which they occur. My data includes a total of 71,218 private-sector accidents reported to OSHA between 1990 and 2010 in the contiguous United States. States included in each Census Division are given in parentheses above.

and widespread compliance with those requirements.

To investigate the extent to which variation in accident incidence across the United States reflects variation in accident reporting across regions, rather than true differences in accident incidence, I compare the proportion of accidents in my sample that occurred in each census division with the proportion of fatalities tabulated in CFOI that occurred in the same region. CFOI tabulates all workplace fatalities, with little concern of underreporting, so it should allow us to compare the true incidence of accidents across regions. The ratios of CFOI fatalities by census region are given in the light blue columns in Figure 7.2.1. Based on these sample ratios and CFOI ratios, it appears that the preponderance of accidents from the Pacific census division in my sample primarily reflects a difference in accident reporting norms or regulations, rather than a difference in accident incidence or risk.

In general, these differences in accident incidence and reporting across regions should not pose a problem for my analysis. In particular, I will control for time-invariant characteristics that impact accident reporting, including reporting regulations and norms, industrial composition, and labor force size, with county fixed effects. These fixed effects allow us to identify the impacts of temperature on accident incidence using deviations from mean accident incidence in particular geographic areas, rather than from differences in accident incidence across geographic areas. I also perform robustness analysis excluding accidents from California to ensure that my results hold elsewhere in the country.

Ultimately, I merge the records for the 71,218 accidents in my sample with weather data, generating daily counts of all accidents and accidents in temperature-sensitive industries by county from 1990 to 2010.

3.3 Employment Data

My data on employment comes from the Quarterly Census of Employment and Wages (QCEW), which is a cooperative program between the Bureau of Labor Statistics and the State Employment Security Agencies (SESAs). The QCEW monitors employment and wage information for workers covered under state unemployment insurance laws and federal workers covered under the Unemployment Compensation for Federal Employees program. Thus, the QCEW produces monthly data on number of establishments, monthly employment, and quarterly wages by industry and by county for the entire United States. While the QCEW delineates industries by SIC codes for the years 1975 through 2000, it exclusively delineates industries by NAICS codes for years following 2000. I am primarily interested in QCEW employment data, which includes the number of covered workers who worked or received pay during the pay period including the 12th of each month. It does not include members of the armed forces, the self-employed, proprietors, domestic workers, unpaid family workers, and railroad workers covered by the railroad unemployment insurance system. While OSHA does not cover several large industries that are covered by the QCEW, like mining, both program have similar coverage within industries. That is, for those industries that OSHA covers, QCEW employment data provides a good measure of the population of workers among which accidents appear in my data.

One might hope that we could use this QCEW data to construct total employment levels for each county and industry so that accident incidence could be expressed as an industry-specific accident rate. In particular, we could calculate employment totals for those industries that I have classified as "temperature-sensitive," generating data on the rate of accidents per number of workers in temperature-sensitive industries. (See the Appendix for a list of SIC and NAICS codes that I classify as temperature-sensitive.) This employment data could then serve as preliminary "denominator data" like that which I describe in Section 1.4, measuring county-level exposure to accident risk. Unfortunately, data suppression makes this approach unworkable. However, QCEW data nondisclosure practices makes this impossible; the QCEW suppresses employment data whenever disclosing it might reveal information on specific employers, so much industry-specific employment data is undisclosed at the county level.

Thus, I have developed a new method of estimating temperature-sensitive employment that pulls industry-specific data from any years in which it is available. Given county A and a particular temperature-sensitive industry, I identify all of the years for which data on employment in that county and in that industry is available. For each such year, I calculate the proportion of total employment falling in that industry, and then average this proportion across all years for which we have data on that industry in that county. Thus, for each temperature-sensitive industry, I use all available data to calculate the average proportion of employment in that industry over total employment for each county. Then, I sum these proportions over all temperature-sensitive industries to generate a measure of the average proportion of total employment in a county found in temperature-sensitive industries over my study period, 1989 to 2010. This employment ratio takes an average value of 0.419 in our sample. See Figure 12.

Throughout this process, I also create an indicator denoting the quality of this data for the temperature-sensitive employment ratio. For a county with employment data for every temperature-sensitive industry at some point in our sample period, this quality indicator takes the value of 0. Then, I add 1 to a county's quality indicator for each temperature-sensitive SIC division with no available employment data, and then add progressively smaller values to the quality indicator to

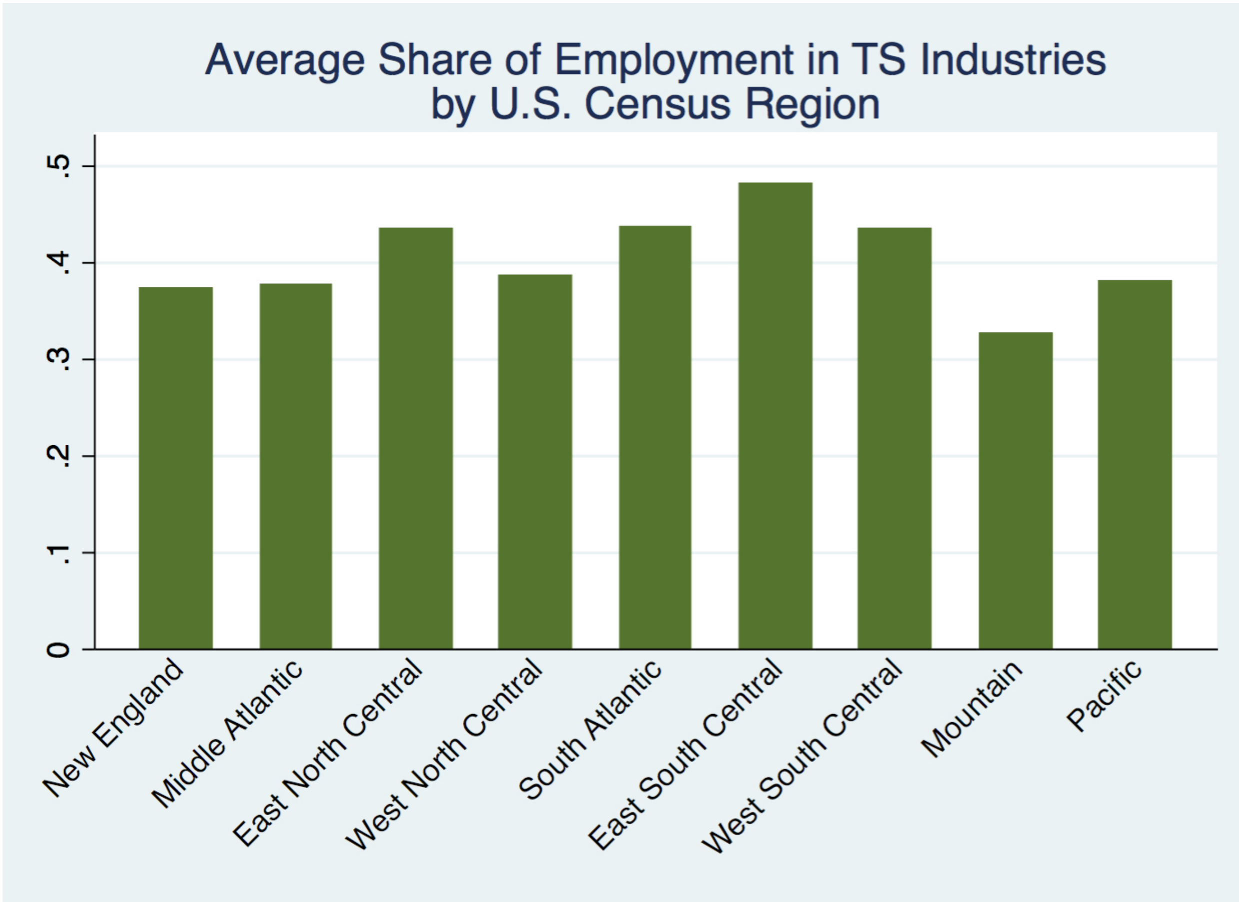

Figure 12: Average share of employment in temperature-sensitive industries by U.S. Census Division. This ratio gives the average proportion of employment in a county that is in temperature-sensitive industries. While I calculate this employment ratio separately by month in analysis, here I average across months to generate an average annual employment ratio. See text for a more detailed description of how I generate data on employment ratio. Census division are defined as follows: New England (CT, ME, MA, NH, RI, VT), Middle Atlantic (NJ, NY, PA), East North Central (IN, IL, MI, OH, WI), West North Central (IA, KS, MN, MO, NE, ND, SD), South Atlantic (DE, DC, FL, GA, MD, NC, SC, VA, WV), East South Central (AL, KY, MS, TN), West South Central (AR, LA, OK, TX), Mountain (AZ, CO, ID, NM, MT, UT, NV, WY), and Pacific (CA, OR, WA).

mark missing data for increasingly granular industries. While this approach is imprecise, I have not found alternative methods of measuring SIC data quality. This quality indicator takes an average value of 0.385 in my sample. For certain analysis, I will exclude data from counties with quality indicators for this temperature-sensitive employment data of over 1; this excludes only about 4.5% of the data.

I then generate estimates of temperature-sensitive employment by month for each county by multiplying total employment in a county, which is almost never undisclosed, by this ratio of temperature-sensitive employment to total employment. It is important to note that this method of generating data on temperature-sensitive employment assumes that local industrial structure stays constant between 1989 and 2010, even if employment totals are changing. In particular, we assume that the total proportion of county-level employment based in temperature-sensitive industries stays constant over our study period.

4 Econometric Strategy

My econometric strategy draws from both the growing economic literature on the impacts of climate change and from previous work in public health on the impacts of temperature on accident incidence.

My primary identification of the causal impacts of temperature on accident incidence draws from the empirical strategy of the growing climate-economy literature in economics. In particular, I use a set of temporal and geographic fixed effects to identify the impact of short-term, plausibly exogenous variation in weather on the incidence of occupational accidents (Dell et al., 2014). See Section 2.1 for an overview of several papers from this literature. Throughout my analysis, I fit variations of the following equation:

$$accidents_{it} \ = \ \sum_{j=1}^{12} \beta_j \ tmax_{itj} + \sum_{k=1}^{13} \lambda_k \ prcp_{itk} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it}$$

In this equation, the letter i denotes county, t denotes calendar day, y denotes year, r denotes U.S. census region, m denotes month, and d denotes day of week. I regress the temperature-sensitive accident count in county i on day t on a series of twelve 10-degree bins for maximum temperature on day t, each denoted by $tmax_j$ for j in $\{1, 2, 3, \ . \ . \ . \ 12\}$. The coefficients on these temperature indicator variables are of primary interest. These bins range from daily maximum temperature below 0°F to daily maximum temperature above 100°F. Thus, a day with maximum temperature of 92°F will take a value of 1 for $tmax_{11}$, the temperature bin for maximum temperature between 90 and 100°F, while all other temperature bin dummies will take the value 0. Using these temperature bins allows us to avoid assuming a particular functional relationship between temperature and accidents. Our only functional form restriction is that the impact of daily maximum temperature on accident incidence is constant within these 10-degree ranges. In all regressions, I omit the temperature bin corresponding to daily maximum temperature between 60 and 70°F, which physiology research suggests may be ideal working conditions (Hancock and Vasmatzidis, 2003; Ramsey and Kown, 1992; Pilcher et al., 2002). This strategy allows us to estimate the change in the incidence of accidents on a day with maximum temperature in some range relative to a day with maximum temperature between 60 and 70 °F. I control for precipitation using a similar set of indicator bins. First, I define a dummy variable indicating zero precipitation. Next, I create eleven precipitation bins divided by the values for the 10^{th}, 20^{th}, 30^{th}, 40^{th}, 50^{th}, 60^{th}, 70^{th}, 80^{th}, 90^{th}, and 95^{th} percentiles of the distribution of non-zero precipitation in my sample.

In addition to variables for precipitation and temperature, I also control for a suite of temporal and geographic fixed effects, which should allow me to identify exogenous short-term weather variation. Again, the letter i denotes county, t denotes calendar day, y denotes year, r denotes U.S. census region, m denotes month, and d denotes day of week. Thus, my most saturated fixed effect specification includes county fixed effects, α_i, year fixed effects, γ_y, region-by-month fixed effects, θ_{r*m}, and day-of-week fixed effects, ν_d. The inclusion of county fixed effects controls for time-invariant county characteristics that could impact accident rates and correlate with climate, and year fixed effects control for nation-wide average temperatures and accident counts. In particular, these year fixed effects should adjust for nation-wide changes in accident-reporting, worker safety, or climate. Next, the region-by-month fixed effects control for region-specific seasonal cycles, and the day-of-week fixed effects control for different levels of economic activity on different days of the week. In this most saturated model, I exploit county-specific variation in weather and accident rates about long-term county averages, after controlling for national averages by year and for region-specific seasonal cycles. I cluster standard errors at the county level throughout my analysis to account for potential autocorrelation in county errors.

In this panel fixed effect structure, these estimates reflect the impact of short-term local temperature fluctuations within a county on daily accident incidence. It is plausible that this variation in weather is exogenous with respect to unobserved determinants of accident incidence, so estimation of my primary regression equation will give unbiased estimates of the causal impact of each temperature bin on accident rate. The most likely exceptions to this assumption are temporal dependence within counties and cross-sectional dependence across counties. First, temperature today may correlate with temperature on recent days, which may also impact the volume of work done, and thus the incidence of accidents, today, creating cross-sectional dependence. Next, temperature is likely correlated across nearby counties on any given day. If shared forces influence the incidence of accidents in nearby counties as well, this cross-sectional dependence could create omitted variable bias. I will explore these potential sources of omitted variable bias in sensitivity analysis.

In general, the recent public health research on accident incidence and temperature uses a similar strategy to identify the causal impacts of temperature on accidents incidence. In particular, past research has used controls for seasonality and long-term trends to isolate the impacts of plausibly exogenous short-term variation in weather (Xiang et al., 2014b, 2014c; Adam-Poupart et al., 2015). However, my identification strategy drawn from the recent climate-economy literature in economics improves on the recent public health research in several key respects. First, my specifications for temperature and precipitation improve substantially on those used in recent public health studies of temperature and accident incidence. Indeed, none of the recent studies on workplace accidents or illnesses and temperature control for precipitation (Xiang et al., 2014b, 2014c; Adam-Poupart et al., 2015). Research suggests that temperature and precipitation are closely related in the climate system (IPCC AR4 3.3.5, 2007), and it is plausible that precipitation might influence accident incidence both by making outdoor work impossible and by increasing the risk of accidents for a given amount of work. By failing to control for precipitation, these public health studies may confound the impacts of temperature on accident incidence with the impact of precipitation on accident incidence.

Next, my temperature specification is more flexible than that used in the recent public health studies of accident incidence. Adam-Poupart et al. (2015) impose a linear relationship between temperature and worker compensations for heat-related illnesses, though they restrict analysis to summer months and thus only impose this linear relationship between heat illnesses and warm temperatures. Xiang et al. (2014c) allow for only limited nonlinearity in the relationship between temperature and accident incidence using a piecewise linear spline function with a threshold at 37.7°C. In contrast, my structure of 10°F temperature bins allows me to flexibly model the nonlinear relationship between temperature and accident incidence, imposing weaker parametric assumptions. I explore the implications of using even fewer parametric assumptions in sensitivity analysis with 5°F temperature bins.

Next, my panel fixed effect structure allows me to analyze a broader geographic study area than do the recent public health studies of the relationship between temperature and accident incidence. So far, each of these studies has analyzed city- or region-level accident counts; Morabito et al. (2006) study Tuscany, Italy, Xiang et al. (2014b, 2014c) study Adelaide, Australia, and Adam-Poupart et al. (2015) study the region of Quebec. In contrast, my fixed effect structure allows me to analyze accident incidence across the contiguous United States. The inclusion of county fixed effects allows me to control for time-invariant county characteristics, preventing confounding by cross-county differences as my geographic sample expands. This broader geographic scale allows me to make use of greater variation in my weather variables and improves the external validity of my estimates for the impact of temperature on accident incidence (Deschênes, 2014).

While my identification strategy is largely modeled on the recent body of work in economics on the impacts of climate change, it also draws from the recent public health research on temperature and accidents in several important ways. First, like the recent public health research on temperature and accident incidence, I analyze daily accident counts. In contrast, most economic analysis of the impacts of climate change uses annual or monthly data on economic outcomes (Deschênes, 2014). In studies of temperature and mortality, this longer exposure window helps to prevent confounding by harvesting, or the accelerated death of those already nearing death due to chronic conditions. To the extent that accidents occur randomly within a population of workers, and not among a subset of workers that were "due" to have an accident regardless of weather conditions, harvesting is unlikely to be relevant to the incidence of workplace accidents. Then, daily-level analysis allows me to more carefully characterize the dynamics of the relationship between temperature and accident incidence (Deschênes, 2014), as in recent public health studies of temperature and accidents (Xiang et al., 2014b, 2014c; Adam-Poupart et al., 2015).

Next, while much of the climate-economy literature uses a standard OLS regression framework, I also incorporate the Poisson regression framework common to public health analysis. Adam-

Poupart et al. (2015) and Xiang et al. (2014a) use negative binomial regression, a variant of Poisson analysis that allows for overdispersion in the outcome variable. Poisson regression is well-suited to analysis of count data, like that on accident incidence, which takes only the values of natural numbers and in which zero-valued observations are common. This form of analysis allows us to estimate the relative risk of accidents associated with different temperature conditions (Deschênes, 2014), and it is a more suitable framework for analysis of accident incidence than are the standard models of the climate-economy literature. See Section 5.2.1 for a more detailed discussion of Poisson regression.

In general, then, my econometric strategy allows me to estimate the causal impact of temperature on the incidence of occupational accidents across the United States. My analysis draws from both the body of climate-economy research in economics and recent public health analysis of accidents and temperature. In doing so, I present the first analysis of accident incidence and temperature to make use of a broad geographic spatial scale and to flexibly model the impacts of temperature, while retaining a framework of daily-level Poisson analysis that is well-suited to analysis of accident counts.

5 Primary Results

5.1 Primary OLS Analysis

I begin by performing a simple OLS regression of daily accident count in a county on the twelve 10-degree bins for daily maximum temperature, precipitation controls, and fixed effects. The results of this regression are plotted graphically in Figure 13 and appear in a regression table in the Appendix. The coefficients on these temperature bins are broadly significant, and they trace out an interesting U-shaped curve between temperature and accident incidence. Recall that in this regression and throughout my analysis, I omit the bin for maximum temperature between 60 and 70 °F, so coefficients on temperature bins give the number of additional accidents in temperature-sensitive industries in a county on a day with maximum temperature in that bin relative to the a day with maximum temperature with maximum temperature between 60 and 70 °F. Thus, we see that days with maximum temperature in all temperature bins above 70°F and all temperature bins below 50°F

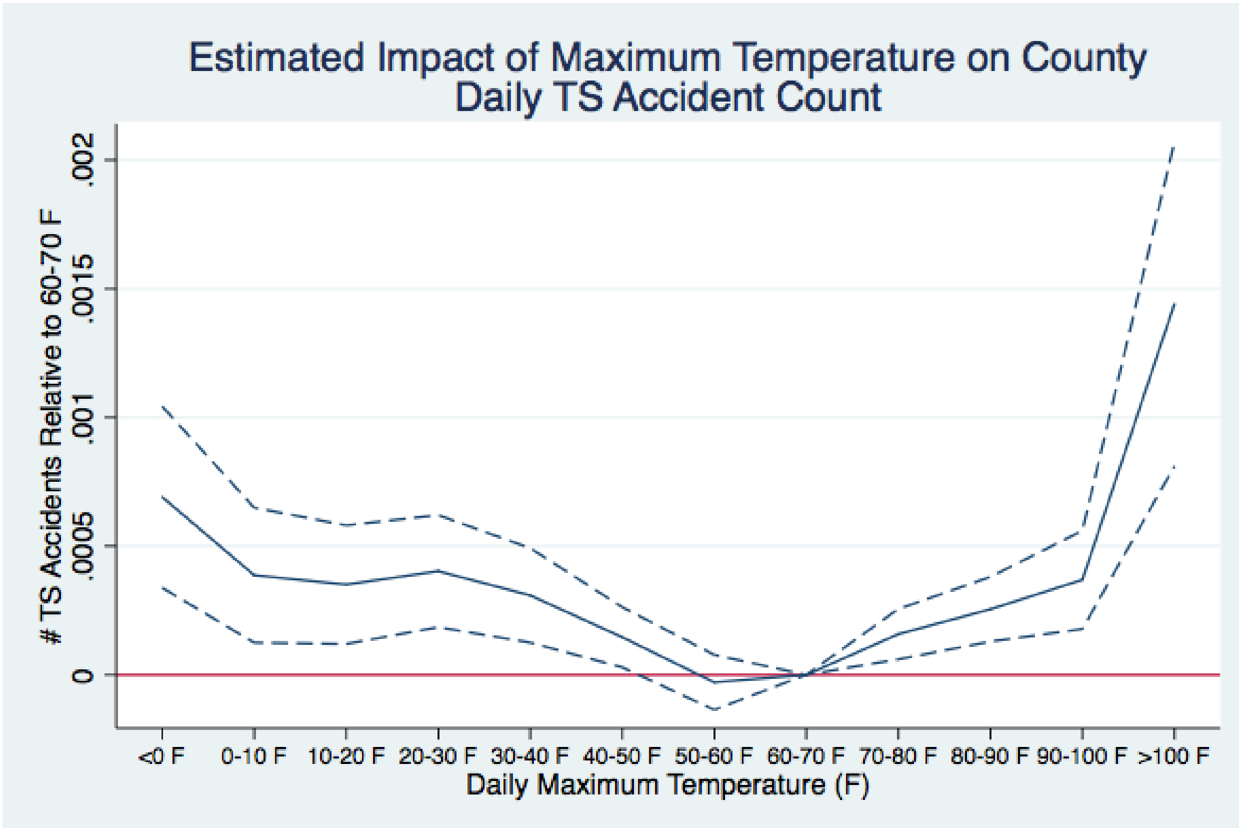

Figure 13: OLS estimates of the impact of daily maximum temperature and county daily accident count in temperature-sensitive industries. N=23,614,877. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects.

are associated with significantly more accidents than days with maximum temperature between 60 and 70°F, where the magnitude of this increase rises as temperatures become more extreme. For example, a day with maximum temperature over 100°F is associated with an additional 0.00144 accidents in temperature-sensitive industries per county on that day and a day with maximum temperature below 0°F is associated with an additional 0.000691 accidents in temperature-sensitive industries per county on that day, all relative to a day with maximum temperature between 60 and 70°F.

Thus, this OLS regression suggests that very hot and very cold days significantly increase the incidence of accidents in temperature-sensitive industries, with the magnitude of this increase rising as temperatures become more extreme. While these coefficients appear to be quite small in magnitude, they are large relative to the mean number of accidents per county per day. Our data only includes particularly severe accidents resulting in worker hospitalization or death, so they are rare events. On average, counties in our sample have an average of about 0.00250 OSHA-reported accidents per day, so an additional 0.00144 accidents on days with maximum temperature over 100°F is about a 57.6% increase over mean accident incidence. If each of the 3108 county-equivalents in the contiguous United States were to experience an additional day with maximum temperature over 100°F instead of a day with maximum temperature between 60 and 70°F, we would expect to see about 4.5 additional severe nonfatal or fatal accidents reported to OSHA on that day.

5.1.1 Issues with Primary OLS Analysis

While these results provide suggestive evidence in support of my hypothesis that extreme temperatures increase the incidence of workplace accidents, this OLS regression with county-level accident count as the dependent variable has several limitations. First, OLS regression is generally inappropriate for analysis of count data. While OLS models assume that the outcome variable can take any real value, including negative and fractional values, count data only takes the values of natural numbers. Thus, OLS may predict values for the outcome variable that are theoretically impossible (Wooldridge, 2001). Next, OLS assumes a linear relationship between the independent and dependent variables of the form $E[y|x] = \beta_0 + \beta_1 X_1 + \beta_2 X_2 ...$, making the assumption that the difference between the occurrence of zero and one event in a time interval is equivalent to the difference between the occurrence of 15 and 16 events in the same time interval. This assumption is likely to be inappropriate for count data like that for occupational accidents (King, 1988). Furthermore, OLS is an inefficient estimator with counts, since it does not incorporate information on the particular form of heteroskedasticity of the count data, which tend to be quite right-skewed, or the functional form of the underlying Poisson distribution of events (King, 1988).

Even if these estimates for the impacts of temperature on accident incidence are accurate, however, these regressions with counts are only somewhat useful in that their estimated coefficients depend closely on the size of the working population used in analysis. Consider county A and county B, where county A has 1 million people working in outdoor industries and where county

B has only 100 people working in outdoor industries. Say that individuals have a 10% chance of having a workplace accident on extremely hot days and a 5% change of having an accident on average-temperature days. relative to average-temperature days. Then, on an extremely hot day, we would expect to see 50,000 additional accidents in county A, while we only expect to see 5 additional accidents in county B. Thus, even after controlling for counties' average accident counts and temperature conditions with county fixed effects, we would expect to see larger magnitude deviations from those mean accident counts in response to extreme weather in large counties than in small counties. In the simple regression sketched out above, we would expect to see substantially larger values for β_i in large counties than in small counties.

Then, my estimates of these coefficients give average regression correlations. Say we take a simple fixed effect model for a panel of counties indexed by i, where x is a dummy variable indicating an extremely hot day and where we include county fixed effects:

$$y_{it} = \alpha_i + \beta x_{it} + \epsilon_{it}$$

We could equivalently express this equation as

$$y_{it}* = \beta x_{it} * + \epsilon(it) \quad \text{where} \quad y_{it}* = y_{it} - \bar{y_i} \quad \text{for} \quad \bar{y_i} = T^{-1} \sum_{t=1}^{T} y_{it}$$
$$\text{and} \quad x_{it}* = x_{it} - \bar{x_i} \quad \text{for} \quad \bar{x_i} = T^{-1} \sum_{t=1}^{T} x_{it}$$

Then, the pooled fixed effects estimator is given by

$$\hat{\beta}_{pool} = \frac{\frac{1}{N}\sum_{i=1}^{N}(\frac{1}{T}\sum_{t=1}^{T} x_{it}^{*} y_{it}^{*})}{\frac{1}{N}\sum_{i=1}^{N}(\frac{1}{T}\sum_{t=1}^{T} x_{it}^{*2})} \rightarrow \frac{\frac{1}{N}\sum_{i=1}^{N}\Omega_{i,xy}}{\frac{1}{N}\sum_{i=1}^{N}\Omega_{i,xx}}$$

$$\hat{\beta}_{pool} \rightarrow \frac{\frac{1}{N}\sum_{i=1}^{N}\Omega_{i,xy}}{\frac{1}{N}\sum_{i=1}^{N}\Omega_{i,xx}} = \frac{\bar{\Omega}_{xy}}{\bar{\Omega}_{xx}}$$

In other words, the β_{pool} estimator gives the average regression correlation, or the ratio of the average county covariance over the average county variance.

While these coefficients are still useful as a sort of average impact across counties of different sizes, their dependence on county size means that they do not allow for meaningful comparisons across different areas of the country. That is, if we estimate that an extremely hot day has a larger magnitude impact in the West than in the Northeast, it may be simply because there is more employment in temperature-sensitive industries or because more accidents are reported in the West than the Northeast, rather than reflecting true differences in the impact of temperature on the risk of accidents across regions. Using a dependent variable that scales accident incidence by a measure of *opportunity* for those accidents to occur, like labor force size, would allow me to generate more meaningful estimates for the impact of temperature on accident incidence.

5.1.2 OLS Analysis Accounting for Employment

As a first pass at controlling for the role of employment volume in my results so far, I simply add a variable to my primary OLS regression controlling for the level of temperature-sensitive employment in a county. The upper graph in Figure 14 gives the results from my primary OLS regression, and the lower graph presents the results of this regression with the addition of a control for temperature-sensitive employment. The results of both regressions are also available in the Appendix. In the lower graph, then, the coefficients on temperature bins give the impact of days in each temperature category on accident count after removing the impacts of variation in the level of temperature-sensitive employment. We see that the coefficients on the temperature bins are largely robust to this addition, though they change slightly in magnitude. Thus, it appears that the impacts of temperature on accident count that we observe are not driven by monthly-level variation in employment. In the absence of daily-level employment data, we cannot investigate the possibility that temperature impacts are driven by within-month employment changes, which could be relevant for industries staffed by day laborers. It is important to note that this regression does not solve the principal issues in our primary count-based OLS regression, in that, even after controlling for employment, we would still expect to estimate very different impacts of temperature on accident count in areas of different sizes.

To more fully address this concern, I attempt to estimate the impact of temperature on the rate of accidents by county, where I calculate this rate first as accidents over total employment by county and then as accidents over temperature-sensitive employment by county. See section 2.3 for a detailed explanation of how I generate this data for temperature-sensitive employment.

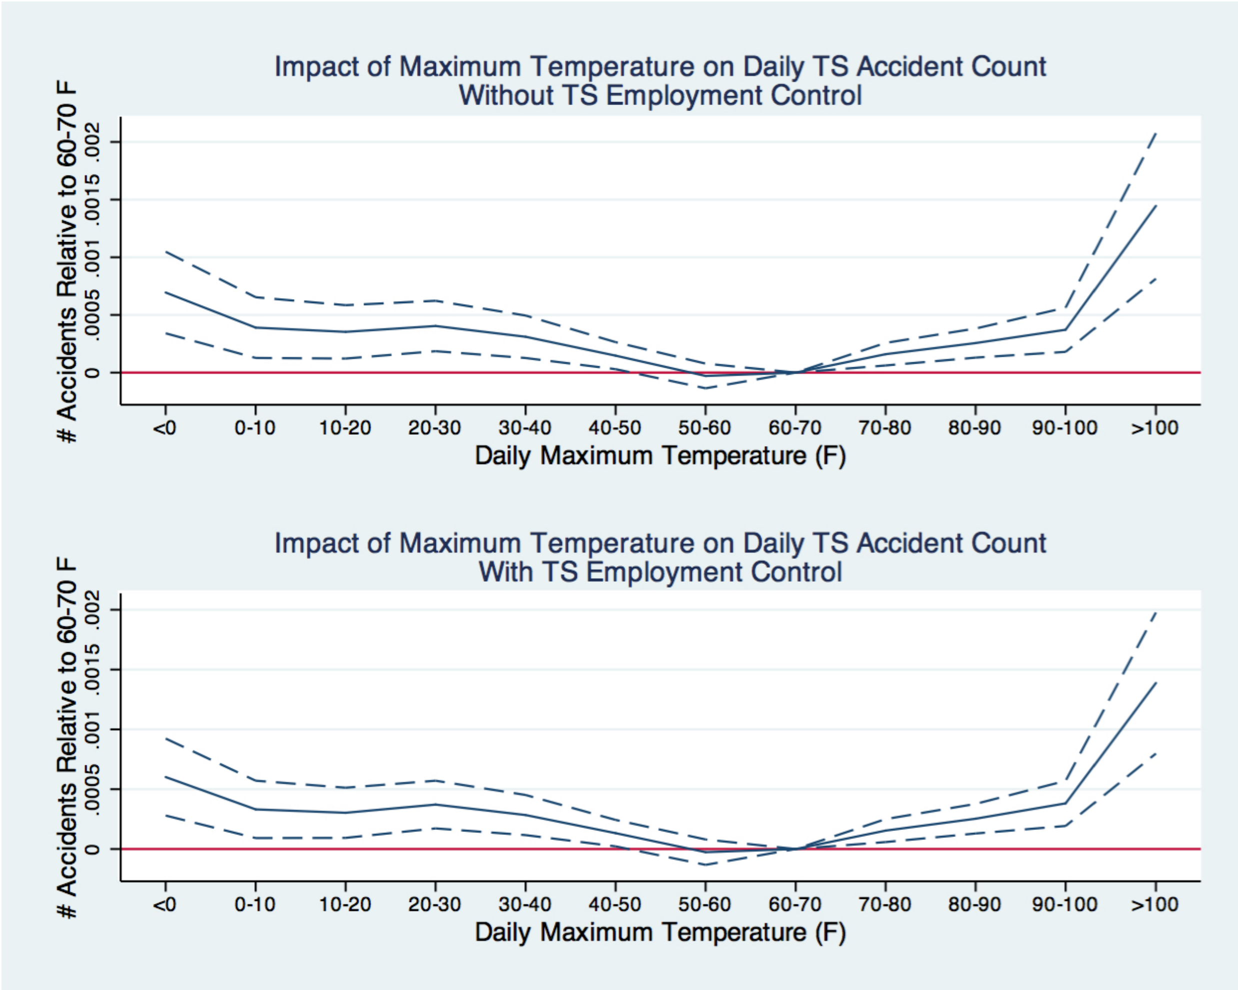

Figure 14: OLS estimates of the impact of daily maximum temperature on county-level daily accident count in temperature-sensitive industries, both with and without a linear control for temperature-sensitive employment. Coefficients give the number of additional accidents relative to a day with maximum temperature between 60 and 70°F. N=23,527,479 in both regressions. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects.

Theoretically, we might expect the impact of temperature on accident rate to be constant across counties of different sizes. In Figure 15, I present the results of a regression where I regress dependent variables of temperature-sensitive accidents over total employment and temperature-sensitive accidents over temperature-sensitive employment, respectively, on the set of independent variables from my primary OLS regression. In these regressions, the coefficients on the temperature bins become largely insignificant, and many of them have opposite sign to the corresponding coefficients

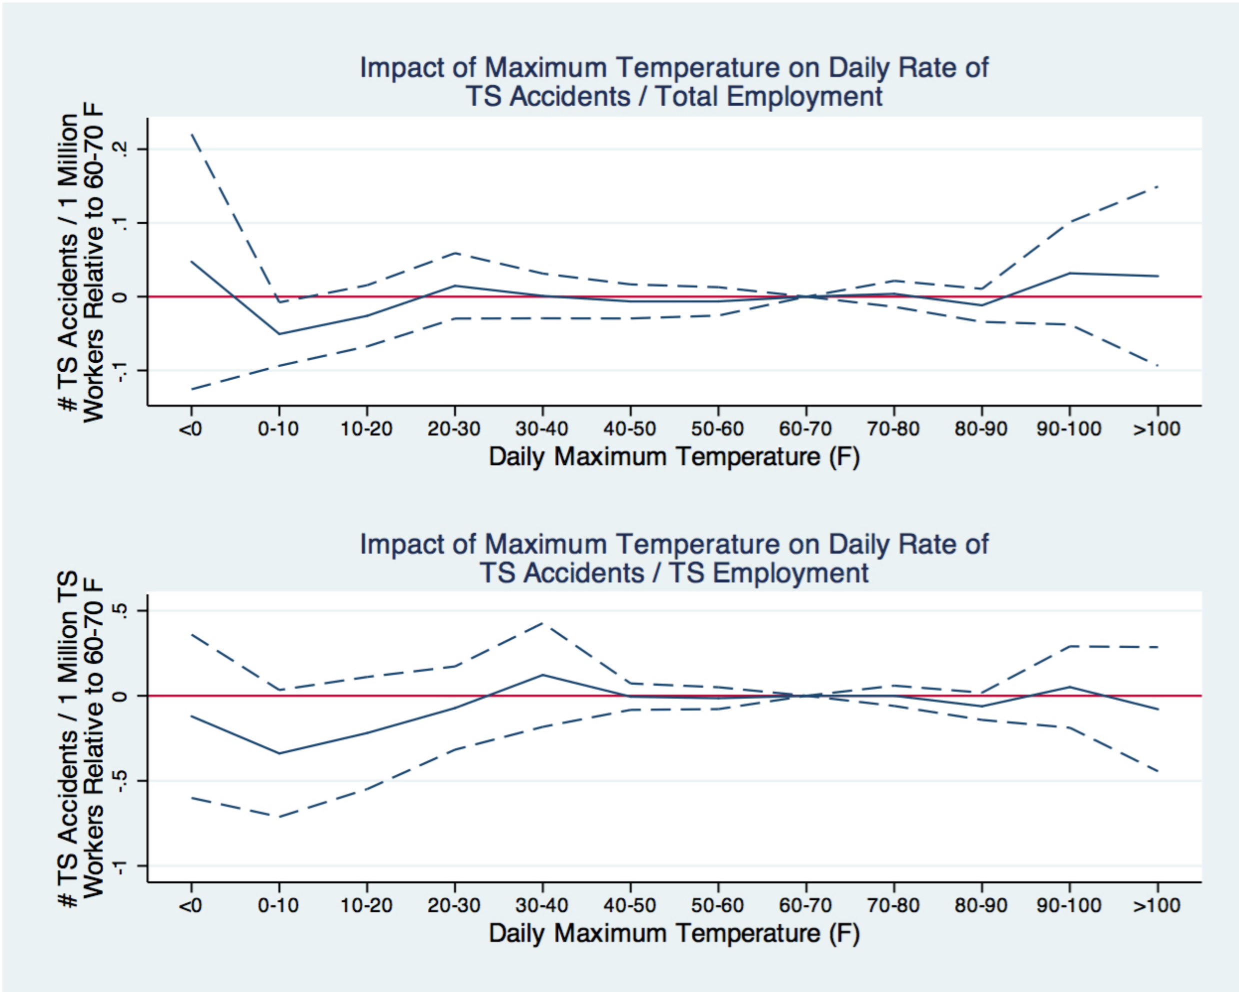

Figure 15: OLS estimates of the impacts of daily maximum temperature on the rate of temperature-sensitive accidents over total employment and on the rate of temperature-sensitive accidents over temperature-sensitive employment at the county level. Coefficients give change in accident rate relative to a day with maximum temperature between 60 and 70°F. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects.

in Figure 14. Again, the results of these regressions are available in tabular form in the Appendix.

Why do my results largely disappear when I attempt to use an accident rate as my dependent variable in an OLS regression? First, this may be because the relationship we are seeing between temperature and accidents is actually explained by the relationship between temperature and employment. Perhaps increases in the number of accidents at extreme temperatures are actually just the result of an increase in the number of people working at more extreme temperatures, after controlling for seasonal cycles. My regression in the lower graph of Figure 14 suggests that this explanation does not hold, since my estimates of the impacts of temperature bins on accident count survive the addition of a control for temperature-sensitive employment. However, I more fully investigate this possibility by regressing the natural log of monthly employment in a county on my temperature bin variables; I perform this regression using both temperature-sensitive and total employment as my dependent variable. I run these regression at the monthly level, redefining each temperature bin variable as the number of days in each temperature bin in a month. I then control for precipitation with the same percentile-based bin variables that I employ in my primary regressions, and include county, year, and region-by-month fixed effects.

I present the results of this regression for temperature-sensitive employment graphically in Figure 16. The Appendix also includes parallel results for a regression of total employment in Table 11. Here, we see that both monthly total and temperature-sensitive employment appear to decrease with an additional day in any temperature bin in a month, relative to a day with maximum temperature between 60 and 70°F. While these coefficients are small in magnitude–an additional day with maximum temperature over 100°F is associated with a decrease of 0.107% in temperature-sensitive employment, for example–they are broadly significant and increase in magnitude as temperature becomes more extreme. Thus, these results suggest that the impact of temperature on the level of employment is not explaining the increase in accident counts at extreme temperatures, since accident count increases even as employment decreases at extreme temperatures. In fact, these results suggest that our estimates for the impact of temperature on accidents may actually understate the impact of temperature on the risk of workplace accidents.

It is important to note that these regressions of employment on temperature do not capture the impact of temperature on industrial composition, so it is possible that changes in industrial composition that correlate with temperature might actually give rise to our findings that temperature increases the incidence of occupational accidents. For example, in a contrived world where people switch from working as secretaries to the dangerous occupation of leading white-water rafting tours in extremely hot weather, one would obtain the spurious conclusion that hot weather directly causes accidents. While this possibility ties closely to my broad concerns over the role of

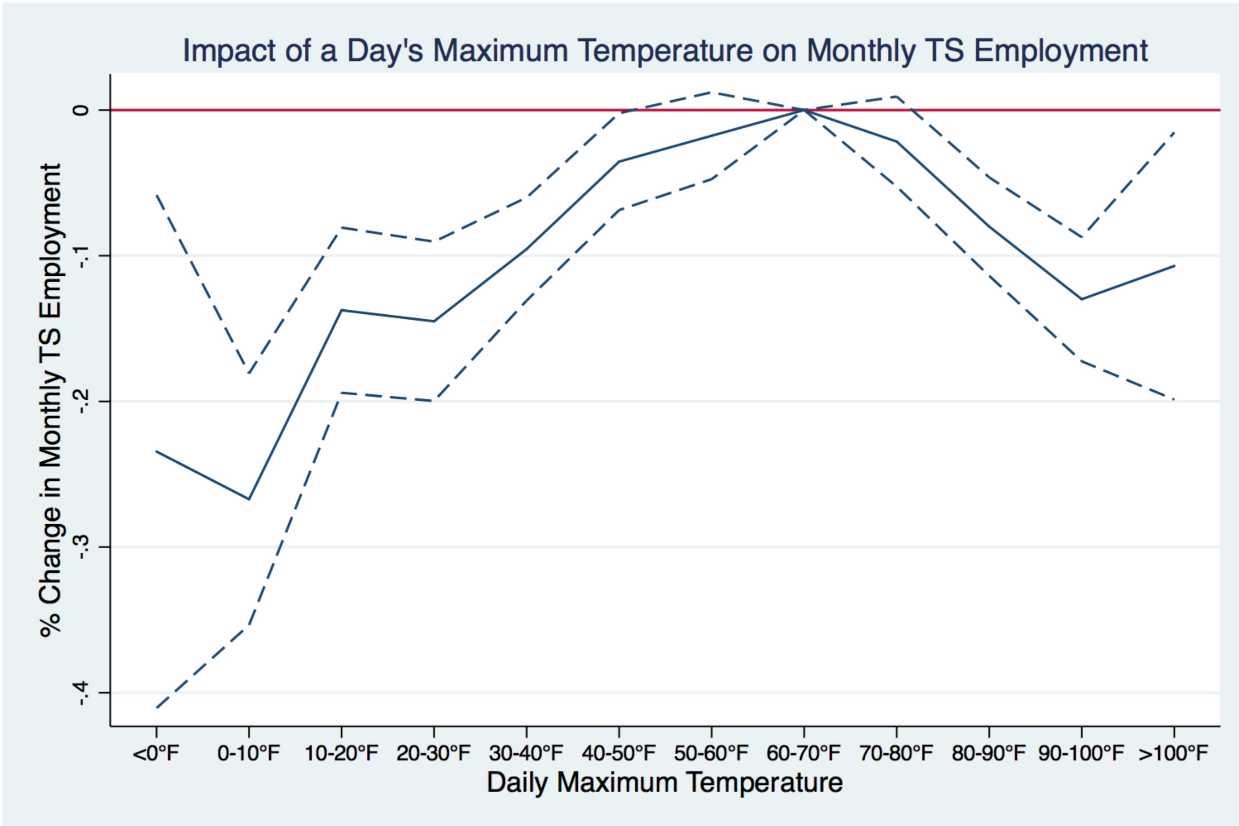

Figure 16: OLS estimates of the impact of a day's maximum temperature on county-level monthly employment in temperature-sensitive industries. Coefficients give the % change in employment level from replacing a day in a month with maximum temperature between 60 and 70°F with a day in another temperature bin. N=773,001 in both regressions. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes precipitation bins and county, year, and region-by-month fixed effects.

behavioral changes in my estimates of the impact of temperature on accident incidence, it would not explain the disappearance of my results when I use accident rate as my dependent variable. These estimates also do not capture the impacts of temperature on within-month employment, which could drive the impacts of temperature on accident incidence that we observe. The impacts of within-month employment level changes would likely be relevant, if at all, only in industries staffed by day laborers.

Next, it is possible that my results for the impact of temperature on accident incidence disappear when I divide by employment because this specification introduces irrelevant noise. First, labor force size may not be a relevant measure of the number of opportunities for heat-related occupational accidents. For example, while New York City has a very large labor force, most of those workers are working indoors and would not contribute to our counts of temperature-sensitive accidents. By using temperature-sensitive accidents divided by total employment as our dependent variable, we impose a strict relationship between accidents and employment. This specification may be inappropriate, and may therefore produce noise that obscures the true relationship between temperature and accident incidence.

This concern should be less relevant in the regression presented in column 4, since we would expect temperature-sensitive employment to be a decent measure of the volume of workers susceptible to temperature-sensitive occupational accidents. However, using a dependent variable of temperature-sensitive accidents over temperature-sensitive employment may also introduce noise in that my data for temperature-sensitive employment is based on approximations. In particular, I created my data for temperature-sensitive accidents by multiplying total employment by a time-invariant ratio of temperature-sensitive to total employment by county. Thus, we may have that these total and temperature-sensitive employment levels are either inappropriate or poorly-measured estimates of the number of opportunities for temperature-sensitive accidents, and that the noise they create prevents us from estimating the true relationship between temperature and accident incidence.

In any case, it appears that OLS regressions with accident rate as the dependent variable will not allow me to produce estimates of the impact of temperature on accident incidence that sensibly account for the number of opportunities for temperature-sensitive accidents across counties of different sizes. Thus, I proceed with analysis using accident count as my dependent variable.

5.2 Primary Poisson Analysis

However, OLS regression is ill-equipped to analyze count models, which are discrete and typically nonlinear. Thus, I will proceed with this analysis using the Poisson model, which is well-suited to analysis of count data.

5.2.1 The Poisson Model

A random variable Y takes a Poisson distribution with rate or intensity parameter $\mu > 0$ if it takes integer values $y = 0, 1, 2, ...$ with probability given by

$$Pr(Y = y) = \frac{e^{-\mu}\mu^{y}}{y!}, \quad y = 0, 1, 2, 3, ...$$

Then, we have that $E(Y) = \mu$ and $Var(Y) = \mu$, so the model is intrinsically heteroskedastic (Cameron 2009).

We can think of the univariate Poisson distribution as describing the number of occurrences of the event y over a fixed time window. We suppose that events occur randomly in time such that the probability of at least one occurrence of the event in a given time interval is proportional to the length of the interval, the probability of more than one event occurrences in a very small time period is negligible, and the number of occurrences of the event in disjoint time intervals are mutually independent (Rodríguez, 2007). Then, we can describe the probability distribution of the number of event occurrences in a fixed time interval as a Poisson distribution with mean $\mu = \lambda t$, where λ gives that rate of event occurrences per unit of time and t gives the length of the time interval. Thus, Poisson estimation is commonly used for analysis of "Poisson processes," like the arrivals of cars to an intersection, the number of bomb strikes in particular areas, or counts of radioactive disintegrations (Rodríguez, 2007).

One particularly useful feature of the Poisson distribution is that the sum of independent Poisson random variables is also Poisson. In particular, say that we have a group of n individuals with identical covariate values, let Y_i give the number of events experienced by individual i, and let Y give the number of events experienced by the full group of n individuals. Then, assuming that these individuals are independent, if $Y_i \sim P(\mu)$ for $i = 1, 2, ..., n$, then $Y \sim P(n\mu)$. That is, the group total Y is Poisson with mean $n\mu$, so we will estimate the same likelihood function if we work with individual or group counts (Rodríguez, 2007).

In the Poisson model, we model the exponential mean parameterization as

$$E[y_i|x_i] = \mu_i = \exp(x_i'\beta) = \exp(\beta_1 + \beta_2 x_{2i} + ... + \beta_k x_{ki}).$$

Based on this specification, μ is restricted to positive values. Here, the exponentiated regression coefficient $\exp(\beta_{\mathrm{j}})$ gives the multiplicative effect of the j-th predictor on the mean, so increasing x_j by one unit multiplies the mean by $\exp(\beta_j)$ (Cameron, 2009).

5.2.2 Primary Poisson Model Results

Here, I perform a panel Poisson regression of daily accident count in a county on the same suite of explanatory variables that I used in my primary OLS analysis. That is, I regress accident count on twelve 10-degree bins for daily maximum temperature, precipitation controls, and year, region-by-month, day-of-week, and county fixed effects. My basic model is

$$E[\text{accidents}_{it}|x_{it}] = \exp(\sum_{j=1}^{12} \beta_j \text{ tmax}_{itj} + \sum_{k=1}^{13} \lambda_k \text{ prcp}_{itk} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it})$$

where the letter i denotes county, t denotes calendar day, y denotes year, r denotes U.S. census region, m denotes month, and d denotes day of week. As in my OLS analysis, I omit the dummy variable for maximum temperature between 60 and 70°F. I make one key modification to this Poisson model in that I account for temperature-sensitive employment by county as an "exposure" variable in the Poisson model. In general, an exposure variable measures the volume of opportunities for an event to occur, where we expect the volume of events to increase proportionally with this exposure variable. For example, we would expect to see twice as many accidents among twice as many workers, all else equal. If we refer to this variable for temperature-sensitive employment as z, we estimate $E[\text{accidents}_{i,t}|x_{it}, z] = z \exp(x'_{it}\beta) = \exp(lnz + x'_{it}\beta)$. If we take the natural log of this expression, we get

$$ln(\text{accidents}_{it}) = ln(z_{it}) + x'_{it}\beta \iff ln(\text{accidents}_{it}) - ln(z_{it}) = x'_{it}\beta \iff$$
$$ln\left(\frac{\text{accidents}_{it}}{z_{it}}\right) = x'_{it}\beta$$

Thus, the coefficient on $ln(z)$ is constrained to be 1, allowing us to effectively estimate the impact of our independent variables on accident rate. Then, I could rewrite my primary Poisson model as

$$E[\text{accidents}_{it}|x_{it}] = \exp(ln(z_{it}) + \sum_{j=1}^{12} \beta_j \text{ tmax}_{itj} + \sum_{k=1}^{13} \lambda_k \text{ prcp}_{itk} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it})$$

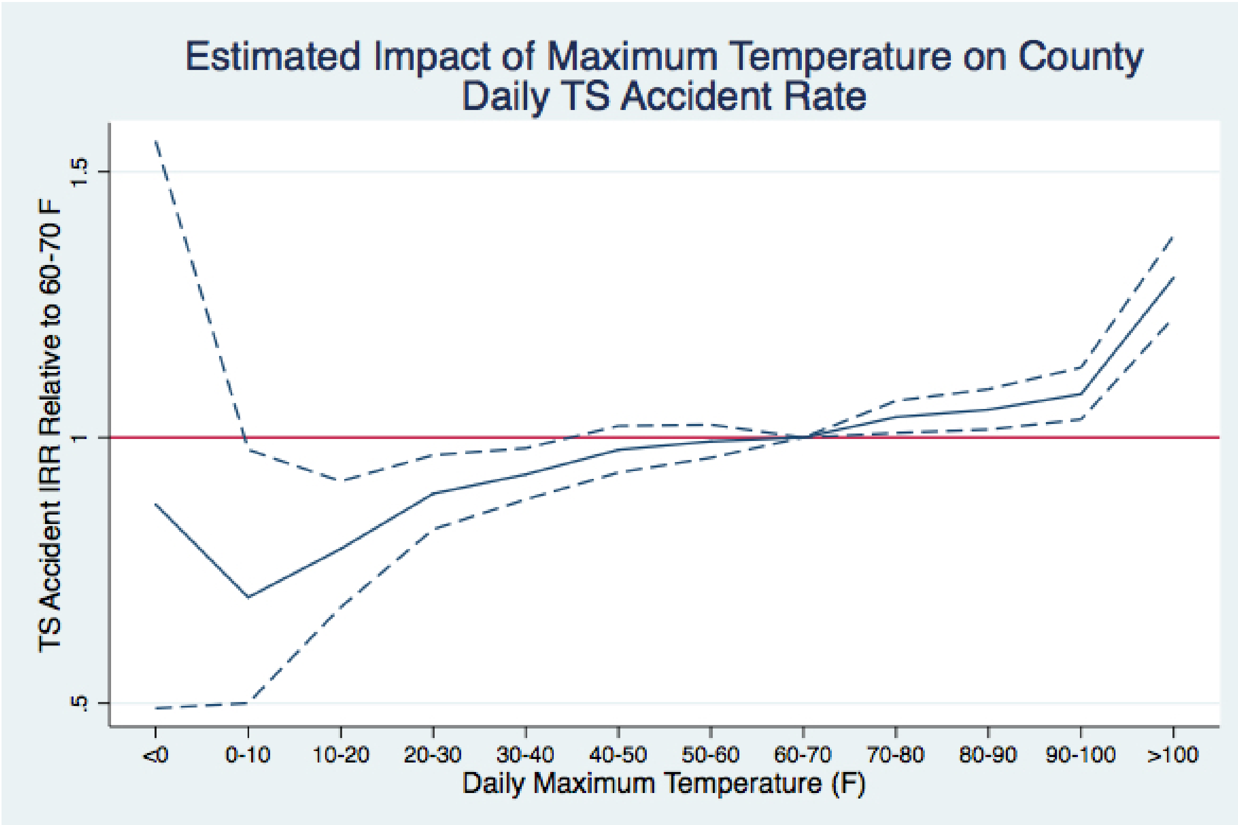

Figure 17: Poisson estimates of the impact of daily maximum temperature on the daily rate of accidents in temperature-sensitive industries. Coefficients give incident rate ratios relative to maximum temperature between 60 and 70°F. N=20,955,316. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes precipitation bins and county, year, and region-by-month fixed effects.

As in my primary OLS analysis, I cluster standard errors at the county level. I present the results of this regression in graphical form in Figure 17 and in a regression table in Table 12 in the Appendix. Here, coefficients on temperature bins give incident rate ratios, or the ratio of the rate of accidents on a day with maximum temperature in that range over the rate of accidents on a day with maximum temperature between 60 and 70°F, our omitted category. For example, the coefficient of 1.038 on the dummy variable for maximum temperature between 70 and 80°F means that days with maximum temperature in that range have 3.8% more accidents than do days with maximum temperature between 60 and 70°F. As in my primary OLS analysis, we see that hot days are associated with significantly higher accident rates. In particular, days with maximum temperature between 70 and 80°F have 3.8% more accidents, days with maximum temperature between 80 and 90°F have 5.2% more accidents, days with maximum temperature between 90 and 100°F have 8.2% more accidents, and days with maximum temperature over 100°F have 30.0% more accidents, all relative to days with maximum temperature between 60 and 70°F. Thus, this impact appears to be nonlinear, increasing slowly with rising temperature before spiking up at extremely high temperatures.

These results are qualitatively similar to my OLS estimates of the impact of high temperature on accident count, where we also see nonlinear increases in accident count associated with very hot days. These results are also of comparable magnitude. For example, my OLS analysis suggests that days with maximum temperature over 100°F have about 0.00146 additional accidents relative to days with maximum temperature between 60 and 70°F. Since days with maximum temperature have an average of 0.00288 accidents in my sample, this represents an increase of about 50.7% in accident incidence. This estimate is roughly similar to my corresponding Poisson estimate, 30.0%. Similarly, while my OLS regression estimates that days with maximum temperature between 90 and 100°F have about 12.4% more accidents than days between 60 and 70°F, my corresponding Poisson estimate is 8.2%.

However, my Poisson estimates for the impact of low temperatures on accident incidence differ sharply from those of my OLS analysis. In particular, while my OLS regressions estimated that very cold days are associated with significantly more accidents than days with maximum temperature between 60 and 70°F, my Poisson analysis suggests that cold days tend to have no significant impact on accident incidence or significantly reduced accident incidence. That is, I estimate that days with maximum temperature between 30 and 40°F have 7.0% fewer accidents, days with maximum temperature between 20 and 30°F have 10.6% fewer accidents, days with maximum temperature between 10 and 20°F have 21.0% fewer accidents, and days with maximum temperature between 0 and 10°F have 30.1% fewer accidents, all relative to a day with maximum temperature between 60 and 70°F. Since Poisson analysis better accounts for the discreteness and nonlinearities of count data, it is likely that these results are more accurate than those of my OLS analysis.

Overall, then, my results suggest that while particularly cold days are associated with reduced accident incidence in temperature-sensitive industries, hot days are associated with high accident incidence. These impacts become more pronounced as temperature becomes more extreme, rising to a 30.0% increase in accident counts on days with maximum temperature between 100°F and to a 30.1% reduction in accident counts on days with maximum temperature between 0 and 10°F, both relative to days with maximum temperature between 60 and 70°F. Since we are identifying these impacts from presumably exogenous local variation in weather, which we isolate with a collection of seasonal, temporal, and geographic fixed effects, we can plausibly interpret these coefficients as the causal impact of temperature on accident incidence. Thus, these results suggest that we might expect to see more accidents in temperature-sensitive industries as extremely hot days become more frequent and cold days become less frequent under climate change.

Why might we see that the incidence of accidents rises at high temperatures but falls at low temperatures? Recall that the impact of temperature on accident rate is a combination of the impact of temperature on the risk of accidents for a given volume of work and the impact of temperature on the volume and type of work being performed. See Section 1.4 for a discussion of these behavioral changes. In Section 1.3.1, I outline a series of physiological mechanisms by which the risk of accidents for a given type and duration of work would rise both at very high and very low temperatures. For example, research in physiology has found that performance of attention- and perception-based, mathematical, and reaction time-based tasks are impaired both at high and low temperature extremes. However, we may have that high temperatures increase accident risk by more than do low temperatures. If so, the same level of behavioral adaptation might reduce accident rate below baseline at low temperatures while failing to adapt away elevated accident rates in hot conditions.

On the other hand, perhaps behavioral adaptation is more effective or possible at low temperatures than at high temperatures. Then, behavioral adaptations might be sufficient to actually reduce the risk of accidents below that at more temperate conditions in very cold weather, while failing to adapt away elevated accident risk in hot weather. For example, workers can add more layers than they can remove, and while removing layers might mean removing protective gear, additional clothing could protect workers from injury. Next, industries typically operating in winter might allow for postponement of work to better weather conditions, while those industries operating in the summer might be less flexible; then, we could expect to see that cold temperatures have smaller impacts on accident rate than do hot temperatures, or that accident incidence even falls on particularly cold days. Previous studies have found similar decreases in accident incidence at temperature extremes, which they attribute to a reduction in work volume or protective measures taken under those conditions (Morabito et al., 2006, Xiang et al., 2014c).

5.2.3 Selected Initial Robustness Checks

As I describe more fully in Section 4, my primary identification strategy relies on the assumption that my suite of temporal and geographic fixed effects allow me to identify the impacts of exogenous short-term, localized variation in weather. Then, my regressions produce unbiased causal estimates of the impacts of temperature on accident incidence. As I have mentioned previously, the most likely sources of bias in my estimates are within-county temporal dependence and across-county cross-sectional dependence. I explore the role of temporal dependence more fully in Section 5.3, and I will discuss cross-sectional dependence briefly here. It is likely that weather conditions are correlated across nearby counties, producing cross-sectional dependence that could contaminate my standard errors. Then, nearby counties are likely also subject to shared forces impacting accident incidence, so this cross-sectional dependence could induce omitted variable bias.

To evaluate the role of cross-sectional dependence in my results, I estimate my primary model at higher levels of geographic aggregation than the county level. As we move to higher levels of geographic aggregation, weather and accident incidence should correlate less strongly across panel members. In the most extreme case, we could aggregate data to the national level and perform a simple time-series analysis; this regression should be entirely free of the impacts of cross-sectional dependence. Thus, I estimate the impacts of temperature on accident rate first at the state level and then in simple time-series analysis at the national level. I present the results of these regressions in Tables 14 and 15 in the Appendix. In my state-level analysis, I include state, region-by-month, year, and day-of-week fixed effects, and I estimate the impact of state-wide average maximum temperature on state total accident incidence. My results here are very similar to my primary county-level results, with significantly higher accident rate at high temperature extremes and significantly lower accident rate at low temperature extremes. In my national time-series analysis, I include month, year, and day-of-week fixed effects, and I estimate the impact of nationwide average maximum temperature on total national accident rate; I perform this analysis both with OLS and Poisson regression. In both analyses, very low temperatures significantly decrease accident incidence and very high temperatures significantly increase accident incidence. Together, these preliminary analyses suggest that cross-sectional dependence does not drive my county-level estimates for the impacts of temperature on accident rate.

Now, recall that almost 40% of accidents in my sample occurred in California. To ensure that my results are not driven just by the relationship between temperature and accidents in California, I estimate my primary model excluding data from California. The results of this regression are broadly similar to those of my primary analysis including California. Again, we find that accident rate falls at low temperatures and increases at high temperatures, with coefficients of similar magnitude. I present these results in the Appendix.

6 Extensions of Primary Poisson Analysis

In this section, I present the results of additional Poisson analysis as I investigate various robustness checks and alternative specifications.

6.1 Impacts of 5-Degree Temperature Bins

Throughout this primary analysis, I have used a series of twelve 10-degree temperature bins to estimate the impact of temperature on accident incidence. This specification is useful in that it imposes minimal functional form on the relationship between temperature and accidents; that is, we do not assume that accident count moves linearly or quadratically with temperature, for example. However, this structure does require the impact of temperature on accidents to be constant within these ten-degree intervals.

To explore whether this constraint obscures additional nonlinearities in the relationship between temperature and accident incidence, I modify my primary Poisson model by replacing these twelve 10-degree temperature bins with twenty-three 5-degree temperature bins ranging from daily maximum temperature below -5°F to daily maximum temperature over 105°F. Then, my regression specification is given by

$$E[\text{accidents}_{it}|x_{it}] = \exp(\sum_{j=1}^{23} \beta_j \text{ tmax}_{itj} + \sum_{k=1}^{13} \lambda_k \text{ prcp}_{itk} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it})$$

where the letter i denotes county, t denotes calendar day, y denotes year, r denotes U.S. census region, m denotes month, d denotes day of week, and j indexes my 23 temperature bins. Thus, I am including county, year, region-by-month, and day-of-week fixed effects. Just as I exclude the temperature bin for maximum temperature between 60 and 70°F in my primary analysis, here I exclude the dummy variable for maximum temperature between 60 and 65°F, which physiology research suggests may be close to ideal working conditions.

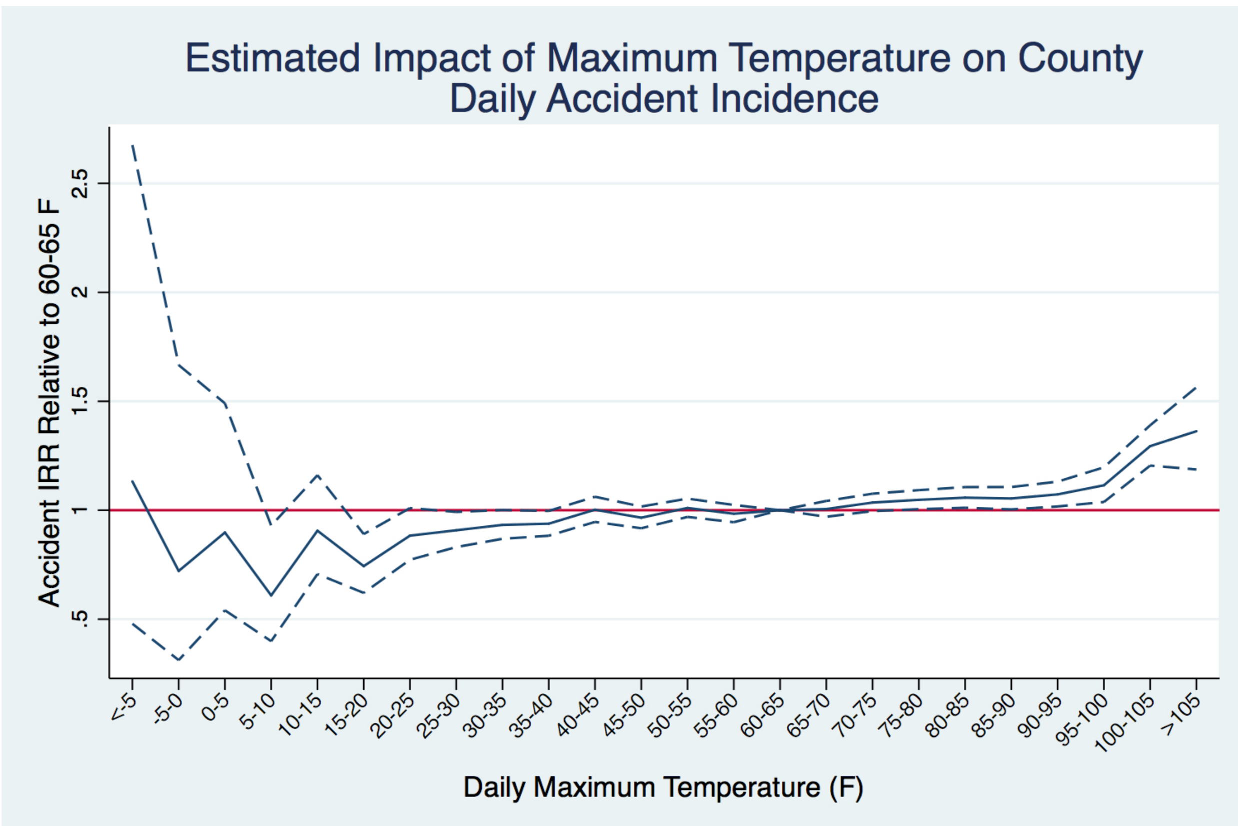

Figure 18: Poisson estimates of the impact of daily maximum temperature on the daily rate of accidents in temperature-sensitive industries. Coefficients give incident rate ratios relative to maximum temperature between 60 and 65°F. N=19,917,285. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 5°F intervals, my regression assumes that temperature impacts are constant within 5°ranges. Regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects.

Figure 18 presents the results of this regression. We see that while many coefficients on high temperature bins are significant, I am unable to identify significant impacts of a number of low temperature bins on accident incidence, likely because those days are less frequent in my sample. In general, I find comparable estimates for the impacts of temperature bins between my primary analysis with 10-degree bins and this analysis with 5-degree bins. These similarities extend throughout most of the temperature distribution. For example, I find here that maximum temperatures between 15 and 2-°F reduce accident rate by about 25.7% relative to maximum temperatures between 60 and 70°F, while I find in my primary analysis that maximum temperatures between 10 and 20°F reduce accident rate by about 21%. I find the same sorts of similarities at high temperatures. For example, I find here that maximum temperature between 80 and 85°F increases accident rate by about 5.8% and maximum temperature between 85 and 90°F increases accident rate by about 5.4%, both relative to maximum temperatures between 60 and 70°F. In my primary analysis, I find that maximum temperatures between 80 and 90°F increase accident rate by about 5.2%.

My primary regression specification constrains the impact of temperature to be constant within 10°F bins; do these results with 5°F temperature bins suggest that this assumption is reasonable? For most of the temperature distribution, the impacts of temperature appear to be largely consistent within 10°F bins. Among the few significant coefficients at low temperatures, the bin for maximum temperatures between 30 and 35°F has an IRR of 0.933 and the bin for maximum temperature between 35 and 40°F has an IRR of 0.938. At high temperatures, incident rate ratios are are largely consistent between 70 and 80°F and between 80 and 90°F: maximum temperatures between 70 and 75°F have 3.5% more accidents while maximum temperatures between 75 and 80°F have 4.8% more accidents, and maximum temperatures between 80 and 85°F have 5.8% more accidents while maximum temperatures between 85 and 90°F have 5.4% more accidents, all relative to maximum temperatures between 60 and 70°F.

However, the assumption that the impacts of temperature are constant within 10°F bins appears to be less appropriate at very high temperatures. In particular, my analysis with 5°F temperature bins suggests that while days with maximum temperature between 90 and 95°F have 7.3% more accidents, days with maximum temperature between 95 and 100°F have 11.4% more accidents, both relative to days with maximum temperature between 60 and 70°F. Then, while days with maximum temperature between 100 and 105°F have 29.4% more accidents, days with maximum temperature over 105°F have 36.3% more accidents, both relative to days with maximum temperature between 60 and 70°F. This divergence is obscured in my primary analysis, where I constrain the impact of temperature to be constant between 90 and 100°F and over 100°F. Thus, these results suggest that the relationship between temperature and accidents is characterized by more significant nonlinearities than my primary analysis suggests. While my regressions here reveal these nonlinearities only at high temperature extremes, they may be equally prevalent at low temperature extremes.

6.2 Dynamics of the Temperature-Accident Relationship

In general, my estimates for the impact of daily maximum temperature on accident incidence are largely free from concerns of omitted variable bias. Using my suite of temporal, seasonal, and geographic fixed effects, I am able to isolate the impacts of short-term localized variation in weather. It is plausible that this variation in weather is exogenous with respect to unobserved determinants of accident incidence. If this assumption holds, my estimation of my primary regression model will produced unbiased estimates of the impact of temperature on accident incidence. Again, some of the most likely contradictions to this assumption of unbiased estimation are within-county temporal dependency and cross-sectional dependence.

In this section, I explore the potential for omitted variable bias from temporal dependence. Temperature today likely correlates with temperature on recent days, which may also impact the incidence of accidents today. For example, we might imagine that if today is particularly hot and an employer notices that her workers seem particularly worn out in the heat, she might plan to perform less work tomorrow. Then, we would expect to see fewer accidents tomorrow; to the extent that today and tomorrow's temperatures are correlated, this mechanism could create omitted variable bias. Similarly, say that an accident occurs among workers at a construction company today, so the foreman and workers decide to be particularly cautious tomorrow, reducing the risk of accidents tomorrow. To the extent that today and yesterday's temperatures correlate, this response could create omitted variable bias. We could think of any number of similar mechanisms by which the correlation between temperature on nearby days could produce omitted variable bias in our estimates for the impact of temperature on accident incidence.

6.2.1 Impact of Heatwaves on Accident Incidence

I first investigate the dynamics of the relationship between temperature and the incidence of accidents by assessing the impact of heatwaves on accident risk. To date, some research on temperature and accident incidence has suggested that prolonged hot periods may have especially severe impacts on worker health. In particular, Xiang et al. (2014b) find that 6.2% more worker injury compensation claims are filed during heatwaves in Adelaide, Australia. We could imagine that several hot days in a row would have a larger effect on accident incidence than would a single hot day. If having a very hot day today is correlated with having a hot day yesterday, the intensification of this "heatwave" effect might be wrapped up in our estimates for the contemporaneous impact of extremely high temperatures on accident incidence. To begin to disentangle this impact, I create a dummy variable indicating whether a day follows a heatwave, independent of that day's temperature. Then, I add this indicator variable to my primary regression, estimating the following Poisson model:

$$E[\text{accidents}_{it}|x_{it}] = \exp(ln(z_{it}) + \rho\ \text{heatwave}_{it} + \sum_{j=1}^{12} \beta_j\ \text{tmax}_{i,t,j} + \sum_{k=1}^{13} \lambda_k\ \text{prcp}_{i,t,k} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d)$$

In this model, I assume that whether or not today follows a string of hot days has some fixed impact on today's accident incidence that is independent of today's temperature. Here, coefficients on temperature bins will give the impact of a day with maximum temperature in that range relative to a day with maximum temperature between 60 and 70°F, holding fixed whether that day follows a heatwave. If heatwaves have a positive impact on today's accident risk and correlate with high temperatures today, we might expect the inclusion of this heatwave dummy variable to reduce our estimate for the contemporaneous impact of high temperatures on accident incidence.

I perform this analysis using several alternative definitions of heatwaves: three days with maximum temperature over 85°F, three previous days with maximum temperature over 90°F, four previous days with maximum temperature over 90°F, and three previous days with maximum temperature over 95°F. Column 2 of Table 3 presents the results of a regression in which I define a day as following a heatwave if the preceding three days had maximum temperature over 85°F and column 3 of Table 3 presents results for which I define a day as following a heatwave if the preceding

Table 3: Heatwaves and Accident Incidence, Poisson

	(1)	(2)	(3)
	Primary	Lag3>85	Lag3>95
VARIABLES	Temperature-Sensitive Accidents		
Max Temp 70-80 F	1.038**	1.037**	1.037**
	(0.0153)	(0.0150)	(0.0151)
Max Temp 80-90 F	1.052***	1.050***	1.056***
	(0.0194)	(0.0199)	(0.0195)
Max Temp 90-100 F	1.082***	1.062**	1.077***
	(0.0249	(0.0269)	(0.0241)
Max Temp >100 F	1.300***	1.280***	1.262***
	(0.0392)	(0.0430)	(0.0442)
3 Preceding Days		1.034**	
Max Temp > 85 F		(0.0175)	
3 Preceding Days			1.063*
Max Temp > 95 F			(0.0355)

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

N=20,955,316 in all regressions. All regressions include a full set of temperature bins, precipitation bins, and county, year, region-by-month, and day-of-week fixed effects. All estimates generated with Poisson regression.

three days had maximum temperature over 95°F. In each regression, we see that the coefficients on the heatwave dummy are positive and significant. In particular, I estimate that a day following three consecutive days with maximum temperatures over 85°F has 3.4% more accidents than a day not following a heatwave of that type, and a day following three consecutive days with maximum temperatures over 95°F has 6.3% more accidents than a day not following a heatwave of that type.

The coefficients on my twelve temperature bins are largely robust to the addition of this heatwave indicator variable in each regression. As expected, the coefficients on low and mid-range temperature bins change only very slightly with these heatwave variables, while the coefficients on very high temperature bins change somewhat more. I only present coefficients on temperature level bins over 70°F in Table 16, but a full regression table is available in the Appendix. While most temperature level bins change little, the coefficients on the very high temperature bins fall somewhat with the addition of a heatwave dummy variable. For example, our estimate for the impact of maximum temperature over 100°F on accidents falls from a 30.0% increase in accidents to a 26.2% increase in accidents when we add a variable indicating whether or not the three previous days had maximum temperature over 95°F. This shift makes sense, since heatwaves of this form increase accident incidence and are presumably positively correlated with hot temperatures today. I also performed this analysis using a definition of heatwaves as three consecutive days with maximum temperature over 90°F and as four consecutive days with maximum temperature over 90°F.The coefficients on these heatwave indicator variables were close in magnitude to those on the heatwave variables in the regressions in columns 2 and 3 of Table 16, but they were not statistically significant. I do not present the results of these regressions.

Overall, then, it appears that accident incidence increases significantly following heatwaves, or strings of hot days. This mechanism may be physiological; a string of hot days could make workers tired or dehydrated, and thus more prone to accidents. However, it may also work through behavioral mechanisms. For example, if workers or employers decide to postpone work during the heatwave, they may reallocate that work to today, increasing today's incidence of workplace accidents. Unfortunately, my analysis here does not permit me to distinguish between these two mechanisms.

My analysis of the impact of heatwaves on accident incidence provides a useful first look at some of the dynamics of the relationship between temperature and accident incidence; in particular, we see that strings of hot days increase the incidence of accidents on the next day, in addition to the contemporaneous effect of today's temperature on today's incidence of accidents. However, this analysis is only a minimal look at the dynamics of temperature and accidents. While we control for a certain sequence of temperatures that we might expect to significantly impact accident risk, namely a certain number of consecutive hot days, we leave all other aspects of temperature dynamics unexplored. For example, this regression does not parse out the impact on accidents today of having a particularly cold day yesterday or of having maximum temperature yesterday between 70 and 80°F. These temperature patterns and their impact on accident incidence might create omitted variable bias that is not addressed in my simple analysis of heatwaves.

6.2.2 Impact of Lagged Daily Temperature

In the following regressions, I seek to more systematically address this potential source of bias in my estimates of the impacts of temperature on accident incidence by controlling for lagged temperature. In particular, I define the same set of twelve 10-degree temperature bins for 1-day lagged, 2-day lagged, and 3-day lagged temperatures that I use for today's temperature throughout my primary analysis. Then, I perform my standard Poisson analysis first with the addition of these 1-day lagged temperature bins, then with both 1- and 2-day lagged temperature bins, and finally, with the addition of 1-, 2-, and 3-day lagged temperature bins.

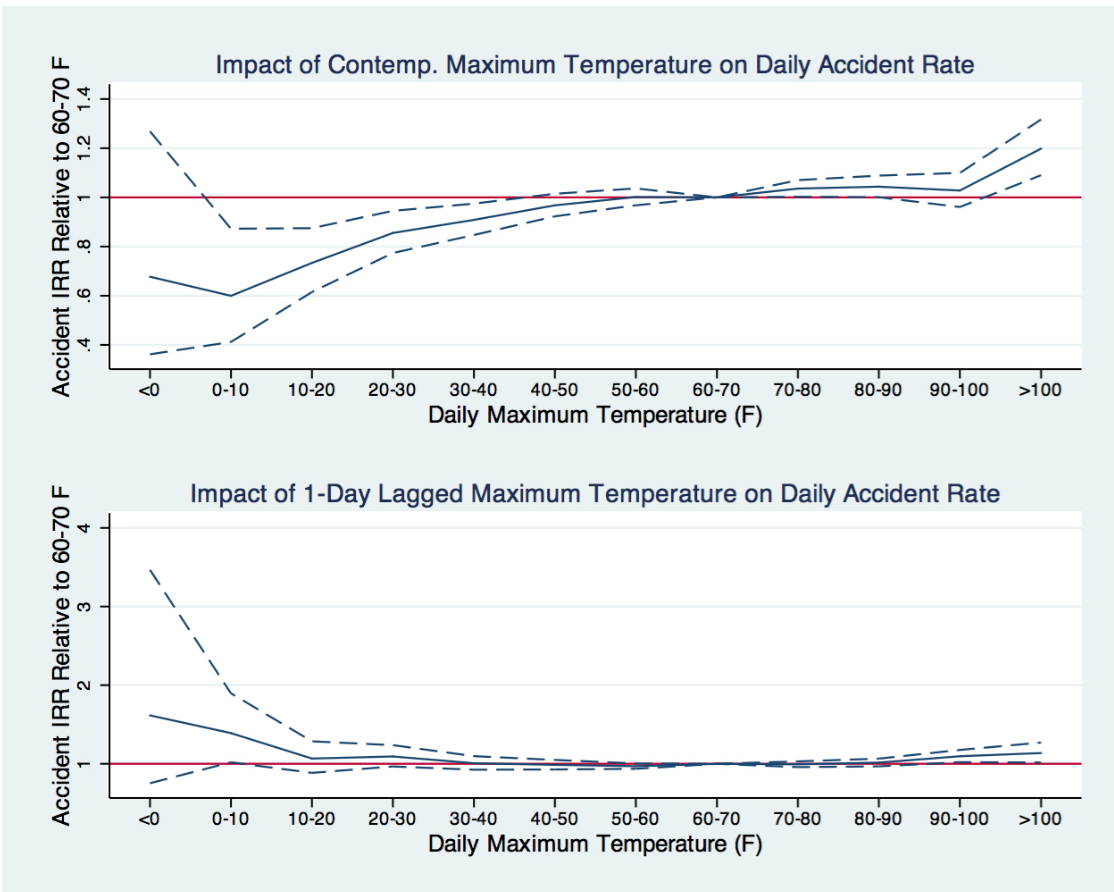

Figure 19: Poisson estimates of the impact of daily maximum temperature on the daily rate of accidents in temperature-sensitive industries in a regression with contemporaneous, 1-day lagged, and 2-day lagged temperature bins. Coefficients give incident rate ratios relative to maximum temperature between 60 and 70°F. N=20,955,316. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Regression includes contemporaneous, 1-day lagged, and 2-day lagged precipitation bins and county, year, and region-by-month fixed effects.

I present some of the results of these regressions in Figures 19 and 20, and present all results in tabular form in the Appendix. Figure 19 gives the results of my primary Poisson model with the addition of twelve 10-degree for both 1- and 2-day lagged temperatures.

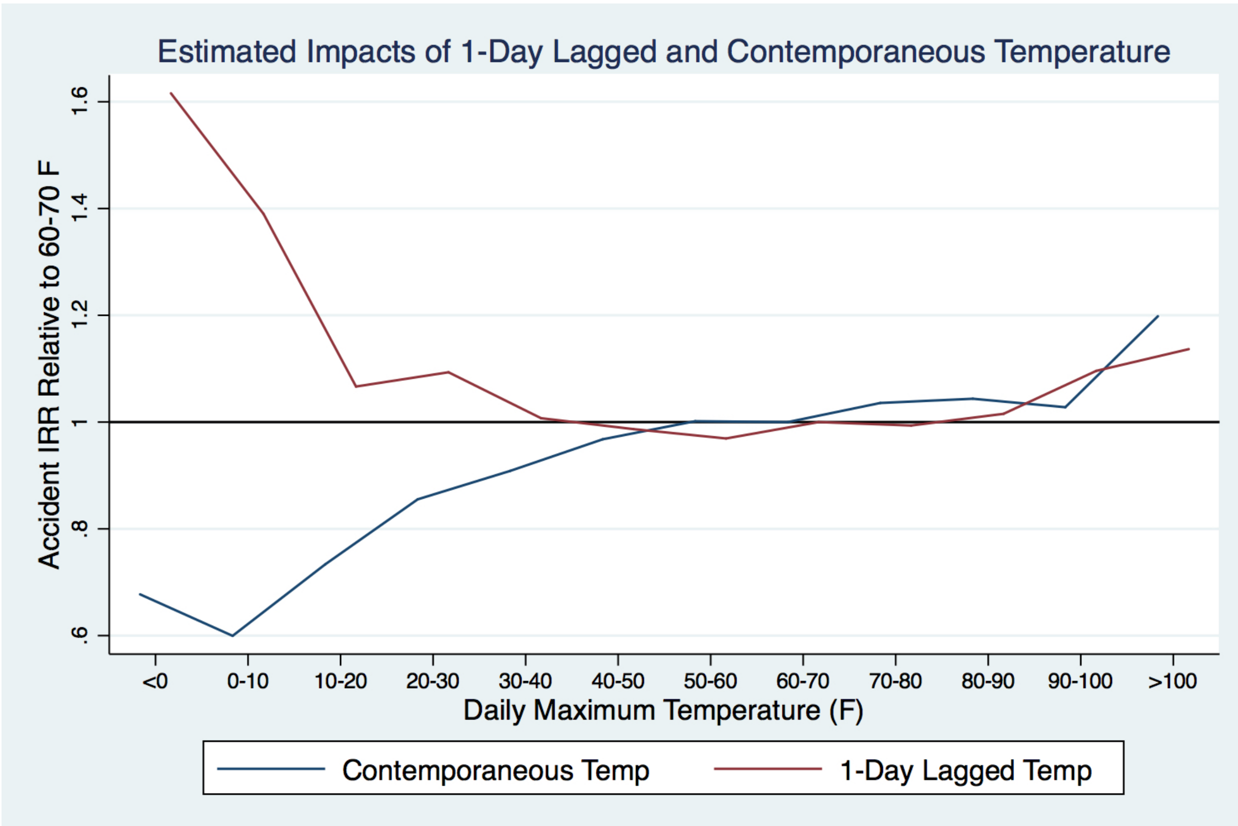

Figure 20: This figure pulls out the coefficient estimates from Figure 19 for easier comparison. 95% confidence intervals are given in Figure 19. These curves trace out the impacts of contemporaneous and 1-day lagged maximum temperature on the rate of accidents in temperature-sensitive industries. These coefficients are pulled from a regression with contemporaneous, 1-day, and 2-day lagged temperature coefficients, as well as precipitation bins and county, year, region-by-month, and day-of-week fixed effects. N=20,955,316.

Based on these results, it appears that 1-day lagged temperature has some impact on today's accident incidence, particularly at temperature extremes. In particular, we see that maximum temperature between 0 and 10°F yesterday is associated with about 33.3% more accidents today, that maximum temperature between 20 and 30°F yesterday is associated with about 9.6% more accidents today, and that maximum temperature over 100°F yesterday is associated with about 13.7% more accidents today, after controlling for today's temperature. While other 1-day lagged temperature bins at temperature extremes are not significantly significant, their point estimates trace out an interesting pattern. In particular, it appears that the impact of yesterday's temperature on today's accident incidence increases in magnitude as yesterday's temperature becomes more extreme. While 1-day lagged temperature appears to have some significant impact on today's accident count, no 2-day lagged temperature bins are significantly associated with today's accident rate. Similarly, when I tried adding in 1-, 2-, and 3-day lagged temperature bins to my primary regression, I found that 3-day lagged temperatures were not significantly associated with today's accident count. Thus, these figures and tables in the Appendix do not present either the coefficients on 2-day lagged temperature bins or any coefficients from the regression including 1-, 2-, and 3-day lagged temperature bins. Comparing columns 1 and 2 of Table 17 in the Appendix, we see that the coefficients on 1-day lagged temperature bins change only slightly when we include controls for 2-day lagged temperature.

In these regressions, we see that my estimates for the contemporaneous impact of temperature on accident incidence are largely robust to the addition of these lagged temperature bins. However, the magnitudes of these coefficients do change somewhat once we control for temperature on previous days. For example, while I estimate in my primary analysis that days with maximum temperature over 100°F have 30.0% more accidents than do days with maximum temperature between 60 and 70°F, my estimate for this impact drops to about 20% once we control for 1- and 2-day temperatures lags. In particular, the coefficients on all contemporaneous temperature bins decrease somewhat in magnitude when we control for 1- and 2-day lagged temperatures. Since today's temperature correlates positively with yesterday's temperatures, we see that high temperature bins today have positive correlation with 1-day lagged high temperature bins and that low temperature bins today have positive correlation with 1-day lagged low temperature bins. Then, since both high and low extreme temperatures yesterday are associated with additional accidents today, our estimates for the contemporaneous impact of temperature appear to be biased upwards somewhat in the absence of 1-day lagged temperature controls.

We would expect that the impacts of yesterday's temperature on the incidence of accidents today would proceed both through physiological and behavioral channels. For example, exposure to hard work in extreme heat or cold yesterday might carry over into physical fatigue today, increasing workers' risk of accidents today. On the other hand, we might expect that if employers postpone work due to extreme temperature yesterday, the volume of work today might be higher, increasing the volume of accidents today. It is interesting to note that while low temperatures today decrease the incidence of accidents today, very cold temperatures yesterday increase the incidence of accidents today. These results would be consistent with work-postponement in response to cold weather; if employers put off work in response to very cold weather today, we would expect to see fewer accidents today and additional accidents today. However, it is not possible to pinpoint which of these mechanisms are at work here in the absence of other data sources.

It is important to note that these results do not entirely remove the potential for omitted variable bias due to interactions with accidents and temperatures on previous days. In particular, it is likely that the incidence of accidents yesterday would impact the risk of accidents today. If an accident occurred among workers at a construction company yesterday, it is likely that workers and supervisors will be more cautious today, reducing today's risk of accident incidence. Then, if the incidence of accidents yesterday correlates both with yesterday and today's temperature and the incidence of accidents today, this mechanism might introduce omitted variable bias to our estimates of the contemporaneous impacts of temperature on accident incidence. However, this omitted variable bias will only be a problem to the extent that today and yesterday's accident incidence is correlated. Here, accidents are sufficiently rare and random that their incidence is not highly correlated from day-to-day within a county. Indeed, the correlation coefficient between today and yesterday's accident incidence is only 0.107 in my data. Thus, I do not expect this source of potential omitted variable bias to pose a significant problem in my estimates of the contemporaneous impacts of temperature on accident incidence.

6.3 Alternative Temperature Shock Specifications

So far, my analysis has assumed that a day with maximum temperature in a certain range has a particular impact on the risk of accidents that is independent of the temperature on preceding days or the climate of the geographic area. In other words, I assume that temperature levels have uniform impacts on accident incidence across time and space. In this section, I explore the impacts on accident incidence of alternative temperature specifications.

6.3.1 Temperature Shocks vs. Levels

While I have assumed that temperature in certain ranges has a fixed impact on accident incidence, some research in physiology suggests that, given sufficient time, workers could adapt to labor in almost all weather conditions. I describe this possibility in further detail in Section 1.3.2. If workers can acclimate to extended periods of heat and cold, we would expect that a day with maximum temperature over 100°F after a period of days around 70°F would have a significantly larger impact on accident incidence than a day with temperatures over 100°F following a full week of equally hot days. In that case, it may be that temperature shocks or variability, rather than absolute temperature levels, truly impact the risk of accidents in outdoor industries.

To address this possibility, I perform a series of variations on my primary analysis that incorporate variables indicating both temperature levels and temperature shocks, which I take to be some measure of the difference between a day's maximum temperature and recent temperature conditions. I define temperature shock to be the difference between today's maximum temperature and the average maximum temperature over the preceding seven days. I then split the distribution for this measure of temperature shock into seven bins: a temperature shock below -25°F, meaning that today is more than 25°F cooler than the average maximum temperature over the preceding week, temperature shock between -25 and -15 °F, temperature shock between -15 and -5°F, temperature shock between -5 and 5°F, temperature shock between 5 and 15°F, temperature shock between 15 and 25°F, and temperature shock greater than 25°F, meaning that today's maximum temperature is at least 25°F warmer than the average maximum temperature over the preceding week. Then, I estimate the following Poisson regression model:

$$E[\text{accidents}_{it}|x_{it}] = \exp\Big(\sum_{l=1}^{5} \phi_l \ (\text{tmax}_{\text{it}} - \text{avg\_week\_tmax}_{it})_l + \sum_{j=1}^{12} \beta_j \ \text{tmax}_{itj} + \sum_{k=1}^{13} \lambda_k$$

$$\text{prcp}_{itk} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it}\Big).$$

I omit the temperature shock variable for a temperature shock between -5 and 5°F, so each coefficient on one of the other temperature shocks bins can be interpreted as giving the incident rate ratio (IRR) for a day in that bin relative to a day with maximum temperature within 5°F of the preceding week's average maximum temperature, holding temperature level constant. I present the coefficients on temperature shock variables from this regression in the first column of Table 4 and in the first graph in Figure 6.3.1. Note that I only report the coefficients on temperature shock bins; these regressions include a full set of temperature level bins and fixed effects in addition to these temperature shock variables. The addition of temperature shock variables leaves the coefficients on temperature level bins largely unchanged in magnitude and significance, so I do not present them here. See a full table of regression coefficients in the Appendix.

Table 4: Temperature Shocks vs. Levels

VARIABLES	(1) Poisson TS Accident IRR	(2) OLS TS Accident Count
Diff 7-day avg <-25	1.095	-0.000126
	(0.109)	(0.000104)
-25≤ Diff 7-day avg <-15	0.986	-9.43e-05*
	(0.0307)	(5.03e-05)
-15≤ Diff 7-day avg <-5	1.005	-4.81e-05*
	((0.0111)	(3.47e-05)
5≤ Diff 7-day avg <15	1.028**	0.000117***
	(0.0123)	(3.40e-05)
15≤ Diff 7-day avg <25	1.035	0.000151***
	(0.0337)	(5.42e-05)
25≤ Diff 7-day avg	0.967	0.000119***
	(0.112)	(0.000138)
Observations	19,917,285	23,614,877

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Column 1 gives the estimated impacts of daily temperature shock on the daily rate of accidents in temperature-sensitive industries, estimated with Poisson regression. Column 2 gives the estimated impacts of daily temperature shock on the daily county-level accident count in temperature-sensitive industries, estimated via OLS regression. Both sets of coefficients give impacts on accident incidence relative to a day with maximum temperature between 60 and 70°F. Both regressions include a full set of temperature level bins, precipitation bins, and county, year, region-by-month, and day-of-week fixed effects. I define temperature shock to be the difference between maximum temperature and average maximum temperature over the preceding week.

Temperature level bins are largely robust to the addition of temperature shock variables, so

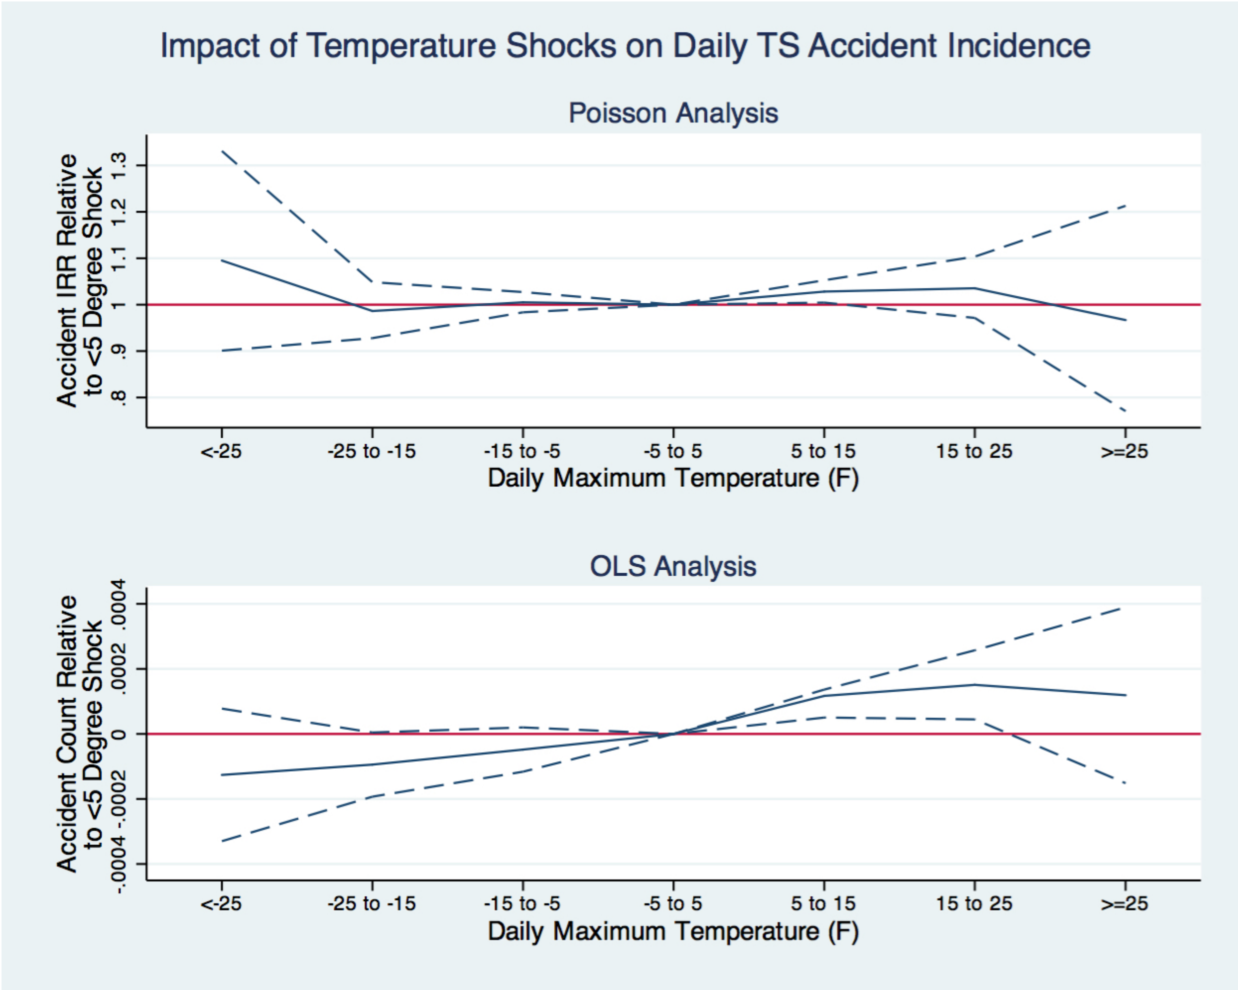

Figure 21: Estimates of the impacts of temperature shocks on daily accident incidence in temperature-sensitive industries at the county level. These figures graphically present coefficients presented in Table 4. The upper graph gives impacts of daily temperature shocks on accident rate, estimated via Poisson regression, while the lower graph gives impacts of daily temperature shocks on daily accident count, estimated via OLS. Both are relative to a day with maximum temperature between 60 and 70°F. N=19,917,285 in the upper graph and N=23,614,877 in the lower graph. Dashed lines give the 95% confidence interval. Both regressions include a full set of temperature and precipitation bins and county, year, region-by-month, and day-of-week fixed effects.

these results suggest that temperature level remains a primary determinant of accident risk in temperature-sensitive industries, independent of temperature variability. However, these regressions provide some suggestive evidence that temperature variability may also impact accident risk. While most coefficients on temperature shock bins are not statistically significant, the coefficients on mid-size positive temperature shocks is significant and positive. In particular, we see that a difference between maximum temperature and average maximum temperature over the past week of between 5 and 15°F is associated with a 2.8% increase in accident incidence, both holding temperature level constant. While this coefficient is the only temperature shock bin coefficient with statistically significant impacts on accident incidence, it is possible that days with larger positive temperature shocks also increase accident risk but are rare enough that I lack the power to estimate their impact.

An OLS regression including both temperature shock and temperature level variables adds to the suggestive evidence that temperature shocks, in addition to temperature levels, may impact accident incidence. In particular, I add to my primary OLS regression models the same temperature shock bin variables that I outline above. I present the coefficients on temperature shock variables from this regression in the second column of Table 4 and in the second graph in Figure 21. Note again that I only report the coefficients on temperature shock bins; these regressions include a full set of temperature level bins and fixed effects in addition to these temperature shock variables. In these OLS regressions, all positive temperature shock bins are associated with statistically significant increases in accident incidence. While we should not place too much weight on these OLS results, they add to the suggestive evidence that positive temperature shocks increase accident incidence.

6.4 County-Specific Temperature Percentiles

Not only does my primary analysis of temperature and accidents implicitly assume that a given temperature level has a fixed impact on accident incidence, regardless of temperature variability, but it also assumes that a given temperature level has a fixed impact on accident incidence across regions with different underlying climates. This assumption may be unrealistic; in particular, we might imagine that a day with maximum temperature over 100°F would have a different impact on accident incidence in counties that regularly experience temperatures over 100°F than in counties for which those temperatures are rare. Then, it may be that days in a certain range of a county's temperature distribution have a particular impact on accident incidence, while the impact of a certain temperature level varies across counties. For example, we could imagine that a day with temperature falling in the bottom 5 percentiles of a county's distribution for maximum temperature would have a constant impact on accident incidence across counties.

Table 5: Average County Temp Percentiles

Percentile	Temperature (°F)
1	24.630
5	34.043
10	39.412
15	43.334
25	49.899
35	56.477
45	63.221
55	69.590
65	75.494
75	80.490
85	84.847
90	87.150
95	90.161
99	95.068

Average values for percentiles of maximum temperature distribution across counties in my sample. Calculated for 1990-2010 across 2,840 counties.

To explore this possibility, I estimate an alternative form of my primary Poisson analysis in which I redefine my temperature variables to identify particular segments of counties' distributions for maximum temperature. For each county, I calculate temperature values corresponding to the 1^{st}, 5^{th}, 10^{th}, 15^{th}, 25^{th}, 35^{th}, 45^{th}, . . . , 75^{th}, 85^{th}, 90^{th}, 95^{th}, and 99^{th} percentile of that county's distribution for maximum temperature. I present the average values for these county-specific percentiles in Table 5. Then, I define a series of fifteen dummy variables indicating whether a day falls in temperature ranges delineated by those percentile temperatures. For instance, if the 35^{th} percentile and 45^{th} percentile for a county's maximum temperature distribution fall at 60 and 65°F, for example, a day with maximum temperature of 63°F would take a value of 1 for the dummy variable indicating maximum temperature between the 35^{th} and 45^{th} percentile and a value of 0 for all other temperature indicator variables. I can express this modified primary regression as

$$E[\text{accidents}_{it}|x_{it}] = \exp(ln(\text{emp}_{it}) + \sum_{j=1}^{15} \beta_j \text{ tpercentile}_{i,t,j} + \sum_{k=1}^{13} \lambda_k \text{ prcp}_{i,t,k} + \alpha_i + \gamma_y + \theta_{r*m} + \nu_d + \epsilon_{it}),$$

where each $tpercentile_j$ variable is a dummy variable corresponding to a certain range of a county's maximum temperature distribution. I omit the temperature percentile bin corresponding to maximum temperature between the 45^{th} and 55^{th} percentile of a county's temperature distribution, so the coefficients on all temperature percentile bins can be interpreted as accident IRRs relative to a day with temperature close to the county's median maximum temperature.

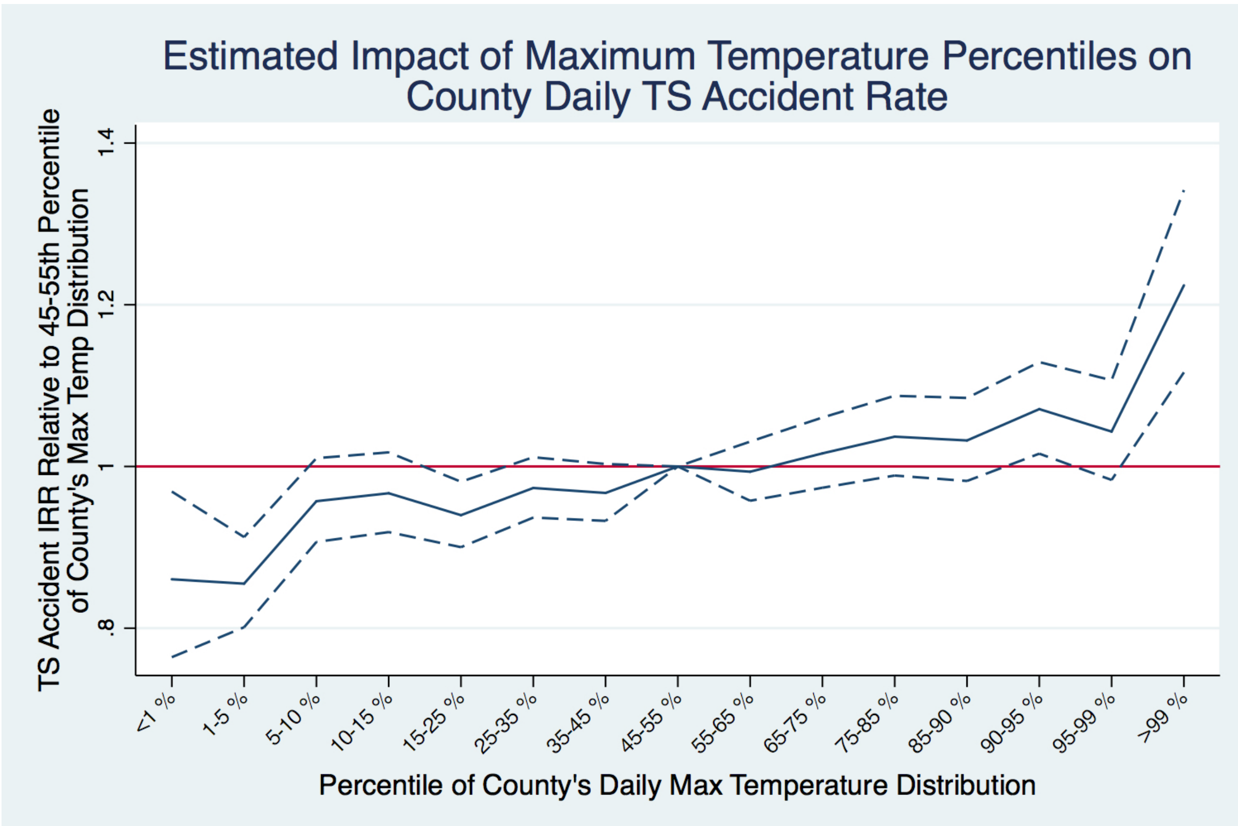

Figure 22: Estimates of the impact of days in each percentile of a county's maximum temperature distribution on daily accident rate in temperature-sensitive industries. Percentile cut-offs are county-specific. Coefficients give incident rate ratios relative to maximum temperature between the 45^{th} and 55^{th} percentiles of a county's maximum temperature distribution. N=19,917,285. Dashed lines give the 95% confidence interval. While this graph interpolates between estimated coefficients, my regression assumes that temperature impacts are constant within temperature percentile bins. Regression includes precipitation bins and county, year, and region-by-month fixed effects.

I present the results of this regression in Figure 22. In these regression results, we see that the coefficients on these county-specific temperature bins trace out a relationship between temperature exposure and accident incidence that is qualitatively similar to that based on absolute temperature bins that do not vary by county. First, we see that low percentile temperatures are associated with lower accident incidence. For example, maximum temperature in the first percentile of a county's maximum temperature distribution is associated with 13.9% fewer accidents than a day with maximum temperature between the 45^{th} and 55^{th} percentiles of this distribution. While some of the coefficients on low temperature bins are not statistically significant in this regression, the point estimates for these coefficients follow a smooth trajectory whereby they approach and then hover around an IRR of 1 as temperature grows less extreme relative, corresponding to no difference in accident incidence relative to a day near a county's median temperature.

As in my primary analysis with county-invariant temperature level bins, relatively high temperatures tend to be associated with increased accident incidence. However, while almost all high temperature bin coefficients are significant in my primary Poisson analysis, only a few high temperature bins are associated with significant changes in accident incidence under this specification with county-specific temperature percentiles. In particular, we see that days with maximum temperature between the 90 and 95^{th} percentiles of a county's maximum temperature distribution have about 7.1% more accidents and days with maximum temperature in the top 99^{th} percentile of a county's maximum temperature distribution have about 22.4% more accidents, both relative to days with maximum temperature between the 45^{th} and 55^{th} percentiles of a county's temperature distribution. While the coefficients on other relatively high temperature bins are not statistically significant, their point estimates generally trace out a pattern in which the accident IRR rises as temperatures grow gradually more extreme relative to the county median maximum temperature.

We see, then, that these temperature bins based on county-specific temperature percentiles are less significant predictors of accident risk than are the absolute temperature bins that I use in my primary analysis, particularly at relatively high temperatures. It is also interesting to note that the coefficients on these county-specific temperature bins tend to be somewhat less extreme in magnitude than those associated with the absolute temperature bins in my primary analysis. This appears to be consistent with a model in which absolute temperature levels have fixed impacts on accident incidence. That is, the absolute temperature values corresponding to these county-specific temperature percentiles tend to be less extreme than the temperatures delineating the bins in my primary models. For example, the average value for the 1st percentile in a county's distribution of maximum temperature is 24.16°F on average, significantly warmer than 0°F, the cut-off for the most extreme low temperature bin from my primary analysis. Similarly, the average value for the 99^{th} percentile in a county's distribution of maximum temperature is 94.95°F, somewhat cooler than 100°F, the cut-off for the most extreme high temperature bin from my primary analysis.

On average, then, these county-specific temperature bins are measuring the impact on accident incidence of more moderate temperatures than those that delineate temperature categories in my primary analysis. Then, if absolute temperature levels do in fact have constant impact on accident incidence, with the impact of temperature on accidents increasing in magnitude as temperature becomes more extreme, we would expect to estimate coefficients on these county-specific temperature bins that are lower in magnitude than those from my primary analysis. Therefore, the results of these regressions are consistent with a world in which accident risk is more closely related to absolute temperature level than to temperature relative to an area's climate.

7 First Looks at Adaptation

I introduce the concept of adaptation in detail in Section 1.3.2. Essentially, adaptation is the process by which agents adjust processes to a new or changing environment. If rising temperatures associated with climate change negatively impact economic processes or outcomes, we would expect agents to gradually make adaptive changes to reduce these impacts. This adaptation may occur at a range of temporal scales. In the short term, workers or employers might reduce the risks associated with high temperatures by postponing work to cooler hours or scheduling additional breaks. In the long term, we might see the development of new worker safety regulations or additional investments in protective technologies. Besides these forms of economic adaptation, where agents make a series of decisions to reduce the impact of temperature on accident incidence, we might also see physical adaptation, where workers' bodies acclimatize to consistently higher temperatures.

Estimating the potential for these various forms of adaptation is critical for fully understanding the relationship between temperature and the incidence of occupational accidents. First, my estimates of the contemporaneous impact of temperature on accident incidence currently incorporate both the impact of temperature on accident risk for a given amount of work and the influence of short-term adaptation; this short-term adaptation would likely be changes in the type or volume of work. To fully understand the impacts of temperature on accidents, we would like to decompose these mechanisms as much as possible. Estimating the role of adaptation becomes particularly crucial as I work to scale my short-run estimates of the impacts of temperature on accident incidence to the changes in accident risk that we would expect to see under climate change. In particular, long-term adaptation like changes in worker safety regulations or norms might make the long-run impacts of temperature lower than the comparable short-run impacts. In this section, I outline some preliminary results that explore some of the adaptive mechanisms that I outlined in Section 1.3.2.

7.1 Physical Acclimatization over Summer Months

As I described in Section 1.3.2, research in physiology suggests that workers can acclimatize to extreme temperature conditions within one to two weeks. Again, this acclimatization advances primarily through increased sweating efficiency, which allows workers to function with a lower core temperature, reduced heart rate, and overall reduced thermoregulatory strain. Thus, physiological acclimatization could substantially reduce the impact of temperature on the incidence of occupational accidents.

I test for the influence of acclimatization in the short-run by comparing the impact of temperature on the incidence of occupational accidents across summer months. We would expect that workers might become better acclimatized to high temperatures as the summer progresses, so that high temperatures would have a larger impact on the incidence of occupational accidents at the beginning of the summer than at the end. I model this analysis after that performed by Graff Zivin and Neidell (2014) in their analysis of the impact of temperature on time spent working in temperature-sensitive industries. Since low temperatures are rare during summer months, I collapse all temperature bins corresponding to maximum temperature below 60°F into a single temperature bin. I then estimate an otherwise-unchanged version of my primary Poisson analysis for June, July, and August. As usual, I omit the temperature bin for maximum temperature between 60 and 70°F, so all coefficients on temperature bins give IRRs relative to temperatures in that range. When performing Poisson regressions, Stata drops panel members with all zero accidents counts, so counties are dropped from the June regression if no accidents have occurred there in June between 1990 and 2010, for example. I standardize my sample across these regressions, keeping only those 614 counties for which an accident has occurred in June, July, and August from 1990 through 2010.

Table 6: Primary analysis across summer months

	(1)	(2)	(3)
	June	July	August
VARIABLES		HS Accident IRR	
Max Temp < 60	1.026	1.770**	0.852
	(0.0978)	(0.474)	(0.159)
Max Temp 70-80 F	1.074	1.028	0.902
fa	(0.0550)	(0.0977)	(0.0801)
Max Temp 80-90 F	1.215***	1.047	0.873
	(0.0782)	(0.104)	(0.0798)
Max Temp 90-100 F	1.255***	1.143	0.892
	(0.105)	(0.127)	(0.0851)
Max Temp >100 F	1.441***	1.522***	0.967
	(0.154)	(0.211)	(0.105)
Observations	385,590	398,350	398,350
Number of counties	614	614	614

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Estimates of the impact of daily maximum temperature on daily accident rate in temperature-sensitive industries across summer months. I modify my primary Poisson analysis by combining all bins for temperatures below 60°F into a single bin. I then estimate this regression separately for June, July, and August. All coefficients give incident rate ratios relative to a day with maximum temperature between 60 and 70°F. All regressions include county, year, region-by-month, and day-of-week fixed effects.

I present the results of these regressions in Table 6.

Here, we see that my estimates for the impacts of temperature on accident incidence change substantially over the course of the summer. The most striking change is that the statistical significance of the relationship between temperature and accident incidence decreases over the course of the summer. While three of the five coefficients on temperature bins are statistically significant at the 1% level in June, only two are significant at least at the 5% level in July, and none are significant even at the 10% level in August. We could imagine that we might see this loss of significance if fewer accidents occurred later in the summer, perhaps because of workers' vacation schedules. Then, we might not have sufficient statistical power to identify the impact of temperature on accidents in August. However, the total volume of accidents reported stays consistent over summer months; about 31.6% of summer-month accidents in my sample occurred in June, about 34.3% occurred in July, and about 34.1% occurred in August. Thus, these results suggest that temperature becomes a progressively less significant determinant of accident incidence as the summer progresses. While this shift could be driven by a number of processes that we do not observe, like changes in the type of work being done in different months, these results would be consistent with the idea that the impact of temperature on accident incidence falls as workers acclimatize to high temperatures over the course of the summer.

Changes in the magnitude of the impacts of temperature on accident incidence between June, July, and August are also weakly suggestive of this acclimatization, though weak statistical significance in July and August prevents robust comparisons. For example, we see that the point estimates for the impact of all temperature bins are larger in magnitude in both June and July than in August, though these coefficients are not statistically significant in August. While most of the coefficients on temperature bins that are statistically significant in June are no longer significant in July, the point estimates of these coefficients tend to be somewhat larger in magnitude in June than in July. For example, while a day with maximum temperatures between 80 and 90°F has about 21.5% more accidents than a day with maximum temperature between 60 and 70°F in June, the point estimate for the same variable in July suggests that the same day has about 4.7% more accidents in July. Similarly, a a day with maximum temperature between 90 and 100°F has about 25.5% more accidents than a day with maximum temperature between 60 and 70°F in June, while the point estimate for the same variable in July suggests that the same day has about 14.3% more accidents in July. My estimates for the impacts of days with maximum temperature over 100°F do not match this pattern, however. While a day with maximum temperature over 100°F has about 44.1% more accidents than a day with maximum temperature between 60 and 70°F in June, the same day has about 52.2% more accidents in July, a much larger impact.

While the impacts of high temperatures on accident incidence do not decrease uniformly in magnitude as the summer progresses, these results suggest that temperature is a substantially stronger predictor of accident incidence early in the summer than later in the summer. Indeed, temperature appears to have no significant impact on the incidence of occupational accidents in August. Altogether, these results are consistent with the hypothesis that physical acclimatization significantly reduces the impact of temperature on the risk of occupational accidents. This potential for physical acclimatization may help to minimize the costs of climate change associated with higher volumes of accidents at high temperatures, since workers could acclimatize to consistently hotter weather.

7.2 Heterogeneity in Impacts Across Historical Climates

In my analysis so far, I have assumed that the impact of temperature on accidents is constant across space and time, after we control for a suite of seasonal, temporal, and geographic fixed effects. However, this assumption is unlikely to hold. We can follow similar reasoning to that which I outline in Section 4.1.1. The inclusion of county and year fixed effects allows us to effectively subtract out county and year averages in temperature and accidents; we then identify the impact of temperature on accidents using variation in temperature and accident incidence about long-run county averages, after controlling for nation-wide year effects, region-specific seasonal cycles, and average accident incidence on each day of the week. This suite of fixed effects allows us to carefully identify the changes in accident incidence that stem from local temperature fluctuations, rather than cross-sectional correlations; even after controlling for these fixed effects, however, the true coefficients on temperature bins may differ across members of our panel. In other words, the extent to which accident incidence increases or decreases in response to temperature likely differs by county. This heterogeneity in the impacts of temperature on accident incidence could stem from many factors, including work safety culture and industrial composition.

While this wide range of possible sources makes it difficult to definitively pin down why the impacts of temperature on accidents differ across space and time, preliminary investigations of the determinants of this heterogeneity may provide a first look at the historical role of adaptation in the impact of temperature on occupational accidents. Dell, Jones, and Olken (2014) outline this form of analysis in their paper, “What Do We Learn From the Weather?: The New Climate-Economy Literature.” They write that we can learn about adaptation by examining how the impacts of temperature on various economic outcomes varies based on areas’ historical climates. Essentially, we would imagine that geographic areas in which certain weather events are common will see lower magnitude impacts from those temperature shocks. For example, Dell, Jones, and Olken cite the example of snow storms that occur in Boston versus those that occur in Washington D.C. Snowstorms are sufficiently common in Boston that we have developed infrastructure and plans to quickly clear snow and restore commercial flows. In contrast, Washington D.C., where snowstorms are less frequent, has not developed the same recovery mechanisms.

In the same way, we might expect the impact of temperature on accidents to be lower in areas that routinely experience extreme temperature conditions. In these areas, workers might have achieved more full physiological acclimatization to extreme temperatures, and norms or regulations might be in place to reduce the impacts of these temperatures on worker safety. By evaluating the extent of this attenuation, we can get a sense for the level of historical adaptation to the impacts of temperature on accident incidence. In turn, these results may provide insight into the extent to which adaptation could reduce the impact of temperature on accident incidence under the long timescales of climate change.

7.2.1 Impacts of Temperature on Accidents by Region

As a first pass at this sort of analysis, I investigate how the impact of temperature on accident incidence varies by region. I classify counties by broad Census Regions, and then perform my primary Poisson analysis separately within each of these regions. I present the results of this analysis in Table 7. Based on this analysis, it is clear that the impact of temperature varies significantly by region. First, more temperature bins are significantly associated with accident

Table 7: Impacts of Temperature on Accident Rate by Region

	(1) Northeast	(2) Midwest	(3) South	(4) West
VARIABLES			HS Accident IRR	
Max Temp <0 F	0.891	0.969	7.85e-09***	4.81e-06***
	(0.911)	(0.300)	(2.37e-09)	(9.68e-07)
Max Temp 0-10 F	0.914	0.714*	1.56e-08***	0.536
	(0.331)	(0.143)	(2.23e-09)	(0.379)
Max Temp 10-20 F	0.708*	0.804**	1.066	1.029
	(0.128)	(0.0877)	(0.299)	(0.207)
Max Temp 20-30 F	0.914	0.963	0.789**	0.823***
	(0.0844)	(0.0705)	(0.0871)	(0.0604)
Max Temp 30-40 F	0.937	0.959	0.899*	0.934
	(0.0607)	(0.0587)	(0.0496)	(0.0392)
Max Temp 40-50 F	0.930	1.024	1.043	0.936*
	(0.0509)	(0.0543)	(0.0384)	(0.0328)
Max Temp 50-60 F	1.009	1.004	1.004	0.975
	(0.0558)	(0.0488)	(0.0331)	(0.0184)
Max Temp 70-80 F	1.131**	1.047	0.984	1.060***
	(0.0588)	(0.0433)	(0.0291)	(0.0210)
Max Temp 80-90 F	1.142**	1.090*	1.010	1.067**
	(0.0754)	(0.0510)	(0.0353)	(0.0291)
Max Temp 90-100 F	1.142	1.149*	1.060	1.082**
	(0.135)	(0.0900)	(0.0466)	(0.0330)
Max Temp >100 F	1.25e-05***	1.529***	1.353***	1.275***
	(4.66e-06)	(0.215)	(0.101)	(0.0439)
Observations	1,599,287	6,097,977	9,567,971	2,652,050
Number of counties	211	813	1,266	357

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Estimated impact of daily maximum temperature on daily accident rate in temperature-sensitive industries. I perform my primary Poisson analysis separately by Census Region. Northeast: CT, ME, MA, NH, RI, VT, NJ, NY, PA. Midwest: IL, IN, MI, OH, WI, IO, KS, MN, MO, NE, ND, SD. South: DE, D.C., FL, GA, MD, NC, SC, VA, WV, AL, KY, MS, TE, AR, LA, OK, TX. West: CO, ID, MT, NV, NM, UT, WY, CA, OR, WA. All regressions include precipitation bins and county, year, month, and day-of-week fixed effects.

incidence in the West than in the other regions. This discrepancy may be driven by the fact that the vast majority of accidents reported to OSHA come from the West, giving us better power to estimate the impacts of temperature in the West, or it could reflect a world in which accidents truly impact accident incidence more in the West than elsewhere in the country.

I also see substantial differences in the magnitudes of coefficients on these temperature bins across geographic regions. First, it appears that the impact of moderately high temperatures is larger in magnitude in the Northeast than in other areas of the country. For example, a day with maximum temperature between 70 and 80°F has about 13.1% more accidents relative to a day with maximum temperature between 60 and 70°F in the Northeast, while it has only 6.0% more accidents in the West. Then, a day with maximum temperature between 80 and 90°F has about 14.2% more accidents than a day with maximum temperature between 60 and 70°F in the Northeast, while a day in that range in the Midwest has only 9.0% more accidents and a day in that range in the West has only 6.7% more accidents.

At extremely high temperatures, though, we see that accident incidence increases less in the Northeast than elsewhere in the country. The impact of a day with maximum temperature over 100°F is quite large in the Midwest, South, and West, with about 52.9% more accidents in the Midwest, about 35.3% more accidents in the South, and about 27.5% more accidents in the West, all relative to a day with maximum temperature between 60 and 70°F. However, maximum temperature over 100°F has a negligible impact on accident incidence in the Northeast, where a day with temperature in that range has only 1.15e-05% of the accidents on a day with maximum temperature between 60 and 70°F.

What can these results tell us about adaptation? Based on the hypothesis that historically hot areas have had the opportunity to adapt to the impacts of temperature on accidents and thus reduce this impact, we would expect to see that temperature has a relatively small impact on accident incidence in the South, which is typically hot, and a relatively large impact on accident incidence in the Northeast, which is typically cooler. We see those patterns to a limited extent; for example, days with moderately high maximum temperature have a relatively large impact on accident incidence in the Northeast. However, we also see that days with maximum temperature

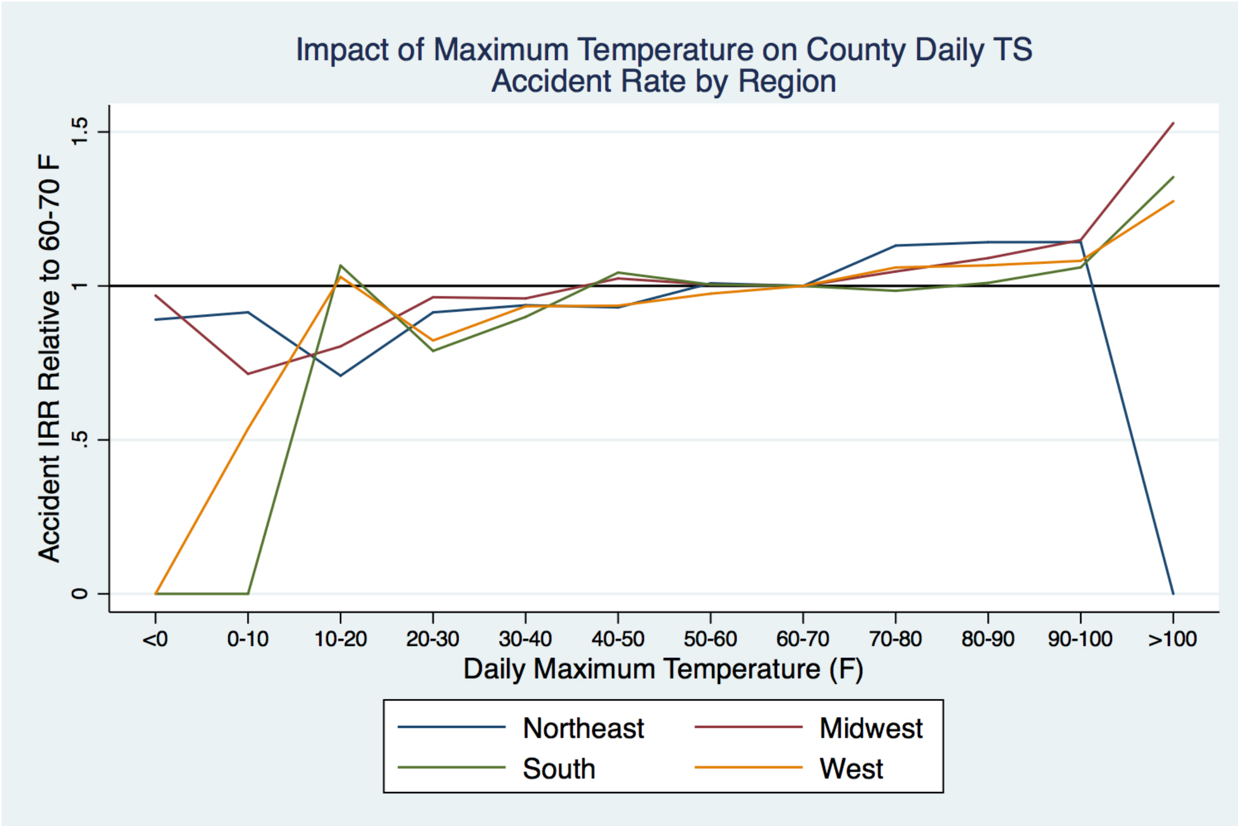

Figure 23: Estimated impact of daily maximum temperature on daily accident rate in temperature-sensitive industries by Census region. This figure plots the estimated coefficients presented in Table 7. Here, I estimate my primary Poisson analysis separately by region. All regressions include precipitation bins, and county, year, month, and day-of-week fixed effects. Northeast: CT, ME, MA, NH, RI, VT, NJ, NY, PA. Midwest: IL, IN, MI, OH, WI, IO, KS, MN, MO, NE, ND, SD. South: DE, D.C., FL, GA, MD, NC, SC, VA, WV, AL, KY, MS, TE, AR, LA, OK, TX. West: CO, ID, MT, NV, NM, UT, WY, CA, OR, WA.

over 100°F have a relatively small impact on accident incidence in the Northeast and a relatively large impact on accident incidence in the South. Thus, these results do not provide significant evidence for the role of adaptation in the impacts of temperature on accident incidence.

However, these broad geographic regions differ in many characteristics other than their climate. Besides varying in their historical weather conditions, these regions differ in industrial composition, cultures, and worker safety laws, for example. Thus, we cannot claim that these differences in the impact of temperature on accident incidence across regions are due to differences in climate, so these regional results allow us to draw only very limited inferences about adaptation.

7.2.2 Impacts of Temperature on Accidents by County Temperature Groups

In the previous section, I use region as a proxy for historical climate to get a first look at historical adaptation to the impacts of extreme temperature. However, historical climates are only loosely unified within broad Census regions. Thus, we could perhaps better look at heterogeneity in the impact of temperature on accident incidence across climates by more rigorously classifying counties by historical weather conditions. I calculate average maximum temperature for each county, and then split the counties in my sample into four groups based on the quartiles of this distribution for average maximum temperature. Then, I run my primary regression separately within these groups of groups, each of which contains about 650 counties. I present the results of these regressions in Table 8.

These regressions provide some limited evidence that historically hot counties have adapted away some of the impacts of temperature on accident incidence. Unfortunately, the coefficients on some high temperature bins are not statistically significant among the 1st and 2nd group of counties, likely because very hot days are not sufficiently common to provide the statistical power to estimate the true relationship between temperature and accidents. I will primarily compare the magnitudes of the impacts of temperature on accident incidence across the remaining statistically significant coefficients.

There is some limited evidence that the magnitudes of the impact of very high temperature on accident incidence declines among historically hotter counties. For example, we see that the impact of days with maximum temperature between 90 and 100 declines somewhat across these county groups, with a 15.1% increase in accidents in the 2nd quartile of counties, an 8.3% increase in accidents in the 3rd quartile, and a 6.2% increase in accidents in the 4th quartile of counties, all relative to a day with maximum temperature between 60 and 70°F. These results suggest that the impact of days with maximum temperatures between 90 and 100°F has attenuated somewhat in historically hot counties, which have had opportunities for long-run adaptation. We see somewhat of the same pattern in the impacts of a day with maximum temperature over 100°F, where days in that range are associated with 57.3% more accidents in the third quartile of the temperature distribution and with only 25.2% more accidents in the highest quartile of the temperature distribution. While

Table 8: Impacts of Temperature by County Temperature Groups

	(1) 1st Quartile	(2) 2nd Quartile	(3) 3rd Quartile	(4) 4th Quartile
VARIABLES		TS Accident IRR		
Max Temp <0 F	0.889	1.640	0	
	(0.260)	(0.956)	(0)	
Max Temp 0-10 F	0.748*	0.841	0***	0
	(0.131)	(0.288)	(0)	(0)
Max Temp 10-20 F	0.829**	0.805	0.923	2.016
	(0.0765)	(0.113)	(0.292)	(1.846)
Max Temp 20-30 F	0.919	0.929	0.877	0.795
	(0.0537)	(0.0598)	(0.122)	(0.314)
Max Temp 30-40 F	0.959	0.932	0.885*	0.895
	(0.0434)	(0.0414)	(0.0642)	(0.112)
Max Temp 40-50 F	0.962	1.018	1.009	0.911*
	(0.0383)	(0.0377)	(0.0358)	(0.0511)
Max Temp 50-60 F	1.000	1.002	1.030	0.961**
	(0.0353)	(0.0344)	(0.0329)	(0.0180)
Max Temp 70-80 F	1.064*	1.036	1.038	1.033
	(0.0389)	(0.0363)	(0.0297)	(0.0251)
Max Temp 80-90 F	1.072*	1.082*	1.062	1.050
	(0.0439)	(0.0477)	(0.0403)	(0.0328)
Max Temp 90-100 F	1.029	1.151**	1.083*	1.062*
	(0.0957)	(0.0778)	(0.0494)	(0.0354)
Max Temp >100 F	1.730	1.215	1.573***	1.252***
	(0.683)	(0.252)	(0.173)	(0.0455)
Observations	5171459	5087197	5199006	5497654
Number of counties	664	649	656	692

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Estimated impact of daily maximum temperature on daily accident rate in temperature-sensitive industries, estimated separately by county temperature groups. Counties are classified into quartiles of the distribution for average daily maximum temperature. Each regression includes county, year, region-by -month, and day-of-week fixed effects, as well as precipitation bins.

the imprecise coefficients on the bin for days with maximum temperature over 100°F for counties in quartiles 1 and 2 prevent us from meaningfully considering whether this pattern continues in other quartiles of the temperature distribution, these results provide some suggestive evidence that the impact of days with maximum temperature over 100°F is substantially lower in hotter counties.

We do not see the same sort of decrease in the impacts of temperature on accident incidence across county quartiles for days with maximum temperature between 80 and 90°F. Here, such days are associated with increases in accident incidence of between 5 and 8% across all county quartiles, though these coefficients are statistically insignificant in the third and fourth quartiles of counties. It is likely that days with temperature in that range are sufficiently common throughout all quartiles of our temperature distribution that hot counties would not have adapted to them significantly more than have cooler counties. Overall, these regressions provide some suggestive evidence that historically hot counties have adapted away some of the impacts of very hot days on accident incidence. Thus, these results provide hope that the impact of temperature on accident incidence might bend down as temperatures rise overall under the long timescales of climate change.

7.3 Heterogeneity in Impacts Over Time

Next, I take another look at the role of long-run adaptation in the relationship between temperature and accident incidence by looking at how the magnitude of this impact has changed over time. If people have historically adapted to the impacts of temperature on accident rate, we would expect the magnitude of this relationship falling over time. I split my sample into three time periods of equal duration: 1990 through 1996, 1997 through 2003, and 2004 through 2010. Then, I perform my primary Poisson analysis separately within each of these time periods. I present the results of these separate regressions graphically in Figure 24. In Figure 25, I combine the estimated coefficients of these regressions on a single graph for better comparison. Finally, I present the results of these regressions in Table 20 in the Appendix.

First, we note that slightly more of the coefficients on temperature bin variables are significant in the two more recent periods than in the first period. This discrepancy is not likely to be driven by differences in the number of accidents in my sample over the three periods; about 30.0% of accidents reported to OSHA in my study period occurred between 1990 and 1996, 35.4% occurred

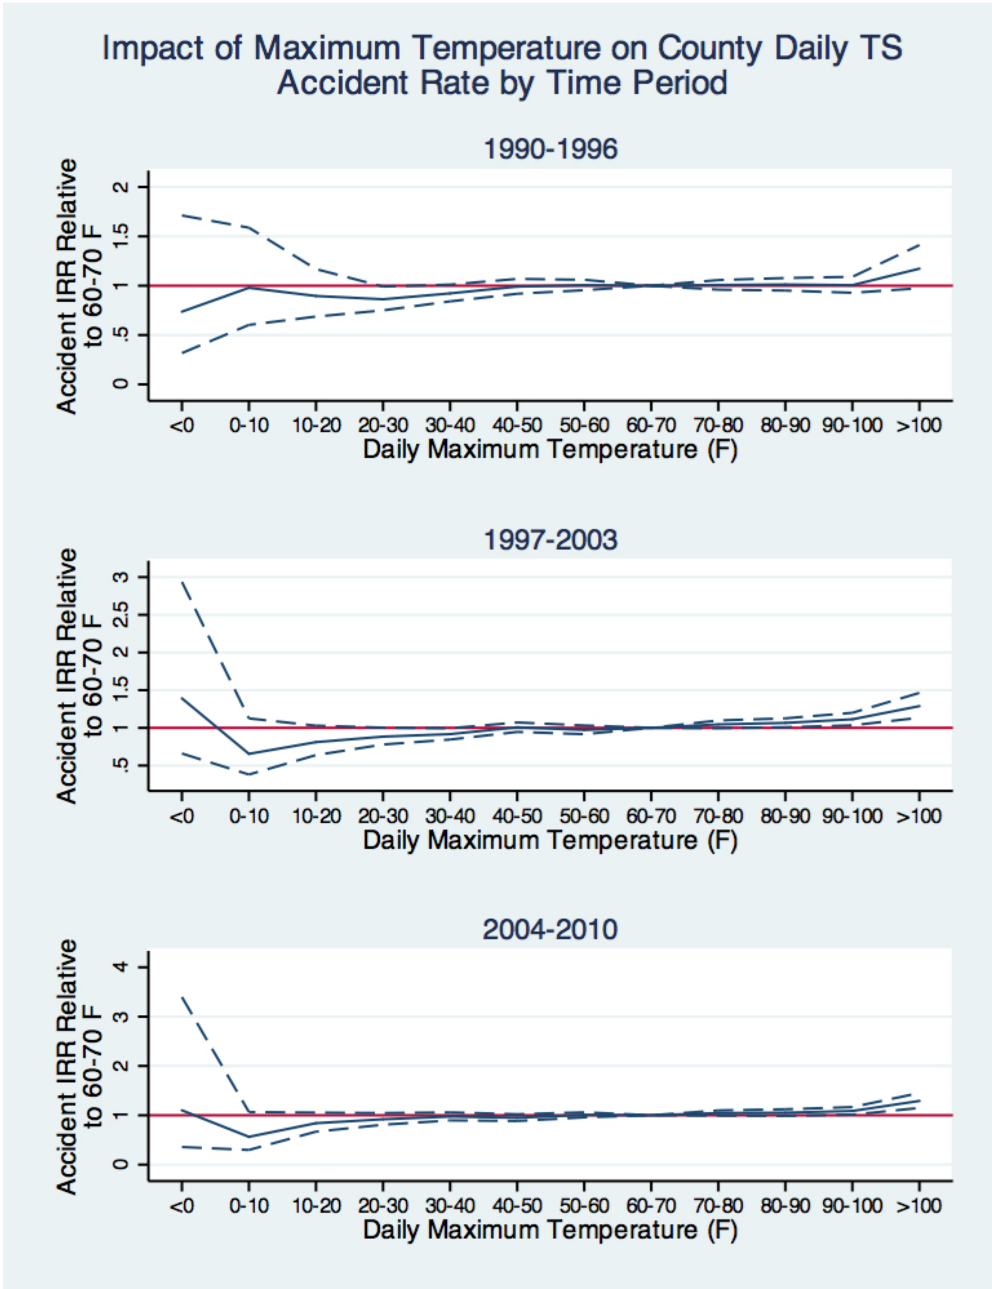

Figure 24: These figures present estimates for the impacts of daily maximum temperature on daily accident rate in temperature-sensitive industries for 1990-1996, 1997-2003, and 2004-2010. I stratify my analysis by time period. Coefficients give incident rate ratios relative to maximum temperature between 60 and 70°F. Dashed lines give the 95% confidence interval. While this graph interpolates between coefficients at 10°F intervals, my regression assumes that temperature impacts are constant within 10°ranges. Each regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects. N=5,165,613 for the 1990-1996, N=5,196,119 for 1997-2003, and N=5,215,335 for 2004-2010.

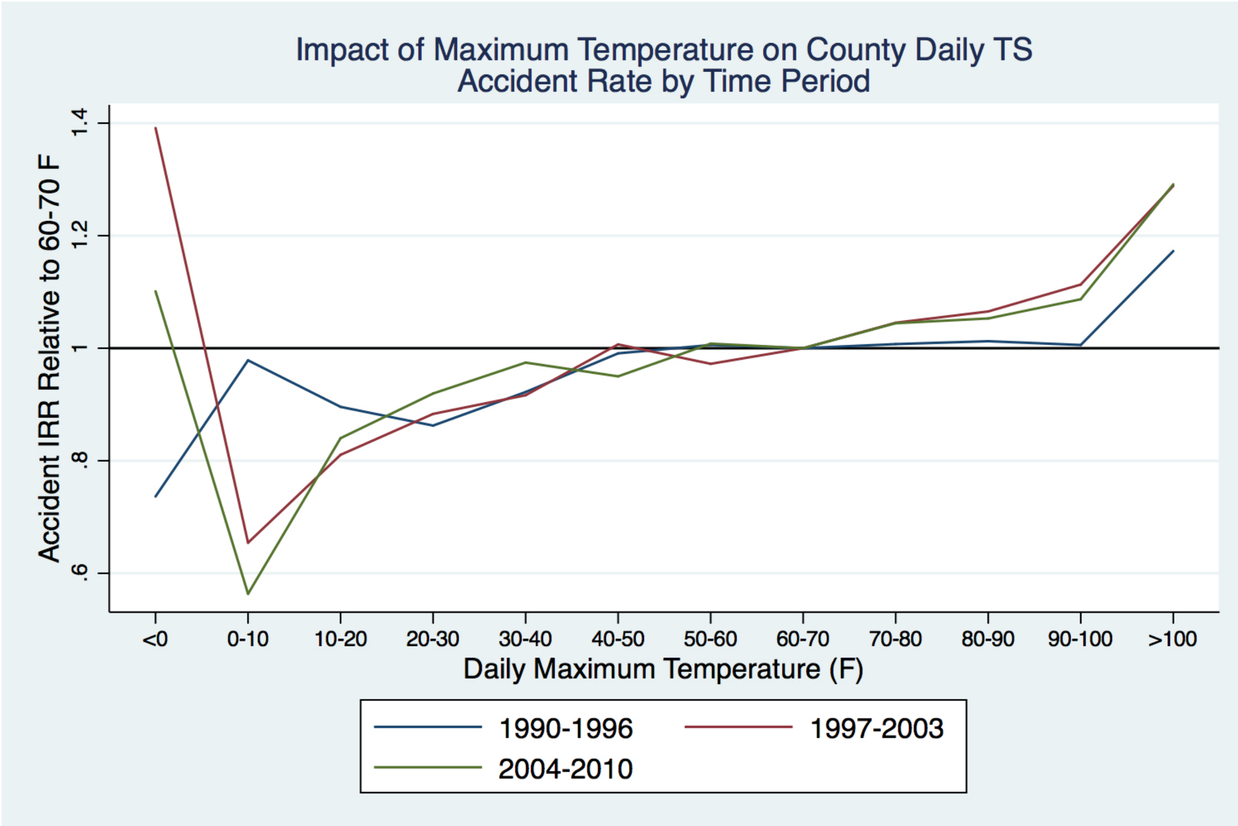

Figure 25: This figure combines my estimates for the impacts of temperature on accident rate by time period. Plots with confidence intervals are given in Figure 24. I stratify my primary Poisson analysis by time period, where each regression includes precipitation bins and county, year, region-by-month, and day-of-week fixed effects. N=5,165,613 for the 1990-1996, N=5,196,119 for 1997-2003, and N=5,215,335 for 2004-2010.

between 1997 and 2004, and 34.6% occurred between and 2004 and 2010. These results suggest that temperature may be a marginally strongly determinant of accident incidence in more recent time periods.

In addition to these differences in the significance of coefficients, there are some differences in the magnitudes of these coefficients across time periods. In general, however, those temperature coefficients that are significant are broadly consistent across time periods. For example, we see that days with maximum temperature between 70 and 80°F had about 4.5% more accidents between 1997 and 2003 and about 4.4% more accidents between 2004 and 2010, both relative to days with maximum temperature between 60 and 70°F. At low temperatures, days with maximum temperature between 30 and 40°F had 7.8% fewer accidents between 1990 and 1996 and 8.4% fewer accidents between 1997 and 2003. The most notable divergence in the impacts of temperature across time periods is for days with maximum temperature above 100°F. The impact of days over 100°F is similar in the two more recent time periods, with a 28.9% increase in accident rates between 1997 and 2003 and a 29.1% increase in accident rates between 2004 and 2010, both relative to days with maximum temperature between 60 and 70°F. In contrast, days with maximum temperature over 100°F have only 17.3% more accidents in the first period.

What can these results tell us about the role of adaptation in the impact of temperature on accident incidence? As I suggested, we might expect that the impact of temperature on accident incidence would attenuate over time if there has historically been adaptation in the relationship between temperature and accidents. However, we see here that there appears to be little difference in the impact of temperature on accident incidence across the three periods I have outlined here. Where there are differences, most notably at temperatures over 100°F, this impact is larger in magnitude in more recent periods. Thus, we do not see evidence of adaptation to the impacts of temperature on accident incidence. However, as in my comparison of the impact of temperature on accident incidence across regions, it is difficult to draw any substantive conclusions about adaptation from this simple comparison.

8 Future Research

Moving forward with my research, there is much work that could be done to more fully investigate how short-run impacts of temperature on accidents map into long-run costs of climate change. In particular, I hope to develop a better understanding of the extent and timing of adaptation. Comparisons of the impacts of the temperature on accident incidence across different historical climates or across time periods, like those which I have presented here, can provide only a basic first look at adaptation in the relationship between temperature and accidents. First, I hope to identify some source of granular data on output or time spent working that would allow me to look more directly at behavioral adaptation in response to temperature. While such data might not be available at a broad geographic scale, even data at the level of a particular factory or facility could provide a first look at behavioral adaptation. Next, several other forms of analysis may better allow me to more fully estimate the longer-run impacts of temperature on accident incidence, incorporating any adaptation or intensification effects. I may be able to estimate these impacts by estimating the same basic panel specification that I use here, but using measures of temperature that are averaged over longer time periods. Since these estimates would be based on a longer timestep, they are likely closer to giving the impact of temperature on accident incidence that we would expect to see under the timescales of climate change.

In addition to estimating my primary panel model at a longer time step, it could be useful to estimate a panel vector autoregressive (VAR) model or event analysis model to estimate a dynamic impulse-response function of the relationship between daily temperatures and accident rates at annual or longer time scales. A panel VAR model incorporates feedback from past shocks to allow the impact of temperature shocks on accident rates to vary over time and according to past accident rates (Bond, 2002). Thus, using a panel VAR would allow me to develop medium-term estimates of the response of worker safety levels to temperature shocks, allowing this response to play out over time. This impulse-response function would incorporate both physiological impacts of temperature and adaptive mechanisms that have played out at longer time scales, and thus may provide reasonable estimates of the longer-term impacts on temperature on accident incidence that we would see under climate change.

If adaptation relies on forward-looking investments, even these medium-run estimates will only incorporate adaptation to the extent that agents were aware of changes in average temperature and perceived those changes to be permanent shifts. If agents did not perceive and respond to shifts in historical data, these longer-term estimates will still fail to give appropriate measures of the potential for adaptation under climate change. I may be able to more rigorously estimate adaptation in the relationship between temperature and accidents using a method outlined by Schrader (2016) that makes better use the role of agent expectations in defining opportunities for adaptation. Schrader argues that since information about the future state of an environmental process will cause behavioral shifts among forward-looking agents, we can estimate forward-looking adaptation by evaluating the impact of forecast information on our outcome variable. Then, the direct impact of temperature on the economic outcome is given by the remaining impact of weather, conditional on those expectations. Studying and adapting this technique to my temperature-accident framework might provide a rigorous method by which to estimate the role of adaptation in my results.

Besides this future work on adaptation in response to temperature, I hope to improve my current analysis by addressing any concerns of cross-sectional dependence with spatial econometric techniques. I also plan to broaden my analysis by looking more closely at the relationship between temperature and accident incidence in non-temperature-sensitive industries. Cachon et al.'s (2013) analysis of the impacts of weather on productivity in the U.S. automobile industry suggests that the impacts of temperature may extend beyond those industries where workers are most directly exposed to weather conditions. If I find that temperature has a similar impact on worker safety in industries that I have classified as non-temperature-sensitive, those results might point to other mechanisms in the temperature-accident relationship than those that I have proposed here.

9 Conclusions

In my research, I have investigated the potential costs of climate change associated with changes in the incidence of occupational accidents among a set of "temperature-sensitive industries," like construction, agriculture, forestry, manufacturing, and utilities servicing. Research in physiology suggests that extreme temperature make workers more tired and impair cognitive functions like coordination, vigilance, and reaction time, thereby increasing their risk of accidents when working in extreme temperature conditions. Thus, we might expect to see increased incidence of workplace accidents as temperatures rise under climate change. At the same time, behavior responses to extreme temperatures might reduce this impact if, for example, workers postpone work to cooler conditions or take additional rest breaks. I estimate the short-term impact of temperature on accident incidence with daily data for a panel of over 2000 counties in the United States between 1990 and 2010. Using a suite of seasonal, temporal, and geographic fixed effects, I am able to isolate the impacts of plausibly exogenous day-to-day fluctuations in temperature. I primarily rely on Poisson regression, and I estimate the accident incident rate ratios (IRR) for a series of twelve ten-degree F bins for daily maximum temperature ranging from maximum temperature below 0°F to maximum temperature above 100°F.

Overall, I find that significantly more accidents occur at high temperature extremes and that significantly fewer accidents occur at low temperature extremes. This relationship is nonlinear, spiking sharply upwards at the highest temperature extremes. At high temperatures, these impacts range from 3.8% more accidents on days with maximum temperature between 70 and 80°F to 30.0% more accidents on days with maximum temperature over 100°F, both relative to a day with maximum temperature between 60 and 70°F. At low temperatures, these impacts range from 7.0% fewer accidents on days with maximum temperature between 30 and 40°F to 30.1% fewer accidents on days with maximum temperature between 0 and 10°F, both relative to a day with maximum temperature between 60 and 70°F.

As the climate change, hot days are projected to become increasingly frequent and cold days increasingly infrequent. Then, my estimates imply that climate change will increase the incidence of occupational accidents in the United States. How large do we expect these impacts to be? Using climate projection data derived from the A1F1 scenario in the Hadley 3 climate model (Deschênes and Greenstone, 2011), I project percentage changes in accident rate averaged across the years 2070 through 2099. Using these rates, I project that we should expect an average of between 22 and 67 additional fatal occupational accidents in temperature-sensitive industries between 2070 and 2099. Using Leigh's (2011) estimates for the ratio of fatal to nonfatal accidents in 2007, we would expect an average of between 33,600 and 102,400 addition nonfatal accidents between 2070 and 2099. Leigh's estimates for average costs of fatal and nonfatal accidents imply that these additional accidents could impose costs of between $750 million and $2.30 billion between 2070 and 2099.

It is important to note that these estimates of climate impacts are highly uncertain; as climate change progresses, many features of work in temperature-sensitive industries will change. These changes may alter workers' exposure to extreme temperatures, modifying the impacts of temperature on worker safety. In particular, my estimates of the short-term impact of temperature on accident incidence might diverge from the long-term impacts of climate change on accident incidence through a combination of adaptation, or the process by which economic actors develop behavioral mechanisms, policies, or technologies that reduce the impacts of those temperature shocks on accident incidence, and intensification, where shifts in weather patterns associated with climate change produce damages to economic outcomes larger than those revealed by transient, short-term weather fluctuations. If there is significant opportunity for adaptation in the relationship between temperature and occupational accidents, we would expect the magnitude of the long-run impact of a permanent change in climate on accident incidence to be smaller than that of the short-run impact of temperature that I have estimated here. In contrast, any intensification would imply that the long-term impacts of climate change are greater than the short-term impacts that I estimate here.

While I can make no definite projections of the magnitude of the costs of climate change associated with changes in the incidence of occupational accidents, my analysis suggests that changes in worker safety may be a significant impact of climate change. If so, future research on the impacts of temperature on accident rate may prove salient both to climate and worker safety policy.

10 Bibliography

Abrigo, Michael R.M., and Inessa Lova. Estimation of Panel Vector Autoregression in Stata: a Package of Programs. For Panel Data Conference, Central European University, Budapest, February 2015.

Acclimatization: Adjusting to the Temperature, University of Iowa Hospitals and Clinics Health Topics. https://www.uihealthcare.org/health-library/acclimatization-adjusting-to-the-temperature/

Adam-Poupart, Ariane, Audrey Smargiassi, Marc-Antoine Busque, Patrice Duguay, Michel Fournier, Joseph Zayed, and France Labreche. Summer Outdoor Temperature and Occupational Heat-Related Illnesses in Quebec (Canada). Environmental Research 134 (2014): 339-344.

Azer, N.Z., P.E. McNall, and H.C. Leung. Effects of Heat Stress on Performance. Ergonomics 15 (1972): 681-691.

Bell, C.R., K.A. Provins, and R.W. Hiorns. Visual and Auditory Vigilance During Exposure to Hot and Humid Conditions. Ergonomics 7 (1964): 279-288.

Bond, Stephen R. Dynamic Panel Data Models: a Guide to Micro Data Methods and Practice. CEMMAP Working Paper, Centre for Microdata Methods and Practice, No. CWP09/02.

Burke, Marshall, Hsiang, Solomon M., and Edward Miguel. Global Non-linear Effect of Temperature on Economic Production. Nature (2015): 235-239.

Cachon, Gerard P., Gallino, Santiago, and Marcelo Olivares. Severe Weather and Automobile Assembly Productivity. (2013). Working Paper.

Canova, Fabio, and Matteo Ciccarelli. Panel Vector Autoregressive Models: A Survey. Working Paper Series, European Central Bank (2013).

Centers for Disease Control and Prevention. Heat-Related Deaths Among Crop WorkersUnited States, 1992-2006. Journal of the American Medical Association 300, no. 9 (2008): 1017-1018.

Chiles, W.D. Effects of Elevated Temperatures on Performance of a Complex Mental Task. Ergonomics 2, no. 1 (1958): 89-96.

Cold Stress. Occupational Health Clinics for Ontario Workers Inc. 2005. *http : //www.ohcow.on.ca/uploads/Clini*

Colquhoun, W.P. Effects of Raised Ambient Temperature and Event Rate on Vigilance Performance. Aerospace Medicine 40 (1969): 413-417.

Deryugina, Tatyana and Solomon M. Hsiang. Does the Environment Still Matter? Daily Temperature and Income in the United States. (2014). NBER Working Paper.

Deschenes, Olivier. Temperature, Human Health, and Adaptation: A Review of the Empirical Literature. Energy Economics 46 (2014): 606-619.

Deschenes, Olivier and Michael Greenstone. Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. American Economic Journal: Applied Economics 3, no. 4 (2011): 152-185.

EPA. Climate Change in the United States: Benefits of Global Action. (2015). United States Environmental Protection Agency, Office of Atmospheric Programs, EPA 430-R-15-001.

Epstein, Y., G. Keren, J. Moisseiev, O. Gasko, and S. Yachin. Psychomotor Deterioration During Exposure to Heat. Aviation, Space and Environmental Medicine 51 (1980): 607-610.

Field, C.B., V.R. Barros, K.J. Mach, M.D. Mastrandrea, M. van Aalst,W.N. Adger, D.J. Arent, J. Barnett, R. Betts, T.E. Bilir, J. Birkmann, J. Carmin, D.D. Chadee, A.J. Challinor, M. Canning, David, and Peter Pedroni. Infrastructure, Long-Run Economic Growth and CausalityTests for Cointegrated Panels. The Manchester School 76, no. 5 (2008): 504-527.

Chatterjee,W. Cramer, D.J. Davidson, Y.O. Estrada, J.-P. Gattuso, Y. Hijioka, O. Hoegh-Guldberg, H.Q. Huang, G.E. Insarov, R.N. Jones, R.S. Kovats, P. Romero-Lankao, J.N. Larsen, I.J. Losada, J.A. Marengo, R.F. McLean, L.O. Mearns, R. Mechler, J.F. Morton, I. Niang, T. Oki, J.M. Olwoch, M. Opondo, E.S. Poloczanska, H.-O. Prtner, M.H. Redsteer, A. Reisinger, A. Revi, D.N. Schmidt, M.R. Shaw, W. Solecki, D.A. Stone, J.M.R. Stone, K.M. Strzepek, A.G. Suarez, P. Tschakert, R. Valentini, S. Vicua, A. Villamizar, K.E. Vincent, R.Warren, L.L.White, T.J.Wilbanks, P.P.Wong, and G.W. Yohe, 2014: Technical summary. In: Climate Change 2014: Impacts,Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Field, C.B., V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T.E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, and L.L.White (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2014.

Dell, Melissa, Jones, Benjamin F., and Benjamin A. Olken. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. American Economic Journal: Macroeconomics 4, no. 3 (2012): 66-95.

Dell, Melissa, Jones, Benjamin F., and Benjamin A. Olken. What Do We Learn From the Weather? The New Climate-Economy Literature. Journal of Economic Literature 52, no. 3 (2014): 740-798.

Dong, Xiuwen S., Fujimoto, Alissa, Ringen, Knut, Stafford, Erich, Platner, James W., Gittleman, Janie L., and Xuanwen Wang. Injury Underreporting Among Small Establishments in the Construction Industry. American Journal of Industrial Medicine 54 (2011): 339-349.

Fine, B. J., and J.L. Kobrick. Effects of Altitude and Heat on Complex Cognitive Tasks. Human Factors 20 (1978): 115-122.

Fogleman, Maxwell, Layla Fakhrazadeh, and Thomas E. Bernard. The Relationship Between Oudoor Thermal Conditions and Acute Injury in an Aluminum Smelter. International Journal of Industrial Ergonomics 35, no. 1 (2005): 47-55.

Fraser, D.C. Environmental Stress and its Effect on Performance. Occupational Psychology 21 (1957): 248-255.

Grether, W.F., Harris, C.S., Mohr, G.C., Nixon, C.W., Ohlbaum, M., Sommer, H.C., Thaler, V.H., Veghte, J.H. Effects of Combined Heat, Noise and Vibration Stress on Human Performance and Physiological Functions. Aerospace Medicine 42 (1971): 1092-1097.

Hancock, P.A. and I. Vasmatzidis. Effects of Heat Stress on Cognitive Performance: the Current State of Knowledge. International Journal of Hyperthermia 19, no. 3 (2003): 355-372.

Hsiang, Solomon M. Temperature and Cyclones Strongly Associated with Economic Production in the Caribbean and Central America. Proceedings of the National Academy of Sciences of the United States of America 107, no. 35 (2010): 15367-15372.

Intergovernment Panel on Climate Change. 3.3.5 Consistency and Relationships between Temperature and Precipitation. Working Group I: The Physical Science Basis (2007).

Jacob, Brian, Lars Lefgren, and Enrico Moretti. The Dynamics of Criminal Behavior: Evidence from Weather Shocks.

Jackson, Larry L. and Howard R. Rosenberg. Preventing Heat-Related Illness Among Agricultural Workers. Journal of Agromedicine 15 (2010): 200-215.

Jones, Benjamin F. and Benjamin A. Olken. Climate Shocks and Exports. American Economic Review 100 (2010); 454-459.

Kopp, Robert, DJ Rasmussen, Michael Mastrandrea, Robert Muir-Wood, Paul Wilson, Solomon Hsiang, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, Kate Larsen, and Trevor Houser. American Climate Prospectus: Economic Risks in the United States. New York: Rhodium Group, LLC, 2014. The Journal of Human Resources 42, no. 3 (2007): 489-527.

Kjellstrom, Tord, Ingvar Holmer, and Bruno Lemke. Workplace Heat Stress, Health and Productivityan Increasing Challenge for Low and Middle-Income Countries During Climate Change. (2009) Global Health Action 2.

Kjellstrom, Tord, Sabine Gabrysch, Bruno Lemke, and Keith Dear. The Hothaps Programme fro Assessing Climate Change Impacts on Occupational Health and Productivity: an Invitation to Carry Out Field Studies. (2009) Global Health Action 2.

Leigh, J. Paul. Economic Burden of Occupational Injury and Illness in the United States. The Milbank Quarterly 89, no. 4 (2011): 728-772.

Lewis, M.Y., G.B. Meese, and R. Kok. The Effects of Moderate Cold and Heat Stress on the Potential Work Performance of Industrial Workers. (Pretoria: National Building Research, Institute Council for Scientific and Industrial Research), CSIR Rep. No. 589.

Lind, A.R. and Bass, D.E. Optimal Exposure Time for Development of Acclimatization to Heat. Proceedings of American Societies for Experimental Biology 22 (1963): 704-708.

Lucas, Rebekah A., Yoram Epstein, and Tord Kjellstrom. Excessive Occupational Heat Exposure: a Significant Ergonomic Challenge and Health Risk for Current and Future Workers. Extreme Physiology and Medicine 23, no. 3: 14.

Lundgren, Karin, Kalev Kuklane, Chuansi Gao, and Ingvar Holmer. Effects of Heat Stress on Working Populations when Facing Climate Change. Industrial Health 51 (2013): 1-13.

Mackworth, N.H. The Breakdown of Vigilance During Prolonged Visual Search. Quarterly Journal of Experimental Psychology 1, no. 1 (1948): 6-21.

Morabito, Marco, Lorenzo Cecchi, Alfonso Crisci, Pietro Amedeo Modesti, and Simone Orlandini. Relationships Between Work-Related Accidents and Weather Conditions in Tuscany (Central Italy). Industrial Health 44, no.3 (2006): 458-464.

Morantz, Alison. Filing Not Found: Which Injuries Go Unreported to Worker Protection Agencies, and Why? Stanford Law School, 2014.

National Climatic Data Center/NESDIS/NOAA/U.S. Department of Commerce, National Weather Service/NOAA/U.S. Department of Commerce, and Federal Aviation Agency/U.S. Department of Transportation. NCDC TD3200 U.S. Cooperative Summary of Day, 1890(1948)-cont. 1981, updated monthly. Distributed by Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory. http://rda.ucar.edu/datasets/ds510.0/.

National Institute for Occupational Safety and Health. Occupational Exposure to Hot Environments: Revised Criteria 1986. U.S. Department of Health and Human Services, Public Health Service, Centers for Disease Control. Washington, D.C., 1986.

Occupational Safety and Health Adminstration (OSHA). OSHA Enforcement Inspection and Accident Investigation Data. April 7, 2010, updated daily. Distributed by U.S. Department of Labor, Enforcement Data. http://ogesdw.dol.gov/homePage.php.

Occupational Safety and Health Administration (OSHA). Cold Stress Guide. United States Department of Labor. https://www.osha.gov/SLTC/emergencypreparedness/guides/cold.html.

Osborne, Ethel E., Horace M. Vernon, and Bernand Muscio. The Influence of Temperature and Other Conditions on the Frequency of Industrial Accidents in Two Contributions to the Study of Accident Causation. London: H.M. Stationery Office, 1922.

Pepler, Richard D. Warmth and Performance: An Investigation in the Tropics. Journal of Comparative and Physiological Psychology 2, no. 1 (1959): 446-475.

Pepler, Richard D. Warmth, Glare and a Background of Quiet Speech: A Comparison of their Effects on Performance. Ergonomics 3 (1960): 68-73.

Pilcher, June J., Eric Nadler, and Caroline Busch. Effects of Hot and Cold Temperature Exposure on Performance: a Meta-Analytic Review. Ergonomics 45, no. 10 (2002): 682-698.

Powell, Philip. 2000 Accidents: A Shop Floor Study of their Causes Based on 42 Months Continuous Observation. London: National Institute of Industrial Psychology, 1971.

Ramsey, Jerry D., and Charles L. Burford. Effects of Workplace Thermal Conditions on Safe Work Behavior. Journal of Safety Research 14, no. 3 (1983): 105-114.

Ramsey, J.D., D. Dayal, and B. Ghahramani. Heat Stress Limits for Sedentary Workers. American Industrial Hygiene Association Journal 36 (1975): 259-265.

Ramsey, Jerry D. and Yeong G. Kwon. Recommended Alert Limits for Perceptual Motor Loss in Hot Environments. International Journal of Industrial Ergonomics 9 (1992): 245-257.

Ramsey, J.D., and S.B. Pai. Sedentary Work by Females in Hot Environments. In Proceedings of the Human Factors Society 19th Annual Meeting (Santa Monica, CA: Human Factors Society), 498-500.

Ranson, Matthew. Crime, Weather, and Climate Change. Journal of Environmental Economics and Management 67, no. 3 (2014): 274-302. Rodrguez, G. Lecture Notes on Generalized Linear Models. (2007). http://data.princeton.edu/wws509/notes/

Rosenman, Kenneth D., Kalush, Alice, Gardiner, Joseph C., Reeves, Matthew, and Zhewui Luo. How Much Work-Related Injury and Illness is Missed by the Current National Surveillance System? Journal of Occupational and Environmental Medicine 48, no. 4 (2006): 357-365.

Stocks, Jodie M., Nigel A.S. Taylor, Michael J. Tipton, and John E. Greenlead. Human Physiological Responses to Cold Exposure. Aviation, Space, and Environmental Medicine 75, no. 5 (2004): 444-457.

Tawatsupa, Benjawan, Vasoontara Yiengprugsawan, Vasoontara, Tord Kjellstrom, Janneke Berecki-Gisolf, Sam-Ang Seubsman, and Adrian Sleight. Association Between Heat Stress and Occupational Injury among Thai Workers: Findings of the Thai Cohort Study. Industrial Health 51, no. 1 (2013):

34-46.

U.S. Department of Labor. Data Enforcement Frequently Asked Questions. $http : //ogesdw.dol.gov/views/faq.php$

U.S. Department of Labor. Occupational Safety and Health Administration. Occupational Injury and Illness Recording and Reporting Requirements. 29 CFR Parts 1904 and 1952. Federal Registry # 61:4029-4067. (1996). Washington D.C.: U.S. Government Printing Office.

U.S. Department of Labor. Occupational Safety and Health Administration. Occupational Injury and Illness Recording and Reporting RequirementsNAICS Update and Reporting Revisions. 29 CFR Part 1904. Federal Registry # 79: 56129-56188. (2014). Washington, D.C.: U.S. Government Printing Office.

U.S. House of Representatives. Committee on Education and Labor. (2008). Hidden Tragedy: Underreporting of Workplace Injuries and Illnesses: A Majority Staff Report

Wagner, J.A., Robinson, S., Tzankoff, S.P., Marino, R.P. Heat Tolerance and Acclimatization to Work in the Heat in Relation to Age. Journal of Applied Physiology 33, no. 5 (1972): 616-622.

WHO. Health Factors Involved in Working Under Conditions of Heat Stress. (1969). Geneva: World Health Organization.

Wooldridge, Jeffrey M. Econometric Analysis of Cross Section and Panel Data. Cambridge: MIT Press, 2001.

Wuellner, Sara E. and David K. Bonauto. Exploring the Relationship Between Employer Recordkeeping and Underreporting in the BLS Survey of Occupational Injuries and Illnesses. American Journal of Industrial Medicine 57 (2014): 1133-1143.

Wyon, D., Wyon, I., Norin, F. Effects of Moderate Heat Stress on Driver Vigilance in a Moving Vehicle. Ergonomics 39 (1996): 61-75.

Xiang, Jianjun, Peng Bi, Dino Pisaniello, and Alana Hansen. Health Impacts of Workplace Heat Exposure: an Epidemiological Review. Industrial Health 52, no. 2 (2014): 91-101.

Xiang, Jianjun, Peng Bi, Dino Pisaniello, and Alana Hansen. The Impact of Heatwaves on Workers Health and Safety in Adelaide, South Australia. Environmental Research 133 (2014): 90-95.

Xiang, Jianjun, Peng Bi, Dino Pisaniello, Alana Hansen, and Thomas Sullivan. Association Between High Temperature and Work-Related Injuries in Adelaide, South Australia, 2001-2010. Occupational and Environmental Medicine 71, no. 4 (2014): 246-252.

Zivin, Joshua Graff and Matthew Neidell. Temperature and the Allocation of Time: Implications for Climate Change. Journal of Labor Economics 32, no. 1 (2014): 1-26.

11 Appendices

Here, I list the industries that I classify as "temperature-sensitive" throughout my analysis. The first table gives associated codes under the Standard Industrial Classification (SIC) system, and the second table gives associated codes under the North American Industry Classification System (NAICS). Note that I do not classify industries that are not covered under OSHA and which therefore do not appear in my accident data. For example, I do not include mining in this list of temperature-sensitive industries.

Temperature-Sensitive Industry	**Standard Industrial Classification (SIC) Code**
SIC Divisions	
Agriculture, Forestry, and Fishing	Division A
Construction	Division C
Manufacturing	Division D
Transportation, Communications, Electric, Gas, and Sanitary Services	Division E
SIC Major Groups (2-digit codes)	
Oil and Gas Extraction	Division B, Major Group 13
Wholesale Trade-durable Goods	Division F, Major Group 50
Automotive Repair, Services, and Parking	Division I, Major Group 75
Miscellaneous Repair Services	Division I, Major Group 76
SIC 4-Digit Codes	
Sporting and Recreational Camps	Division I, 7032
Recreational Vehicle Parks and Campsites	Division I, 7033
Heavy Construction Equipment Rental and Leasing	Division I, 7353
Public Golf Courses	Division I, 7992
Amusement Parks	Division I, 7996
Membership Sports and Recreation Clubs	Division I, 7997
Police Protection	Divison J, 9221
Fire Protection	Division J, 9224

Temperature-Sensitive Industry	North American Industry Classification System (NAICS) Code
NAICS Sectors	
Agriculture, Forestry, Fishing, and Hunting	Sector 11
Utilities	Sector 22
Construction	Sector 23
Manufacturing	Sector 31-33
Wholesale Trade	Sector 42
Transportation and Warehousing	Sector 48-49
NAICS Subsectors	
Oil and Gas Extraction	Subsector 211
Waste Management and Remediation Services	Subsector 562
Repair and Maintenance	Subsector 811
NAICS Industry Groups	
Amusement Parks and Arcades	Industry Group 7131
Other Amusement and Recreation Industries	Industry Group 7139
NAICS Industries	
Police Protection	NAICS Industry 92212
Fire Protection	NAICS Industry 92216
NAICS National Industries	
Drilling Oil and Gas Wells	National Industry 213111
Support Activities for Oil and Gas Operatinos	National Industry 213112

Table 9: Primary OLS analysis

VARIABLES	TS Accident Count
Max Temp <0 F	0.000691***
	(0.000180)
Max Temp 0-10 F	0.000387***
	(0.000134)
Max Temp 10-20 F	0.000350***
	(0.000118)
Max Temp 20-30 F	0.000403***
	(0.000111)
Max Temp 30-40 F	0.000308***
	(0.0000936)
Max Temp 40-50 F	0.000146**
	(0.0000596)
Max Temp 50-60 F	-0.0000293
	(0.0000544)
Max Temp 70-80 F	0.000158***
	(0.0000498)
Max Temp 80-90 F	0.000255***
	(0.0000641)
Max Temp 90-100 F	0.000369***
	(0.0000975)
Max Temp >100 F	0.00144***
	(0.000321)
Mean TS Accident Count	0.00250
Observations	23,614,877
R-squared	0.093
County FE	YES
Month*Region FE	YES
Year FE	YES
Day of Week FE	YES

SEs clustered by county in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 10: OLS regressions accounting for employment

VARIABLES	(1) Heat-Sens Acc	(2) Heat-Sens Acc	(3) HS Acc/Million Workers	(4) HS Acc/Million TS Workers
Max Temp <0 F	0.000694***	0.000601***	0.0474	-0.120
	(0.000180)	(0.000164)	(0.0882)	(0.245)
Max Temp 0-10 F	0.000390***	0.000331***	-0.0507**	-0.339*
	(0.000134)	(0.000122)	(0.0219)	(0.190)
Max Temp 10-20 F	0.000353***	0.000303***	-0.0259	-0.219
	(0.000118)	(0.000107)	(0.0211)	(0.168)
Max Temp 20-30 F	0.000404***	0.000371***	0.0147	-0.0722
	(0.000111)	(0.000102)	(0.0226)	(0.125)
Max Temp 30-40 F	0.000310***	0.000284***	0.00102	0.122
	(9.39E-05)	(8.55E-05)	(0.0155)	(0.156)
Max Temp 40-50 F	0.000147**	0.000133**	-0.00641	-0.00525
	(5.98E-05)	(5.63E-05)	(0.0118)	(0.0396)
Max Temp 50-60 F	-2.92E-05	-2.65E-05	-0.00637	-0.0142
	(5.46E-05)	(5.41E-05)	(0.00981)	(0.0328)
Max Temp 70-80 F	0.000160***	0.000154***	0.00387	-0.00014
	(5.00E-05)	(4.98E-05)	(0.00900)	(0.0304)
Max Temp 80-90 F	0.000256***	0.000253***	-0.0117	-0.0618
	(6.43E-05)	(6.30E-05)	(0.0114)	(0.0409)
Max Temp 90-100 F	0.000372***	0.000381***	0.0317	0.0513
	(9.78E-05)	(9.62E-05)	(0.0354)	(0.122)
Max Temp >100 F	0.00145***	0.00139***	0.0278	-0.0791
	(0.000322)	(0.000301)	(0.0620)	(0.186)
HS Employment		2.90e-07***		
		(8.11E-08)		
Observations	23,527,479	23,527,479	23,527,479	23,527,479
R-squared	0.088	0.088	0.000	0.000
County FE	YES	YES	YES	YES
Month*Region FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Day of Week FE	YES	YES	YES	YES

SEs clustered by county in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 11: OLS regression of monthly employment on temperature bins

	(1)	(2)
VARIABLES	Log(Employment)	Log(HS Employment)
# Days Max Temp <0 F	-0.000788	-0.00234***
	(0.000811)	(0.000898)
# Days Max Temp 0-10 F	-0.00104***	-0.00267***
	(0.000402)	(0.000440)
# Days Max Temp 10-20 F	-0.000679**	-0.00137***
	(0.000276)	(0.000289)
# Days Max Temp 20-30 F	-0.000981***	-0.00145***
	(0.000237)	(0.000279)
# Days Max Temp 30-40 F	-0.000699***	-0.000954***
	(0.000159)	(0.000180)
# Days Max Temp 40-50 F	-0.000557***	-0.000354**
	(0.000148)	(0.000170)
# Days Max Temp 50-60 F	-0.000139	-0.000176
	(0.000147)	(0.000152)
# Days Max Temp 70-80 F	-0.000227	-0.000216
	(0.000166)	(0.000158)
# Days Max Temp 80-90 F	-0.000879***	-0.000800***
	(0.000187)	(0.000172)
# Days Max Temp 90-100 F	-0.00154***	-0.00130***
	(0.000225)	(0.000217)
# Days Max Temp >100 F	-0.00128***	-0.00107**
	(0.000353)	(0.000468)
Observations	773,001	773,001
R-squared	0.993	0.992
County FE	YES	YES
Year FE	YES	YES
Month*Region FE	YES	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 12: Primary Poisson analysis

	(1) Full Sample	(2) Excluding CA
VARIABLES	HS Accidents IRR	
Max Temp <0 F	0.874	0.867
	(0.258)	(0.256)
Max Temp 0-10 F	0.699**	0.709**
	(0.119)	(0.121)
Max Temp 10-20 F	0.790***	0.806***
	(0.0602)	(0.0619)
Max Temp 20-30 F	0.894***	0.907**
	(0.0356)	(0.0374)
Max Temp 30-40 F	0.930***	0.944**
	(0.0246)	(0.0274)
Max Temp 40-50 F	0.977	1.002
	(0.0224)	(0.0235)
Max Temp 50-60 F	0.993	1.003
	(0.0158)	(0.0208)
Max Temp 70-80 F	1.038**	1.021
	(0.0153)	(0.0198)
Max Temp 80-90 F	1.052***	1.056**
	(0.0194)	(0.0242)
Max Temp 90-100 F	1.082***	1.098***
	(0.0249)	(0.0333)
Max Temp >100 F	1.300***	1.385***
	(0.0392)	(0.0730)
Observations	20,955,316	19,480,917
Number of Counties	2661	2590
County FE	YES	YES
Month*Region FE	YES	YES
Year FE	YES	YES
Day of Week FE	YES	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

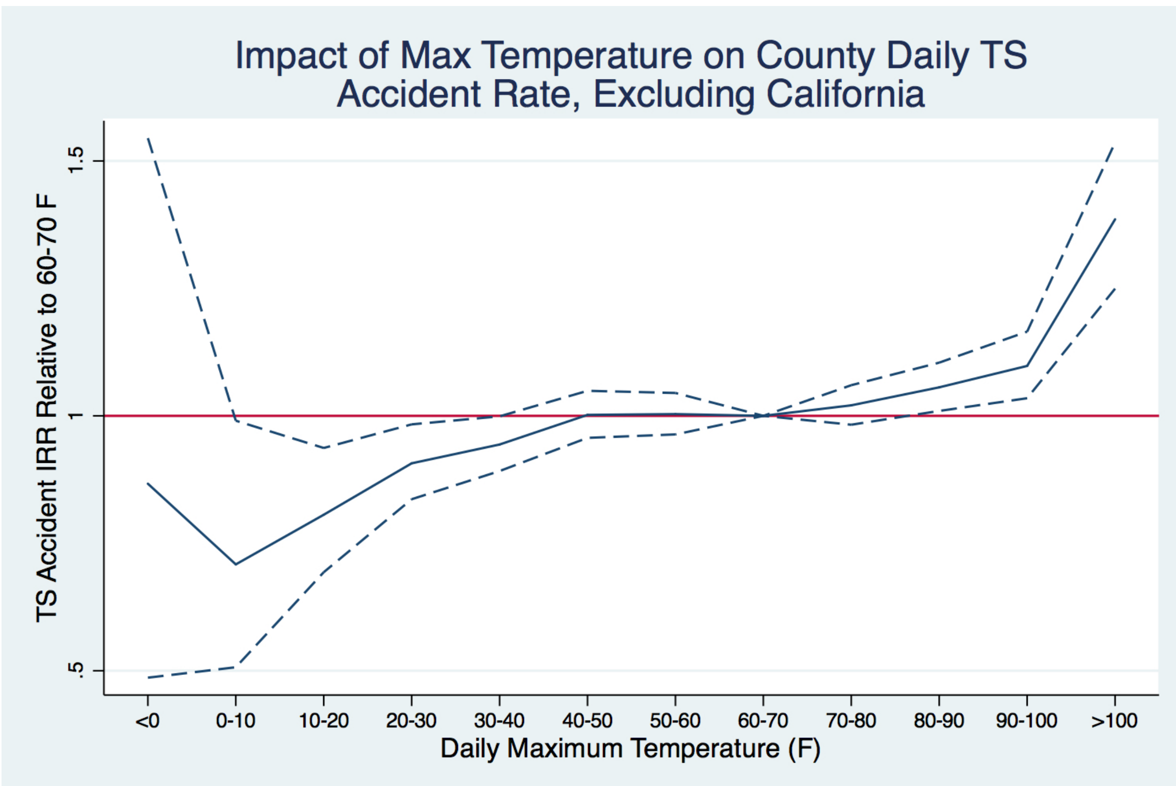
Impact of Max Temperature on County Daily TS
Accident Rate, Excluding California
TS Accident IRR Relative to 60-70 F
.5
1
1.5
<0
0-10
10-20
20-30
30-40
40-50
50-60
60-70
70-80
80-90
90-100
>100
Daily Maximum Temperature (F)

Table 13: 5-degree bin Poisson analysis

VARIABLES	TS Accident IRR	VARIABLES	TS Accident IRR
Max Temp < -5 F	1.132	Max Temp 70-75 F	1.035*
	(0.497)		(0.0204)
Max Temp -5-0 F	0.721	Max Temp 75-80 F	1.048**
	(0.308)		(0.0223)
Max Temp 0-5 F	0.898	Max Temp 80-85 F	1.058**
	(0.232)		(0.0242)
Max Temp 5-10 F	0.609**	Max Temp 85-90 F	1.054**
	(0.131)		(0.0262)
Max Temp 10-15 F	0.906	Max Temp 90-95 F	1.073***
	(0.115)		(0.0290)
Max Temp 15-20 F	0.743***	Max Temp 95-100 F	1.114***
	(0.0685)		(0.0406)
Max Temp 20-25 F	0.883*	Max Temp 100-105 F	1.294***
	(0.0604)		(0.0472)
Max Temp 25-30 F	0.908**	Max Temp ≥ 105 F	1.363***
	(0.0411)		(0.0959)
Max Temp 30-35 F	0.933*		
	(0.0337)		
Max Temp 35-40 F	0.938**		
	(0.0291)		
Max Temp 40-45 F	1.002		
	(0.0294)		
Max Temp 45-50 F	0.966		
	(0.0254)		
Max Temp 50-55 F	1.010		
	(0.0214)		
Max Temp 55-60 F	0.984		
	(0.0202)		
Max Temp 65-70 F	1.005		
	(0.0185)		
	Observations	19,917,285	
	Number of county	2647	
	Standard FE	YES	

SEs clustered by county in parentheses
*** p¡0.01, ** p¡0.05, * p¡0.1

Table 14: State-level primary Poisson analysis

VARIABLES	TS Accident IRR
Max Temp <0 F	0.662
	(0.299)
Max Temp 0-10 F	0.696***
	(0.0863)
Max Temp 10-20 F	0.788***
	(0.0593)
Max Temp 20-30 F	0.919**
	(0.0365)
Max Temp 30-40 F	0.927**
	(0.0285)
Max Temp 40-50 F	0.960
	(0.0289)
Max Temp 50-60 F	0.998
	(0.0134)
Max Temp 70-80 F	1.023*
	(0.0124)
Max Temp 80-90 F	1.064***
	(0.0163)
Max Temp 90-100 F	1.144***
	(0.0249)
Max Temp >100 F	1.233***
	(0.0891)
Observations	375,649
Number of states	49
State FE	YES
Month*Region FE	YES
Year FE	YES
Day of Week	YES

SEs clustered at state level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 15: Nationwide primary OLS and Poisson analysis

VARIABLES	(1) Poisson TS Accident IRR	(2) OLS TS Accident Count
Max Temp 10-20 F	0.357***	-4.224
	(0.111)	(2.649)
Max Temp 20-30 F	0.738***	-1.869***
	(0.0641)	(0.643)
Max Temp 30-40 F	0.875***	-0.907***
	(0.039)	(0.334)
Max Temp 40-50 F	0.910***	-0.678**
	(0.0322)	(0.273)
Max Temp 50-60 F	0.937**	-0.487**
	(0.0282)	(0.230)
Max Temp 70-80 F	1.010	0.0948
	(0.0299)	(0.233)
Max Temp 80-90 F	1.028	0.250
	(0.0414)	(0.314)
Max Temp 90-100 F	1.143*	1.280**
	(0.0783)	(0.571)
Observations	7670	7670
Month FE	YES	YES
Year FE	YES	YES
Day of Week	YES	YES

Robust SEs in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 16: Heatwaves and Accident Incidence, Poisson

	(1)	(2)	(3)
	Primary	Lag3>85	Lag3>95
VARIABLES	Temperature-Sensitive Accidents		
Max Temp <0 F	1.025	1.022	1.024
	(0.269)	(0.268)	(0.268)
Max Temp 0-10 F	0.755*	0.753*	0.755*
	(0.116)	(0.116)	(0.116)
Max Temp 10-20 F	0.824***	0.821***	0.823***
	(0.0582)	(0.0581)	(0.0582)
Max Temp 20-30 F	0.907***	0.904***	0.906***
	(0.0345)	(0.0344)	(0.0344)
Max Temp 30-40 F	0.927***	0.924***	0.926***
	(0.024)	(0.0240)	(0.0240)
Max Temp 40-50 F	0.971	0.97	0.971
	(0.0219)	(0.0219)	(0.0218)
Max Temp 50-60 F	0.991	0.99	0.991
	(0.0154)	(0.0154)	(0.0153)
Max Temp 70-80 F	1.038**	1.037**	1.037**
	(0.0153)	(0.0150)	(0.0151)
Max Temp 80-90 F	1.052***	1.050***	1.056***
	(0.0194)	(0.0199)	(0.0195)
Max Temp 90-100 F	1.082***	1.062**	1.077***
	(0.0249)	(0.0269)	(0.0241)
Max Temp >100 F	1.300***	1.280***	1.262***
	(0.0392)	(0.0430)	(0.0442)
3 Preceding Days		1.034**	
Max Temp > 85 F		(0.0175)	
3 Preceding Days			1.063*
Max Temp > 95 F			(0.0355)
Observations	20955316	20955316	20955316
Number of counties	2661	2661	2661
County FE	YES	YES	YES
Month*Region FE	YES	YES	YES
Year FE	YES	YES	YES
Day of Week	YES	YES	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 17: Impacts of Lagged Daily Temperature, Poisson

VARIABLES	1-Day Lags TS Accident IRR	2-Day Lags
Max Temp <0 F	0.670	0.677
	(0.214)	(0.217)
Max Temp 0-10 F	0.597***	0.600***
	(0.114)	(0.115)
Max Temp 10-20 F	0.732***	0.733***
	(0.0658)	(0.0659)
Max Temp 20-30 F	0.852***	0.855***
	(0.0438)	(0.0437)
Max Temp 30-40 F	0.908***	0.909***
	(0.0327)	(0.0325)
Max Temp 40-50 F	0.971	0.968
	(0.0232)	(0.0234)
Max Temp 50-60 F	1.001	1.002
	(0.0175)	(0.0175)
Max Temp 70-80 F	1.037**	1.036**
	(0.0174)	(0.0173)
Max Temp 80-90 F	1.043**	1.044**
	(0.0224)	(0.0224)
Max Temp 90-100 F	1.026	1.028
	(0.0346)	(0.0353)
Max Temp >100 F	1.195***	1.198***
	(0.0562)	(0.0577)
Lag1 Max Temp <0F	1.487	1.616
	(0.474)	(0.629)
Lag1 Max Temp 0-10F	1.325*	1.390**
	(0.195)	(0.220)
Lag1 Max Temp 10-20F	1.035	1.067
	(0.0903)	(0.102)
Lag1 Max Temp 20-30F	1.067	1.093
	(0.0546)	(0.0689)
Lag1 Max Temp 30-40F	1.014	1.007
	(0.0364)	(0.0439)
Lag1 Max Temp 40-50F	1.005	0.987
	(0.0264)	(0.0309)
Lag1 Max Temp 50-60F	0.968**	0.969*
	(0.0153)	(0.0170)
Lag1 Max Temp 70-80F	0.994	0.994
	(0.0165)	(0.0186)
Lag1 Max Temp 80-90F	1.002	1.016
	(0.0201)	(0.0250)
Lag1 Max Temp 90-100F	1.070**	1.096**
	(0.0353)	(0.0397)

Lag1 Max Temp >100F	1.102*	1.137**
	(0.0569)	(0.0648)
Lag2 Max Temp <0F		0.854
		(0.345)
Lag2 Max Temp 0-10F		0.917
		(0.151)
Lag2 Max Temp 10-20F		0.995
		(0.0892)
Lag2 Max Temp 20-30F		0.934
		(0.0471)
Lag2 Max Temp 30-40F		1.013
		(0.0354)
Lag2 Max Temp 40-50F		1.042
		(0.0262)
Lag2 Max Temp 50-60F		0.982
		(0.0154)
Lag2 Max Temp 70-80F		0.998
		(0.0162)
Lag2 Max Temp 80-90F		0.974
		(0.0209)
Lag2 Max Temp 90-100F		0.961
		(0.0305)
Lag2 Max Temp >100F		0.953
		(0.0445)
Observations	20,955,316	20,955,316
Number of counties	2661	2661
County FE	YES	YES
Month*Region FE	YES	YES
Year FE	YES	YES
Day of Week	YES	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 18: County temp percentiles, Poisson

VARIABLES	HS Accidents
<1 Temp %	0.861**
	(0.0521)
1-5 Temp %	0.855***
	(0.0284)
5-10 Temp %	0.957
	(0.0264)
10-15 Temp %	0.967
	(0.0251)
15-25 Temp %	0.940***
	(0.0206)
25-35 Temp %	0.973
	(0.0190)
35-45 Temp %	0.967*
	(0.0179)
55-65 Temp %	0.994
	(0.0186)
65-75 Temp %	1.016
	(0.0221)
75-85 Temp %	1.037
	(0.0251)
85-90 Temp %	1.032
	(0.0262)
90-95 Temp %	1.071**
	(0.0289)
95-99 Temp %	1.043
	(0.0314)
>99 Temp %	1.224***
	(0.0574)
Observations	19,917,285
Number of counties	2,647
Standard FE	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 19: Temperature shocks analysis, OLS and Poisson

	(1) Primary Poisson	(2) Poisson w/ Temperature Shocks	(3) OLS w/ Temperature Shocks
VARIABLES	TS Accident IRR		TS Accident Count
Max Temp <0 F	1.025	1.058	0.000945***
	(0.269)	(0.280)	(0.000224)
Max Temp 0-10 F	0.755*	0.786	0.000626***
	(0.116)	(0.123)	(0.000175)
Max Temp 10-20 F	0.824***	0.855**	0.000557***
	(0.0582)	(0.0616)	(0.000151)
Max Temp 20-30 F	0.907***	0.935*	0.000557***
	(0.0345)	(0.0370)	(0.000138)
Max Temp 30-40 F	0.927***	0.950*	0.000428***
	(0.024)	(0.0252)	(0.000115)
Max Temp 70-80 F	1.036**	1.026*	0.000134***
	(0.0150)	(0.0151)	(0.0000484)
Max Temp 80-90 F	1.055***	1.036*	0.000199***
	(0.0195)	(0.0201)	(0.0000613)
Max Temp 90-100 F	1.083***	1.057**	0.000305***
	(0.0245)	(0.0260)	(0.0000970)
Max Temp >100 F	1.310***	1.269***	0.00138***
	(0.0402)	(0.0423)	(0.000328)
Diff 7-Day Avg ¡-25		1.099	-0.000126
		(0.103)	(0.000104)
-25≤ Diff 7-Day Avg <-15		0.973	-9.43e-05*
		(0.0290)	(0.0000503)
-15leq Diff 7-Day Avg <-5		0.995	-0.0000481
		(0.0110)	(0.0000347)
5≤ Diff 7-Day Avg <15		1.036***	0.000117***
		(0.0124)	(0.0000340)
15≤ Diff 7-Day Avg <25		1.036	0.000151***
		(0.0332)	(0.0000542)
25≤ Diff 7-Day Avg		1.117	0.000119
		(0.124)	(0.000138)
Observations	20,955,316	20,955,316	23,614,877
Standard FE	YES	YES	YES

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 20: Impacts of Temperature on Accidents by Time Period

	(1) 1989-1996	(2) 1996-2003	(3) 2003-2010
VARIABLES		HS Accident IRR	
Max Temp <0 F	0.736	1.391	1.101
	(0.317)	(0.530)	(0.633)
Max Temp 0-10 F	0.978	0.654	0.563*
	(0.242)	(0.181)	(0.183)
Max Temp 10-20 F	0.896	0.811*	0.840
	(0.121)	(0.0987)	(0.0978)
Max Temp 20-30 F	0.862**	0.883*	0.919
	(0.0615)	(0.0568)	(0.0599)
Max Temp 30-40 F	0.922*	0.916**	0.974
	(0.0433)	(0.0381)	(0.0414)
Max Temp 40-50 F	0.991	1.007	0.950
	(0.0383)	(0.0321)	(0.0346)
Max Temp 50-60 F	1.006	0.972	1.008
	(0.0269)	(0.0294)	(0.0257)
Max Temp 70-80 F	1.007	1.045*	1.044*
	(0.0248)	(0.0261)	(0.0269)
Max Temp 80-90 F	1.012	1.065**	1.053
	(0.0324)	(0.0303)	(0.0343)
Max Temp 90-100 F	1.006	1.113***	1.087**
	(0.0412)	(0.0424)	(0.0393)
Max Temp >100 F	1.173*	1.289***	1.291***
	(0.111)	(0.0839)	(0.0769)
Observations	5,165,613	5,196,119	5,215,335
Number of counties	2,050	2,058	2,066

SEs clustered at county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1