

NBER WORKING PAPER SERIES

KILLER CONGESTION:  
TEMPERATURE, HEALTHCARE UTILIZATION AND PATIENT OUTCOMES

Sandra Aguilar-Gomez  
Joshua S. Graff Zivin  
Matthew J. Neidell

Working Paper 33491  
<http://www.nber.org/papers/w33491>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
February 2025

We are grateful to conference participants at PacDev 2024 at Stanford University, RIDGE Forum at Universidad de Chile, AERE 2024 Summer Conference, LSE-Imperial Workshop in Environmental Economics and the Sustainable Development Research Symposium at Columbia University, as well as to seminar participants at the University of Pennsylvania, London School of Economics, University of Virginia, Hunter College, Universidad de los Andes, Tecnológico de Monterrey, ITAM and Colegio de Mexico. Valentina Castilla-Gutierrez provided outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Sandra Aguilar-Gomez, Joshua S. Graff Zivin, and Matthew J. Neidell. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Killer Congestion: Temperature, Healthcare Utilization and Patient Outcomes  
Sandra Aguilar-Gomez, Joshua S. Graff Zivin, and Matthew J. Neidell  
NBER Working Paper No. 33491  
February 2025  
JEL No. I15, I18, Q54

### **ABSTRACT**

Extreme heat imperils health and results in more emergency department (ED) visits and hospitalizations. Since temperature affects many individuals within a region simultaneously, these health impacts could lead to surges in healthcare demand that generate hospital congestion. Climate change will only exacerbate these challenges. In this paper, we provide the first estimates of the health impacts from extreme heat that unpacks the direct effects from the indirect ones that arise due to hospital congestion. Using data from Mexico's largest healthcare subsystem, we find that ED visits rise by 7.5% and hospitalizations by 4% given daily maximum temperatures above 34 C. As a result, more (and sicker) ED patients are discharged home, and deaths within the hospital increase. While some of those hospital deaths can be directly attributed to extreme heat, our analysis suggests that approximately over half of these excess deaths can be viewed as spillover impacts due to hospital congestion. Additional analyses also reveal an increase in the share of deaths occurring outside hospitals, consistent with congestion-related health harms arising from the discharge of sicker patients from the ED. Our findings highlight an important new avenue of adaptation to climate change. If hospital congestion contributes to excess health damages from a changing climate, then expanding labor and capital investments and improving surge management tools can help reduce those damages.

Sandra Aguilar-Gomez  
Universidad de los Andes  
s.aguilargomez@uniandes.edu.co

Joshua S. Graff Zivin  
University of California, San Diego  
9500 Gilman Drive, MC 0519  
La Jolla, CA 92093-0519  
and NBER  
jgraffzivin@ucsd.edu

Matthew J. Neidell  
Department of Health Policy and Management  
Columbia University  
722 W 168th Street, 6th Floor  
New York, NY 10032  
and NBER  
mn2191@columbia.edu

# Introduction

Extreme heat imperils health, increasing both morbidity and mortality (Deschenes, 2014; Basu, 2009), and results in more emergency room visits and hospitalizations (e.g., Gould *et al.*, 2024; White, 2017; Fritz, 2022). Since temperature impacts many individuals within a region simultaneously, these health impacts could lead to surges in healthcare demand that generate hospital congestion. Hospital capacity constraints may lead to hospitals turning away some patients that require care and may also generate spillover effects for other patients within the healthcare system given the need to spread scarce resources across a larger number of patients. In this paper, we provide the first exploration of the health impacts from extreme heat that unpacks the direct from the indirect effects that arise due to hospital congestion.

Understanding these impacts matters because, unlike idiosyncratic shocks to healthcare demand, we know that climate change will make those shocks associated with extreme heat far more common. Indeed, health-related costs are already estimated to be the largest component of the economic consequences of climate change (Carleton and Greenstone, 2022), and spillover effects will only make those numbers larger. Moreover, the existence of spillover effects is likely to alter our understanding of the distributional impacts of extreme heat and climate change. Those groups directly susceptible to the harms of extreme heat may look very different from the groups harmed by diminished care in other parts of the healthcare system.

This discovery also highlights a potentially important new avenue of adaptation to climate change. If hospital congestion contributes to excess health damages from a changing climate, then expanding labor and capital investments and improving surge management tools can help reduce those damages. The benefits and costs of these investments relative to the protective effect of residential air conditioning (Barreca *et al.*, 2016) will clearly depend on the context. Expanding healthcare infrastructure may be a promising climate adaptation strategy in developing countries, where residential air conditioning is much less prevalent and electricity reliability is a concern. Furthermore,

health system expansions may offer considerable co-benefits in settings where healthcare demand far outstrips supply.

Our empirical work focuses on Mexico, and particularly on those persons with public insurance (Seguro Popular, henceforth SP) and the uninsured. Studying this setting is relevant, since evidence on the relationship between extreme heat and health in developing countries is much more limited (Sapari *et al.*, 2024). Moreover, like many other developing countries, health staffing and infrastructure in Mexico are quite constrained, particularly in the facilities where patients with SP seek care.<sup>1</sup> Furthermore, the incidence of extreme heat days under climate change in Mexico is likely to be disproportionately high relative to the number in countries in higher latitudes (Pörtner *et al.*, 2022; Murray-Tortarolo, 2021), underscoring the potential welfare and policy implications of our findings.

Our primary analysis covers nearly 57% of the Mexican population over an 8-year period. In particular, we employ data on the universe of 2012–2019 emergency department (ED) and hospital visits to Mexico’s Ministry of Health (MoH) hospitals, which serve patients with public insurance under SP and uninsured patients. We use these data to examine the impact of temperature on hospital crowding, patient transitions within the healthcare system, and patient outcomes. Vital statistics data for the entire population supplement our core analyses, allowing us to examine mortality outside of the hospital. Using the exact date from patients’ admissions and discharge records and the municipality of each facility, we link health data to weather data to estimate distributed lag models of temperature impacts. Our core econometric analysis relies on a nonparametric binned model specification for temperature that includes facility fixed effects as well as fixed effects to capture seasonality, intraweek patterns in admissions, and municipality-specific trends.

Our analyses reveal a series of interesting, interconnected results. To begin, we

---

<sup>1</sup>Mexico has a multi-tiered healthcare system where people seek care in different facilities based on their insurance coverage. SP, currently under restructuring, was the largest healthcare provider in the country during our study period.

find that higher temperatures increase ED and hospital admissions. When the daily maximum temperature reaches the highest bin of  $>34^{\circ}\text{C}$ , we estimate an additional 3 ED visits compared to the number under  $22\text{--}24^{\circ}\text{C}$ . From a mean of approximately 40 daily visits, this translates to a 7.5% increase. The corresponding figure for hospitalizations is 0.5 additional admissions, an approximately 4% increase. These results are consistent with prior literature in different contexts (e.g., White, 2017).

The gap between ED and hospital admissions motivates a deeper investigation of transitions within the healthcare system. Our examination of ED discharges is illuminating in this regard. While the absolute number of patients admitted to the hospital from the ED increases, that increase is small, such that the probability that any given patient in the ED is admitted to the hospital decreases. It appears that direct hospital admissions due to extreme heat create congestion within the hospital that limits its ability to admit ED transfers. This congestion, in turn, appears to push the ED toward its capacity constraints, leading to more patients being sent home on extreme heat days. A similar phenomenon happens within the hospital, with more patients being discharged on hot days, suggesting that hospitals respond to this congestion by decreasing the length of stay.

The astute reader will note our choice of the word ‘seems’ when offering congestion and capacity constraints as the mechanism driving our findings. While our findings thus far are consistent with congestion, this pattern of results could also arise due to changes in the composition of patients on extreme heat days. If, for example, less-sick patients show up to the ED on hot days, it is perfectly reasonable for EDs to admit fewer patients to the hospital and to send more home. Fortunately, we can examine this directly based on measures for the severity of illness (Hoe, 2022). In fact, we find that the severity of illness in the ED increases with temperature. Furthermore, those sent home on extreme heat days are significantly less healthy than those being discharged on cooler days. By contrast, while those admitted to the hospital from the ED are among the sickest patients in that department, those transferred to the hospital are no more

infirm than those already admitted there. Congestion is the most plausible explanation for these results.

The final piece of the puzzle is estimating the impacts of this congestion. Inside the hospital, this estimation is relatively straightforward. Since the composition of illness severity does not change, we can simply look at excess mortality for hospitalized patients. We find that experiencing one day with maximum temperature exceeding  $34^{\circ}\text{C}$  leads to a 5% increase in excess deaths. These deaths include the direct physiological impacts of heat as well as any indirect effects arising from hospital congestion. Separating these channels is trickier than it might sound because heat can impact many conditions often not classified as heat-related, ranging from accidents to heart attacks (Park and Pankratz; Drescher and Janzen, 2024).<sup>2</sup> To parse these indirect effects, we examine the impact of heat shocks on patients already admitted to the hospital (conditional on heat when admitted). Our analysis suggests that one day over  $34^{\circ}\text{C}$  also increases mortality for this population by 5%, revealing large spillover effects. As an additional test, we limit our attention to cancer patients, who should not be impacted by temperature conditional on hospital admission, and find a similar increase, underscoring the spillover effects generated by heat-driven overcrowding in the hospital.

Assessing the health consequences for those sent home from the ED is more challenging, as we cannot follow patients once they leave the healthcare system. We can, however, use vital statistics data to measure the number of deaths that occur on hot days in specific settings. Although we find that deaths increase both inside and outside of hospitals, we see a larger relative increase in deaths outside the hospital. While we cannot know whether these deaths occur among patients sent home from the ED or those discharged early from inpatient care, this result is consistent with crowding impacts: more severe patients are sent home from the ED, those admitted into inpatient care have shorter stays, and deaths disproportionately increase outside the hospital sys-

---

<sup>2</sup>Indeed the public health literature emphasizes that there are no standardized protocols for reporting heat-related illness (Vaidyanathan *et al.*, 2019), and heat stress can exacerbate preexisting health conditions, often resulting in conditions being recorded under different diagnoses like renal failure or cardiovascular events (Schulte and Chun, 2009; Bell *et al.*, 2016).

tem. Taken together, these results suggest that the additional crowding caused by higher temperatures leads to increased mortality.

Our paper builds upon the extensive literature focused on the health impacts of climate change (see reviews by Basu, 2009; Ye *et al.*, 2012; Bunker *et al.*, 2016) to reveal a new channel through which those impacts may arise. In so doing, it also illuminates a new approach to climate adaptation. While several studies have explored the adaptive role of air conditioning (see reviews by Kahn, 2016; Fankhauser, 2017; Owen, 2020; Deschênes and Greenstone, 2011; Barreca *et al.*, 2016), our study highlights investments in healthcare infrastructure and patient management as another tool for adapting to warmer temperatures. Our paper is also closely related to recent studies that examine the impact of non-heat-related surges in hospital demand on healthcare quality. For example, Hoe (2022), who focuses on external injuries, finds that shocks to UK hospital trauma and orthopedic departments lead to a reduced quality of care within those departments as measured by readmission rates. Gutierrez and Rubli (2021) utilize data similar to ours from Mexico and show that surges in hospital demand due to an outbreak of H1N1 influenza led to more in-hospital deaths for patients with other conditions. Lastly, Guidetti *et al.* (2024) examine the impacts of health shocks due to spikes in pollution on pediatric hospitalizations in Sao Paulo, Brazil, and find that higher levels of pollution generate congestion that impacts the health of children with respiratory conditions and decreases admissions for non-respiratory conditions. We extend this work by broadening our purview to spillovers between EDs and the entire hospital system, as well as to the health impacts on those outside the hospital system.

# 1 Background and data

## 1.1 The Mexican healthcare system

The Mexican healthcare system is fragmented and multi-tiered.<sup>3</sup> People obtain insurance through the private sector ( $\sim 2\text{--}3\%$  of the population), social insurance from the Mexican Social Security Institute (IMSS) and the Institute for Social Security and Services for State Employees (ISSSTE) ( $\sim 40\text{--}45\%$ ), and the safety net SP program ( $\sim 40\text{--}45\%$ ).<sup>4</sup> Approximately 14% of the population is uninsured, although they are technically eligible for SP but have not enrolled. IMSS and ISSSTE are for formally employed individuals in the private and public sectors, respectively, with funding provided by contributions from employees and employers along with government subsidies. SP is safety net insurance for those who are informally employed or unemployed, with funding provided entirely by the government.

In line with the different tiers of insurance, healthcare services are highly fragmented. Privately insured individuals seek care in a private and largely unregulated market. IMSS and ISSSTE "organize, provide and regulate most of their own health services through vertically integrated, national organizations" (Block *et al.*, 2020). Those with SP or without insurance obtain care from facilities managed directly by the MoH.<sup>5</sup>

Resources at Mexican healthcare facilities are highly constrained, especially in more rural areas and for public facilities. Overall, Mexico spends 5.5% of its GDP on healthcare, compared to an OECD average of 9.2%. Half of this expenditure is out-of-pocket even though fewer than 2% of the population has private insurance. Mexico also faces a notable deficit in its medical workforce, with a ratio of 2.5 physicians per 1,000 inhabitants, which is considerably lower than the OECD's recommended threshold of 3.2. This shortfall extends to nursing personnel, where shortages are even more significant:

---

<sup>3</sup>For more detail, see Block *et al.* (2020).

<sup>4</sup>SP was superseded by the Institute of Health for Welfare (INSABI) in 2020. Our analysis, however, ends in 2019 to avoid confounding due to COVID-19, so we refer to SP throughout.

<sup>5</sup>Care at EDs, however, is not fragmented. Patients can freely enter any ED regardless of insurance status.



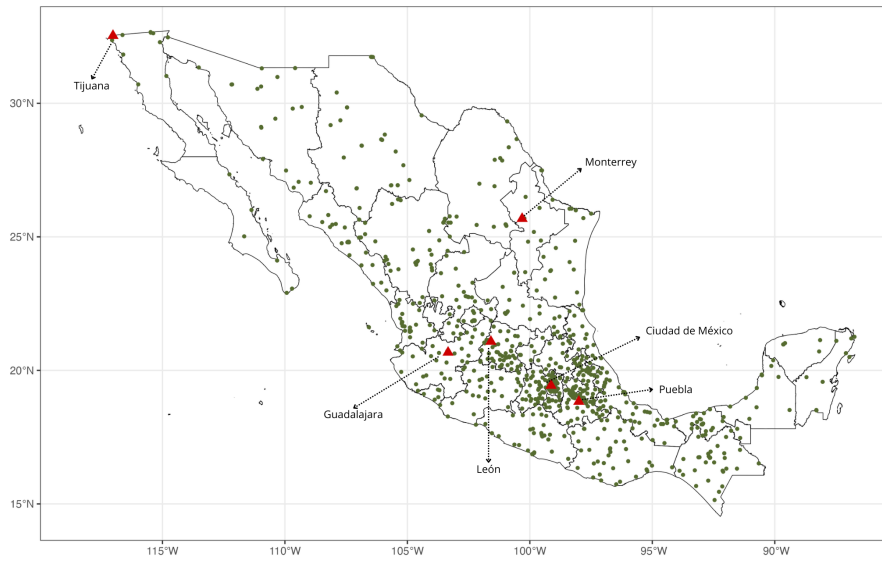
only 2.9 nurses per 1,000 people, significantly below the OECD’s recommendation of 8.8 (OECD, 2023). Access to medical technology is also limited. For instance, Mexico has only 2.6 MRI machines per million inhabitants, compared to an OECD average of 15.6 (Block *et al.*, 2020). As a result, long wait times and a lack of medical supplies are common concerns, and any patient surge is likely to further stress an already constrained healthcare system. It is worth noting that Mexico is quite representative of the type of healthcare access constraints that inhabitants of most countries face. For example, Mexico’s density of doctors per capita ranks 47 out of 96 countries with available data (OWID, 2021).

## 1.2 Data

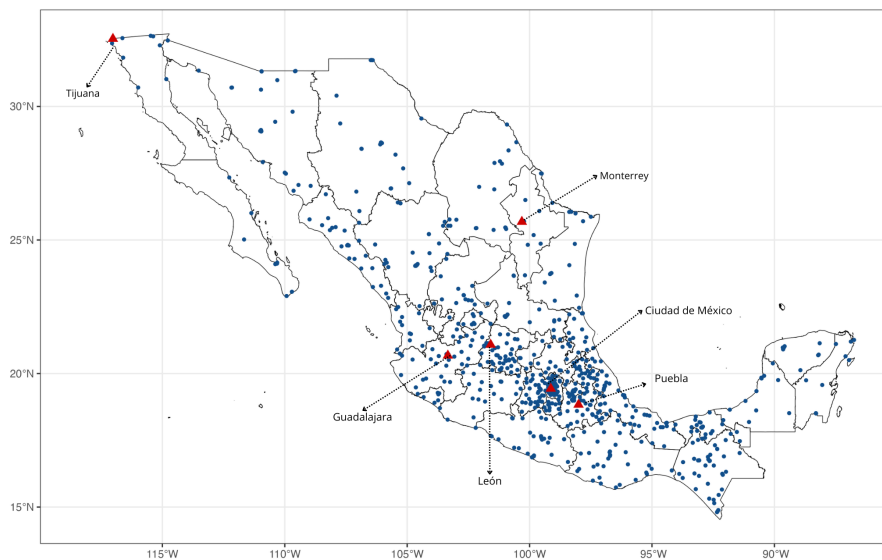
To evaluate the relationship between temperature, healthcare usage, and health outcomes, we combine public hospital administrative records from the MoH and mortality records for the entire population with weather data for the years 2012–2019. Our use of MoH-managed hospitals, which provide healthcare to the uninsured and those enrolled in SP, concentrates our focus on healthcare usage by the lower and middle part of the income distribution. The mortality records come from death certificates and thus reflect the entire population.

### Healthcare utilization and mortality data

Our study draws on patient-level records from the complete network of hospitals in Mexico’s non-contributory public healthcare system, overseen by the MoH, for the years 2012–2019. Although many hospitals in our sample operate an ED that subsequently admits patients for inpatient care, data for these two care phases are recorded separately by distinct information subsystems: the Emergency Department Subsystem and the Hospital Discharge Subsystem. The data come from 916 EDs and 857 hospitals and are plotted in Figure 1; these data account for the healthcare of 57% of the national population.



(a) Emergency departments



(b) Hospitals with inpatient services

Figure 1: Distribution of hospitals

*Notes:* This figure illustrates the distribution of hospitals in our sample. Top: Data from 916 EDs in Mexico for 2019. Bottom: Data from 857 hospitals with inpatient services in Mexico for 2019

In both subsystems, detailed diagnostic information is recorded for each patient, including a primary diagnosis and up to 6 secondary diagnoses. Specialized coders are tasked with converting handwritten diagnostic notes into standardized 4-digit codes based on the International Classification of Diseases, Tenth Revision (ICD-10). ICD-10

comprises over 70,000 possible diagnosis codes, but the Mexican health system utilizes only the first 4 digits, delineating 9,795 distinct categories.

Although we have detailed information on the patient’s diagnosis, we primarily focus our analysis on all ED and hospital visits regardless of cause (we rely on diagnosis to define illness severity as described below). The reason for this choice is twofold. First, it is difficult to identify specific illnesses related to heat. Very few patients are coded as having heat-related illnesses, such as heat stroke or exhaustion.<sup>6</sup> Instead, heat often exacerbates conditions that do not have a heat-related code. For example, heat stresses the cardiovascular system by making the heart work harder, which increases the risk of a heart attack (Sun *et al.*, 2018; Ebi *et al.*, 2021). Epidemiological studies have shown that heat increases ED and hospital visits for conditions related to the heart (Wettstein *et al.*, 2018), kidney (Hansen *et al.*, 2008), lungs (Michelozzi *et al.*, 2009), mental health (Liu *et al.*, 2021), and more (Ebi *et al.*, 2021). Furthermore, as shown in Appendix Figure A.2, excluding patients with explicitly heat-related codes results in nearly identical results. Second, as discussed earlier, our goal is to estimate not only the direct impacts of heat but also the spillover effects that arise through resource constraints. The presence of these spillover effects suggest heat could impact any condition sensitive to resource constraints. By considering all admissions regardless of diagnosis, we capture both the direct and the indirect spillover effects.

Both datasets record the discharge status of each patient, which we use to explore how patients flow through the healthcare system. For the ED, the following five options exist: 1) sent home; 2) hospitalized; 3) referred to another hospital; 4) exit without discharge; and 5) deceased.<sup>7</sup> The following three discharge options exist for the hospitalized: 1) sent home, 2) transferred to another facility, and 3) deceased. The data, however, are not linked, so we are unable to analyze the mortality outcomes from ED admissions.

---

<sup>6</sup>Monitoring of heat-related illness through specific diagnosis codes has been highlighted as a global problem (Harduar Morano and Watkins, 2017; Fox *et al.*, 2019).

<sup>7</sup>For administrative reasons, very few patients die in the ED. They are transferred to the hospital before being pronounced dead.

The municipality of each facility and the date of admission are reported in both sets of data, which enables us to merge the weather data. We aggregate individual visits, which consist of approximately 70 million ED and 22 million hospital visits during our study period,<sup>8</sup> to the facility-day level. This results in 1,865,622 ED-days and 2,141,109 hospital-days.

**Severity of illness.** We use the hospital data to calculate a measure of illness severity to understand the composition of patients admitted to the ED and hospital on any given date. Following Hoe (2022), we leverage the availability of the universe of deaths and admissions within the SP over almost a decade and define severity  $\hat{s}$  as the estimated in-hospital mortality rate at the national level for each diagnosis code, conditional on age and sex.<sup>9</sup> Since there are nearly 10,000 4-digit ICD-10 codes, we calculate this rate based on groupings of the first 3 digits of the ICD-10 codes, following the disease categories listed in the ICD-10 manual. For example, codes A00-A09 indicate "intestinal infectious disease," A15-A19 indicate "tuberculosis," and A20-A28 indicate "certain zoonotic bacterial diseases." This results in a total of 290 disease categories for which we calculate mean mortality rates. We assign this severity measure to each patient based on age, sex, and diagnosis code determined at admission. We aggregate these data to the hospital-date level by taking an average of patients admitted for each facility and date in our data frame.

**Excess deaths.** To investigate the contribution of heat to additional deaths in the hospital, we employ an excess mortality measure comparable to the risk-adjusted mortality measure monitored by the NHS (NHS, 2024). Specifically, we define excess mortality as  $\min\{0, \mathbf{1}(m) - \hat{s}\}$ , where  $\mathbf{1}(m) = 1$  if the patient died and 0 otherwise, and  $\hat{s}$  is the severity assigned to each patient as defined above. In this definition, excess mortality equals zero for patients who do not die in the hospital and  $1 - \hat{s}$  for patients who do. As

---

<sup>8</sup>The ED visits are split evenly between SP and uninsured patients, while approximately 75% of the hospital visits are patients with SP.

<sup>9</sup>More formally, we perform a Poisson regression of the probability of death on the diagnosis categories, with age and sex fixed effects. Severity is the predicted likelihood of death estimated using the coefficients from this regression.

with the ED and hospital data, we aggregate excess deaths to the hospital-date level.<sup>10</sup>

## **Mortality data**

We utilize a comprehensive dataset comprising all death certificates in the country from 2012 to 2019, regardless of insurance status. We continue to use all deaths regardless of cause in order to capture the direct and indirect effects of heat. These records provide the date and location of each death, enabling us to link these observations to the relevant weather data. Furthermore, they provide the setting of each death, such as a private hospital, a public hospital, a public space, or a residence, which we will use to analyze events that happen outside of the healthcare system. This enables us to explore the impacts of heat on mortality outcomes regardless of the setting, capturing patients who may be sent home from the ED or hospital but still succumb to their illness.

## **Weather data**

For the weather data, we use the Daymet reanalysis data product from NASA’s Oak Ridge National Laboratory (ORNL) (Thornton *et al.*, 2021). These data are based on statistical models that interpolate and extrapolate from weather stations, thus enabling weather calculations in areas where stations do not exist. Daymet data include daily measures of minimum and maximum temperature on a 1km x 1km gridded surface. We average all grids within a municipality on a daily basis to obtain the average daily maximum temperature and assign this to each ED and hospital at the daily-municipality level. Figure 2 (a) displays the distribution of daily maximum temperatures for hospitals at the 10<sup>th</sup>, 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup> and 90<sup>th</sup> percentiles of the historical average maximum temperature. Most hospitals experience daily maximum temperatures between 10 and 40 degrees, with considerable variation in median temperatures and year-round variability across percentiles.

---

<sup>10</sup>The risk-adjusted mortality rate (NHS, 2024), calculated as the ratio between the actual and the expected rate given patient case mix as a measure of excess deaths, is designed to capture annual variation. On a given year, there are no hospitals with zero expected deaths, which makes a ratio an adequate measure. To capture daily variation, a ratio would generate numerous missing values because on many days, there are zero expected deaths given patient composition. For this reason, we measure this gap as a subtraction rather than as a ratio.

### 1.3 Descriptive statistics

Summary statistics for ED visits are shown in the first panel of Table A.1. The average number of ED visits is 39.65 daily, but considerable variability exists, with a standard deviation of 40.14. This variability is partly due to changes in visits over time within an ED, as reflected by the standard deviation within EDs of 7.02, but largely due to the sample consisting of EDs of different sizes. This reflects the diverse contexts in which the Mexican healthcare system operates and the different capacities of these facilities. The nature of this variation is highlighted in Figure 2 (b), which shows the distribution of daily ED visits separately for EDs at the 25th, 50th and 75th percentile of the mean daily admissions.

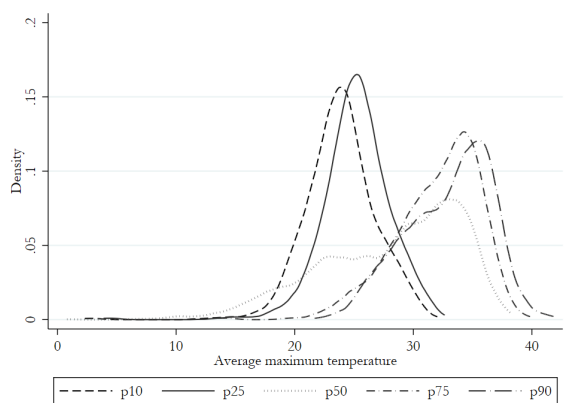
In terms of discharge status, the majority of ED patients (73.43%) are sent home, followed by hospitalization (12.52%) and referral to another hospital (10.77%). Very few patients exit without an official discharge (0.47%) (0.23%) or have missing discharge information. The mean severity of ED patients is 0.02, indicating that the severity of illness for the mix of diagnosis, age and sex of patients admitted to the ED corresponds with a 2% mortality rate for that same mix in the hospital. Appendix Figure A.1 shows that while there are some outliers, the severity for most ED patients ranges between 0 and 0.06.

Shown in the second panel of Table ?? are descriptive statistics on hospital admissions. Admissions average 12.37 per day, again with considerable variability driven by daily factors as well as hospital features, with the third panel of Figure ?? reflecting the latter. Patients stay, on average, 3.5 days in the hospital. When they leave, most patients are deemed sufficiently healthy to be sent home (92.79%), followed by transferred to another facility (2.44%) and deceased (1.61%).<sup>11</sup> The severity score for the inpatient services is 0.02 (standard deviation = 0.03), similar to those in the ED, with an excess mortality average of 0.23 (standard deviation = 0.63). Appendix Figure A.1 shows that the severity for most patients ranges between 0 and 0.1.

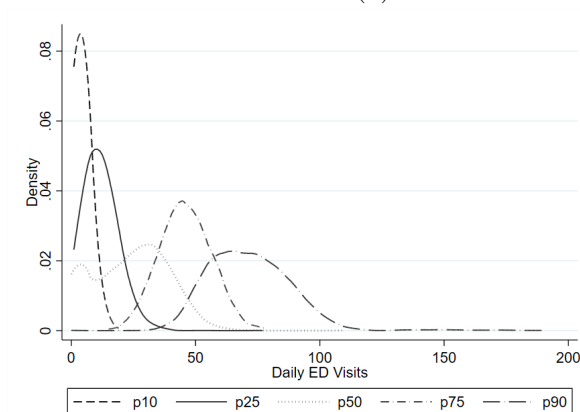
---

<sup>11</sup>The remaining patients leave without a formal discharge.

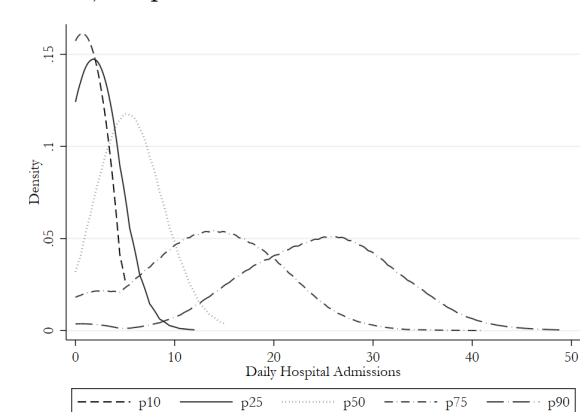
The county-level mortality data indicate a mean of 3.02 deaths per day (standard deviation = 5.06). These deaths are distributed across various settings, with 23.17% occurring in public hospitals, 1% in private hospitals, and the rest occurring outside of hospitals: at home, in a public space or at an unknown location.



(a) Maximum temperature, hospitals



(b) Daily ED visits



(c) Daily hospital admissions

Figure 2: Descriptive statistics

*Notes:* Data from 916 EDs and 857 hospitals with inpatient services in Mexico between 2012 and 2019. Unit of observation: hospital-day, with cohorts based on the date of admission of the patient. Panels a) and b) depict the 10th, 25th, 50th, 75th, and 90th percentile of the distribution of total patient influx for ED visits and hospitalizations, respectively. Panel c) presents the distribution of the daily maximum temperature for the 10th, 25th, 50th, 75th, and 90th percentile of the distribution of average annual temperature for the period of observation.

## 2 Empirical methodology

The initial focus of our empirical strategy is to estimate the causal effect of the temperature realized on a given day on ED and hospital visits in a particular location. Given

that daily ED and hospital visits are count data, we estimate a Poisson pseudo-maximum likelihood model (PPML). The PPML point estimates are consistent as long as the conditional mean is correctly specified, irrespective of the distribution of the outcome or errors (Gourieroux *et al.*, 1984). PPML models also readily accommodate fixed effects without an incidental parameters problem (Correia *et al.*, 2019). Furthermore, the PPML estimator performs well with a large number of zeros and over- or under-dispersion in the data (Silva and Tenreyro, 2011).<sup>12</sup>

Our baseline model for ED and hospital visits is specified according to the conditional exponential mean function in equation 1, where  $Y$  is the number of ED or hospital visits on date  $d$  in hospital  $h$  in municipality  $m$ :

$$E[Y_{hcd} | \sum_{l=0}^7 (f(t_{\max_{hcdl}})), \rho_h, \Omega_d] = \exp \left( \sum_{l=0}^7 \sum_{b=1}^9 \beta_b \mathbb{1}(t_{\max_{hcdl}} \in (\underline{t}_b, \bar{t}_b]) + \rho_h + \Omega_d + \varepsilon_{hcd} \right) \quad (1)$$

The key variable of interest is  $t_{\max}$ , which measures the average daily maximum temperature in a municipality.<sup>13</sup> In line with previous work, we allow for a nonlinear effect of temperature by using a series of indicator variables for each 2°C, with 22–24°C serving as the reference. Thus, we interpret the coefficient  $\beta$  of bin  $b$  as the contemporaneous impact of the maximum temperature in that bin relative to 22–24°C. This allows each 2°C bin to have an independent impact on outcomes. Based on the distribution of temperature, we specify  $\leq 18^\circ\text{C}$  as the lowest bin and  $\geq 34^\circ\text{C}$  as the highest. Since temperature impacts may arrive with delay, we also include 7-day lags

<sup>12</sup>Our results are robust to using a linear fixed effects regression model.

<sup>13</sup>In Section A.2, we show that our results are robust to using wet-bulb temperature, a measure that attempts to account more fully for heat stress. The calculations of wet bulb temperature are based on the variables of temperature (degrees C) and water vapor pressure (Pa) using the PsychroLib library in Python. This library contains functions for calculating the thermodynamic properties of gas-vapor mixtures and standard atmosphere suitable for most engineering, physical and meteorological applications. Most of the functions are an implementation of the formulae found in American Society of Heating and Engineers (2017).



of these temperature bins,  $l \in 1, 7$ .<sup>14</sup> As we demonstrate below, the contemporaneous effects are considerably stronger than any lag, so we only show the full set of temperature bin coefficients for same day temperature.

Our model includes several fixed effects to aid in the identification of temperature impacts. Consistent with the standard approach in the climate economics literature (Dell *et al.*, 2014; Deschenes, 2014; Barreca *et al.*, 2016; Mullins and White, 2019), we include hospital-specific fixed effects ( $\rho_h$ ) to compare changes in temperature within an area with changes in ED or hospital visits. We include several additional fixed effects, denoted by  $\Omega_d$ , to control for other important factors affecting ED and hospital visits. These include day-of-the-year (doy) fixed effects to adjust for seasonality, municipality-by-year fixed effects to control for annual trends specific to each municipality, and day-of-the-week (dow) by year fixed effects to capture evolving intra-week patterns. The identification assumption is that after accounting for baseline factors such as annual trends and seasonality, the remaining fluctuations in daily temperature at a given facility are exogenous. The error term,  $\varepsilon_{hcd}$ , includes both an i.i.d. and a municipality level component. The municipality level component accounts for the assignment of temperature to all hospitals at that level and also allows for serial correlation within municipalities (and the hospitals nested within them). We cluster standard errors at the municipality level to allow for these features.

Although our initial focus is on healthcare utilization, the second, and more novel, focus of our empirical strategy explores transitions within the health care system and their implications for patient health. To do so, we consider several additional outcomes, such as cases transferred from the ED to the hospital and excess deaths inside the hospital, using the same econometric model but with these different dependent variables. When using vital statistics data to explore mortality for the entire nation, the model is modified to include municipality fixed effects instead of hospital fixed effects.

---

<sup>14</sup>We also employ specifications with 30-day lags of temperature and find our results do not appreciably change (results available upon request).

## 3 Results

### 3.1 Emergency department and hospital visits

The first set of results documents the relationship between temperature and either ED or hospital visits. These results set the stage for our subsequent analyses and also serve as a validation exercise, since previous studies document heat impacts on healthcare utilization.

The results show that, consistent with previous studies, higher temperatures lead to increased health care utilization. Figure 3 Panel (a) shows the results for ED visits. We find a monotonically increasing relationship between temperature and ED visits that is statistically significant for all temperature bins. When temperature reaches the highest bin of  $\geq 34^{\circ}\text{C}$ , we estimate an additional 3.01 ED visits compared to the number given a temperature of  $22\text{--}24^{\circ}\text{C}$ . From a mean of  $\sim 40$  daily visits, this translates to a 7.6% increase. The effects decrease as we move to colder temperatures, with an estimate of 2 additional visits at  $30\text{--}32^{\circ}\text{C}$  and 1 at  $26\text{--}28^{\circ}\text{C}$ . The impacts do not level out with cooler temperatures, as a decrease in visits occurs at each temperature bin below  $22\text{--}24^{\circ}\text{C}$ . At first blush, this pattern seems surprising, but it is consistent with recent evidence in California (Gould *et al.*, 2024). The most likely explanation for this pattern is that heat produces a change in health care seeking behaviors (White, 2017), in addition to its acute health impacts.

Turning to hospital visits in Panel (b), we also find statistically significant increases resulting from higher temperatures. At the highest temperature bin of  $\geq 34^{\circ}\text{C}$ , we estimate an 4.2% impact, which translates into roughly .45 additional hospital admissions. The estimates are progressively smaller as temperature decreases, with a convex shape consistent with a U-shaped relationship, noting that we only observe higher temperatures in the Mexican setting.

As previously mentioned, we include 7 lags of the temperature bins in our econometric model. To display the impact from lags 1–30, Figure 4 shows the coefficients

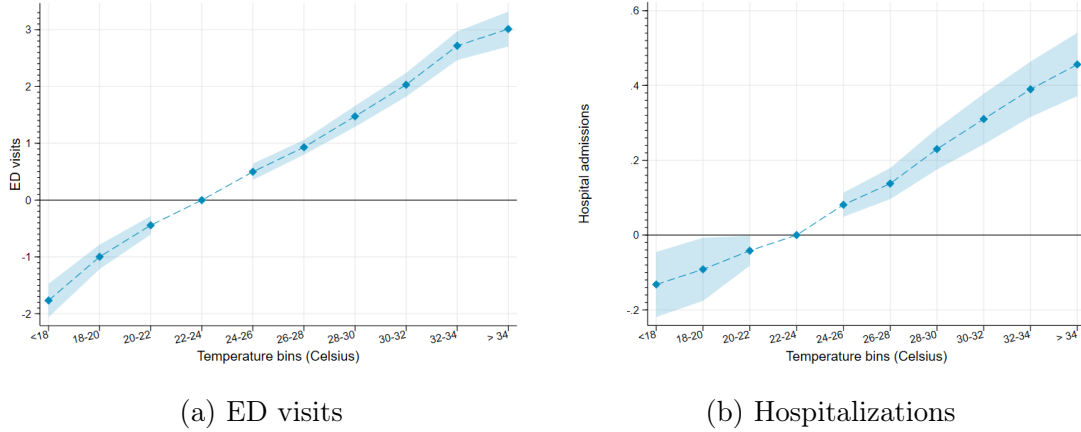


Figure 3: Temperature and number of visits per day

*Notes:* ED visits between 2012–2019. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year, day-of-the-week by year and hospital fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

for contemporaneous temperature and for each lag for the highest temperature bin only. Panel (a) shows that, for ED visits, the impact of contemporaneous temperature far outweighs that of any of the lags, although the impacts from lags 1 to 3 are statistically significant. The coefficient of 3 for same day is followed by a 0.5 coefficient for 1-day lag, a considerable drop in magnitude, and only gets smaller from there. A similar pattern arises for hospital visits, shown in Panel (b), though only the coefficient from lag 1 is statistically significant. For the remaining analyses, we continue to employ the lagged specification but only display the estimates for contemporaneous temperature given its dominance in impacting healthcare utilization.

### 3.2 Transitions within the health care system

The above results demonstrate an increase in both ED and hospital visits when temperature increases but highlight an important gap in care. Focusing solely on the highest temperature bin, an additional 3 patients visit the ED when temperatures exceed 34°C, whereas only 0.5 additional patients enter the hospital. What happens to these extra patients? We investigate this question using the discharge status of ED patients.

One possibility is that hospitals admit fewer patients from outside of the ED (e.g.

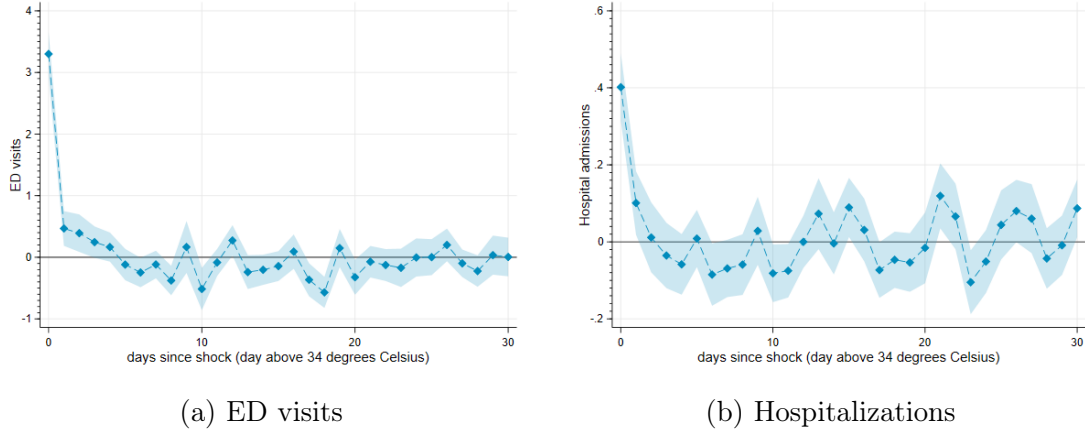


Figure 4: Temperature and number of visits per day

*Notes:* ED visits between 2012–2019. Econometric specification includes 30 lags of 9 bins; only the hottest bin ( $\geq 34^\circ\text{C}$ ) is plotted. Omitted category:  $22\text{--}24^\circ\text{C}$ . Controls include day-of-the-year, municipality-by-year, day-of-the-week by year and hospital fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

direct admissions for elective or non-urgent procedures) to free up space to transfer ED patients into the hospital. We explore this possibility by looking at the number of ED patients discharged directly to the hospital. The results, shown in Panel (a) of Figure 5, demonstrate an increase in discharges to the hospital as temperatures increase. When temperatures exceed  $34^\circ\text{C}$ , we find a statistically significant increase of 0.28 patients admitted to the hospital. This implies that part of the increase in hospital visits comes from patients admitted through the ED but that another proportion of the increase arises from direct admissions, which could be transferred from other EDs or directly admitted from consultation. Moreover, these findings still do not fully explain what happens to the additional patients admitted to the ED.

Although the number of hospital admissions from the ED increases on hot days, the probability of being hospitalized decreases, as shown in Figure 5, Panel (b). At the highest temperature bin, we estimate a .35 percentage point decrease in the probability of being admitted to the hospital. This amounts to a  $\sim 3\%$  decrease from the baseline hospitalization rate of 12% among ED patients. Panel (c) shows that more patients are discharged home as heat increases.<sup>15</sup> When temperatures exceed  $34^\circ\text{C}$ , we estimate an

<sup>15</sup>The third major discharge category is transferred, but we find no effect of temperature on this group.

increase of 0.5 percentage point in the probability of being sent home. In fact, we find that the increase in patients sent home on hotter days (Panel (b)) represents almost all of the additional ED admissions, shown in Figure 3. Together, these results suggest that hospitals are unable to accommodate the additional patients admitted to the ED in their inpatient or outpatient units and end up sending these patients home.

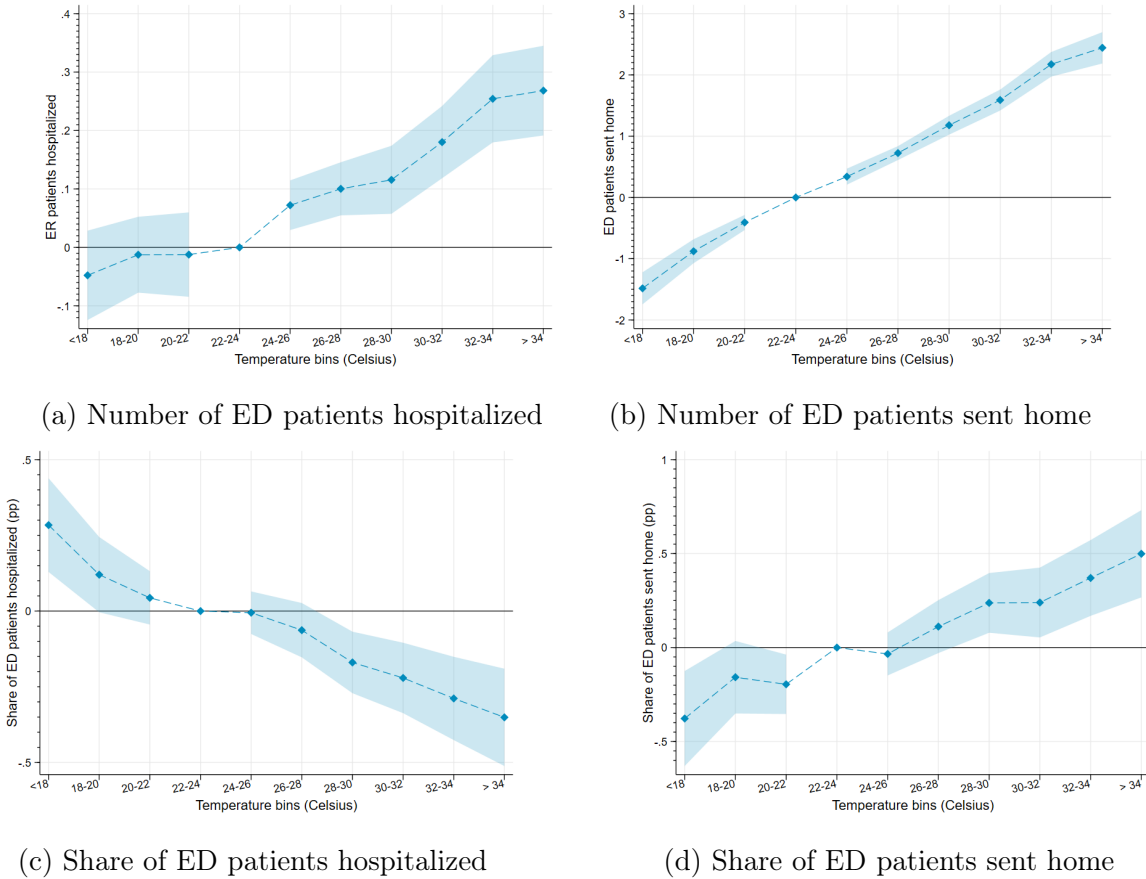


Figure 5: Changes in triage of ED patients

*Notes:* ED visits between 2012–2019. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year, day-of-the-week by year and hospital fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

A similar phenomenon happens within the hospital, with more patients being discharged on hot days, suggesting that hospitals respond to this congestion by decreasing length of stay. Figure 6 shows that the number of discharges increases by  $\sim 0.07$  patients when a day’s maximum temperature reaches above 34°C, an increase that corresponds to  $\sim 1\%$  of daily patient inflow.<sup>16</sup>

<sup>16</sup> $p < 0.1$  for the 28–30°C and 32–34°C.  $p < 0.05$  for 30–32°C.




Figure 6: Temperature and patient discharges

*Notes:* Hospitalizations between 2012–2019. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year, day-of-the-week by year and hospital fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

### 3.3 Patient composition

One explanation for the increase in patients being sent home from the emergency and inpatient departments could be patient composition. If less-sick patients show up to the ED on hot days—perhaps driven by behavioral factors<sup>17</sup>—it is perfectly reasonable for EDs to send more patients home. To explore this possibility, we estimate the impact of temperature on the severity of ED patient illness. Recall that severity is the mortality rate based on the diagnosis code, by age and sex, assigned to each patient at admission. The results, shown in Figure 7, reveal that patient severity increases with temperature. When temperature exceeds 34°C, there is a statistically significant 1% increase in patient severity.<sup>18</sup> Compared to days with milder temperatures, on average, sicker patients show up at the ED as temperatures rise.

Heterogeneity in patient severity could arise whereby more severely ill patients are admitted to the hospital and those with a less severe illness are sent home. When we specifically examine the severity for patients discharged home, shown in Panel (c),

<sup>17</sup>White (2017) finds increased treatment-seeking behavior on hotter days in EDs in California.


<sup>18</sup>We report  $\exp(\beta - 1)$  instead of the marginal effects for ease of interpretation for severity and excess mortality results, as the original units are less intuitive. This transformation represents the percent change in the severity of excess mortality for the maximum temperature falling in a specific bin relative to that for the 22–24°C bin.

however, we produce nearly the same pattern of results. The average patient sent home (instead of being admitted) is also sicker on hotter days. This pattern supports a “quicker and sicker” scenario in which patients are discharged more rapidly despite being sicker, a finding consistent with constrained hospital resources. We also explore the severity of illness among patients admitted to the hospital, shown in Panel (b). We find that patient severity is unrelated to heat, with no clear pattern of results. Combined with the previous results, this suggests that the marginal patient transferred from the ED to the hospital is a more severe case on a hotter day than the average ED patient on a cooler day, but comparable to the average severity for patients in the hospital. It also indicates that the increasing admissions to hospitals on hot days consists of patients with a similar health status, thus placing significant strain on the hospital system.

### 3.4 Mortality impacts


We have thus far established that heat generates an increased threat to the patient who receives health care, either through more crowding at the hospital, which diminishes available resources, or more patients being sent home despite being sicker. To investigate the mortality impacts of this congestion, we conduct three tests of quality of care: i) we look at excess mortality for hospitalized patients, ii) we examine the impact of heat shocks on patients already admitted to the hospital, controlling for heat when admitted, and iii) we focus on excess mortality among cancer patients. Recall that heat leads to an increase in admitted patients but does not change the composition of patients. Therefore, if heat leads to an increase in mortality for hospital admits, the increased crowding from the additional admissions is the likely explanation.

For this test, we examine the impact of temperature on excess deaths. Recall that we define excess deaths as  $\min\{0, \mathbb{1}(m) - \hat{s}\}$ , where  $\mathbb{1}(m) = 1$  if the patient dies and 0 otherwise, where  $\hat{s}$  is the severity assigned to each patient through their diagnosis and demographic characteristics. The results, shown in Figure 8, Panel (a), indicate that excess deaths increase with temperature. Specifically, we find that one day with a



(a) Severity, ED visits

(b) Severity, hospitalizations



(c) Severity, ED visits sent home

Figure 7: Patient composition

*Notes:* ED and hospital visits between 2012–2019. The dependent variable is the average patient severity for (a) all ED visits, (b) all hospitalized patients and (c) all ED patients who were sent home. Severity at the patient-level is measured as the estimated in-hospital mortality rate at the national level for each diagnosis code, conditional on age and sex. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year, day-of-the-week by year and hospital fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

maximum temperature exceeding 34°C leads to a 5% increase in excess deaths. These deaths include the direct physiological impacts of heat as well as any indirect effects arising through hospital congestion.

Disentangling the indirect effects of heat exposure from its direct impacts is particularly challenging because heat exacerbates a wide range of health conditions that are not always classified as heat-related. For instance, elevated temperatures can aggravate preexisting conditions such as cardiovascular diseases or renal failure, often leading to diagnoses unrelated to heat stress (Schulte and Chun, 2009; Bell *et al.*, 2016). Compounding this issue, the lack of standardized protocols for reporting heat-related illnesses



further obscures the link between heat and health outcomes (Vaidyanathan *et al.*, 2019). To address these complexities, we focus on patients already admitted to hospitals and evaluate the effect of heat shocks on their outcomes, controlling for heat exposure at admission. Our findings reveal that one day exceeding 34°C also increases in-hospital mortality by 5% among this group (Panel (b)), indicating that spillover effects amount to at least half of the impact of heat on excess mortality. As an additional test, we limit our attention to cancer patients<sup>19</sup>, who should not be impacted by temperature conditional on hospital admission, and find a similar increase (Panel (c)), underscoring the spillover effects generated by heat-driven overcrowding in the hospital.

One possible explanation for these patterns could be that heat impacts physician care. Just as patients suffer from stress due to heat, physicians may experience the same stress. Although we are unable to directly test this, nearly all hospitals have air conditioning (AC)<sup>20</sup>, which greatly minimizes heat exposure and the impacts of heat on healthcare providers (Barreca *et al.*, 2016). Furthermore, while the level of residential AC ownership is low in Mexico, healthcare providers are drawn from higher socioeconomic strata and are more likely to own them (Davis *et al.*, 2021). Regardless, given the statistically significant contemporaneous effect of temperature, residential AC is less likely to be a contributing factor to the performance of healthcare providers. While we cannot rule out other possible explanations, our results are consistent with crowding as a mechanism behind these results.

### 3.4.1 Mortality outside the hospital system

These results suggest that patients who stay within the healthcare system experience increased mortality due to the heat. What about discharged patients? As shown above, patients with more severe conditions are more likely to be sent home on hotter days. Given that so few SP and uninsured patients are likely to have AC, it is probable that this

---

<sup>19</sup>We classify as cancer patients those who have their main diagnosis code within the ICD categories C00-D49 (Neoplasms).

<sup>20</sup>Using Mexico’s information requests platform (Plataforma Nacional de Transparencia), we sent information requests to each of the 32 states asking about their hospitals’ climate control infrastructure. We obtained information for 200 hospitals, and 197 of them had AC.

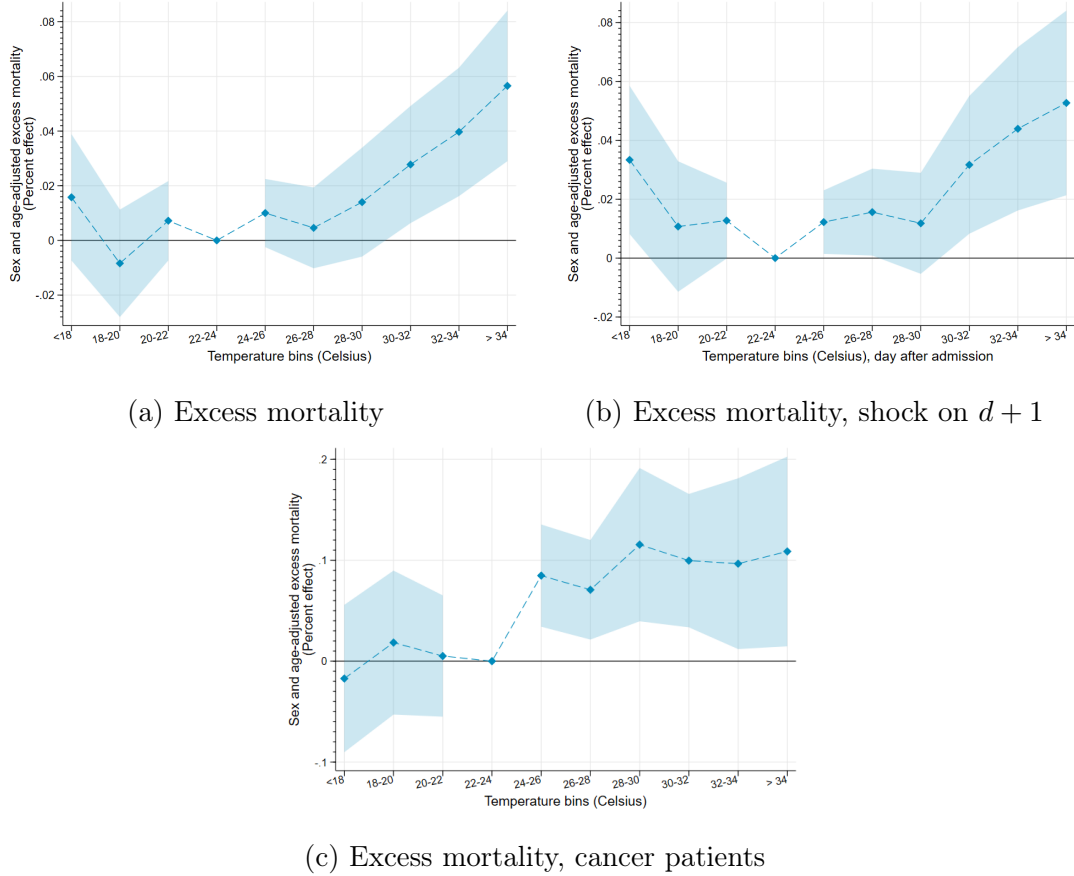


Figure 8: Excess mortality

*Notes:* Hospital visits between 2012–2019. The dependent variable is excess mortality in response to (a) temperature on the day of admission  $d$  and (b) temperature on the day after admission  $d + 1$ , controlling for temperature in  $d$ . Panel (c) shows excess mortality focusing exclusively on cancer patients (ICD categories C00-D49). Severity at the patient level is measured as the estimated in-hospital mortality rate at the national level for each diagnosis code, conditional on age and sex. Excess mortality is defined as the difference between predicted and measured mortality, as defined in Section ???. Econometric specification includes 7 lags of 9 bins. Shaded area represents 95% confidence intervals. Omitted category: 22–24 °C.

lack of treatment increases their mortality. Although we do not observe the mortality specifically for these patients, we observe it for the entire population along with the setting of each death, enabling us to explore the potential role of the healthcare system.

We first estimate the overall municipality-level mortality effects of heat, which serves to validate our model, since many previous studies estimate such a relationship. Consistent with these studies, we find that higher temperatures increase mortality, where a day that exceeds 34°C leads to a statistically significant 0.17 increase in total deaths (Figure 9). From a mean of 3.02, this indicates a 6% increase in mortality. The convex

shape is also consistent with the U-shaped pattern found elsewhere.<sup>21</sup>

Given that we have established an expected relationship between temperature and mortality, we next explore the setting of each death. We have shown that more deaths occur in hospitals, but we have also shown that a disproportionately higher number of patients are discharged home than admitted to the hospital. If this holds true for patients with IMSS and ISSSTE, the other source of social insurance in Mexico, and going home offers less protection from the heat than does the hospital, we would expect a disproportionate increase in deaths at home relative to the increase in the hospital. This pattern is supported by the results shown in Figure 9. Panel (a) indicates that the proportion of deaths in a public hospital decreases by almost 1% when temperatures exceed 34°C, a 4% decrease from a sample mean of 23%. The proportion of deaths outside of the hospital increases by approximately 1% (Panel (b)). Thus, the additional demand on the healthcare system that arises from more patients seeking health care, and the resulting decision to send most of these patients home, appears to have life-threatening implications.

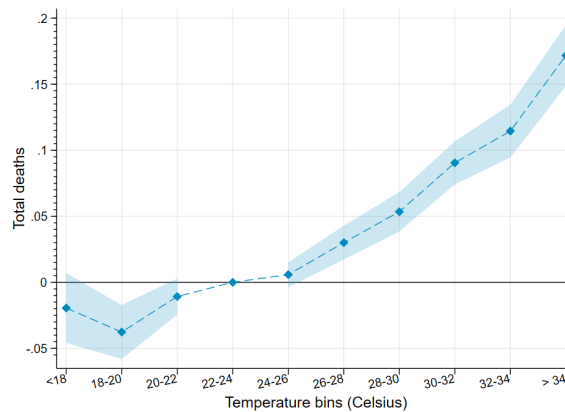


Figure 9: Temperature and daily deaths

*Notes:* Data from the universe of 2012–2019 death certificates. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year and day-of-the-week by year fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

<sup>21</sup>Appendix Figure A.6 also shows that we obtain a U-shaped pattern when we do not control for lagged temperature, consistent with findings from the literature that, overall, mortality responds to cold with some delay (for instance, when it is caused by infectious disease) (White, 2017).

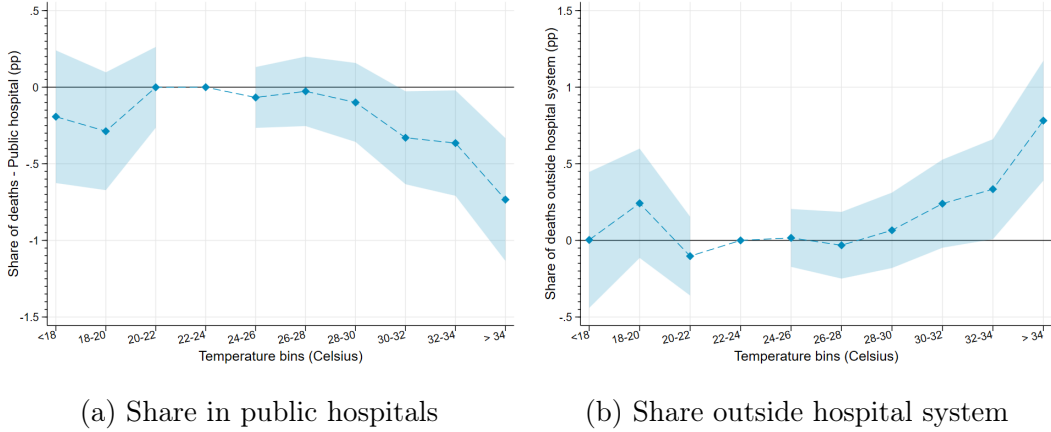


Figure 10: Share of deaths by location

*Notes:* Data from the universe of 2012–2019 death certificates. Econometric specification includes contemporaneous and 7-day lags of 9 temperature bins. Omitted category: 22–24 °C. Controls include day-of-the-year, municipality-by-year and day-of-the-week by year fixed effects. Standard errors are clustered at the municipality level. Shaded area represents 95% confidence intervals.

## 4 Conclusion

Using comprehensive data from Mexico’s largest healthcare subsystem, we provide the first estimates of the impact of extreme heat that parses the physiological effects of heat stress from the spillover effects on health outcomes arising due to hospital congestion. Extreme heat leads to more ED and hospital visits. Due to capacity constraints, some patients are turned away from the system, while those within the system experience deteriorated care. An extra day at which the maximum temperature exceeds 34°C leads to a 6% increase in deaths within the hospital, with over half of these due purely to congestion spillovers, and a 6% increase in deaths at home. Our results are robust to varying lag structures, using wet-bulb temperature as the independent variable, and excluding patients with heat-related diagnoses from the analysis.

To place these numbers in a broader context, we offer a simple back-of-the-envelope calculation of potential impacts under climate change. Under a high-emission, low-mitigation pathway, we project that the Mexican territory in 2050 will experience an annual county-level average increase of 33 days with temperatures exceeding 34°C compared to 2019 (Coupled Model Intercomparison Project Phase 6 (CMIP6), 2017; Haarsma *et al.*, 2016). By 2050, ED visits are projected to rise by 6%, and inpatient hospitaliza-

tions are expected to rise by 16%. This will place increased strain on EDs and hospitals and yield an increase of 1.7 additional deaths outside the hospital system out of 5.33 additional deaths for the average 2050 day. <sup>22</sup>

This pernicious threat from extreme heat comes with a silver lining. Increasing capacity in the healthcare system is a novel arrow that can be added to the rather limited quiver of climate adaptation tools. Moreover, these adaptation strategies can be implemented in the short run by rescheduling patients and reallocating them across hospitals (Gutierrez and Rubli, 2021) as well as in the longer run by creating more facilities for emergency care, adding hospital beds, and increasing the number of healthcare professionals. How labor and capital investments in healthcare systems, as well as better surge management tools, stack up against the rollout of AC, the other principal means for limiting health harms from extreme heat (Barreca *et al.*, 2016), will, of course, depend on many contextual factors, not least of which are income, existing levels of infrastructure and expected changes in the frequency and spread of extreme heat within and across countries. Insofar as AC remains a private good while healthcare systems serve a broader public, they also have very different distributional consequences.

To fix ideas, the starting salary for a nurse in India in 2017 was US \$10,909, while in the US, it was US \$55,969 (George and Rhodes, 2017). Even after adjusting for purchasing power parity, the costs in the US are 1.8 times higher than those in India. At the same time, the feasibility of large-scale residential AC penetration in a place like India is greatly hampered by electricity reliability issues. Indeed, rural households report an average of 11 hours of outages every day (Aklin *et al.*, 2016). Even within Europe, where electricity reliability is not a concern, the cost of household electricity in, e.g., Ireland is more than triple the price in Norway and twice that in Spain (sta, 2023). At the same time, the average Norwegian nurse is paid 10% more than nurses

---

<sup>22</sup>We utilize the SSP5-8.5 scenario, which assumes radiative forcing increases of 8.5 W/m<sup>2</sup> relative to preindustrial levels. Gridded daily temperature estimates for 2050 were produced using the delta method for bias correction. Implicit in our extrapolation using 2019 data except for temperature is the assumption that healthcare infrastructure grows in line with population while adaptation levels, such as AC adoption, remain unchanged from 2012–2019. Appendix Figure A.8 shows that despite a handful of regions observing minimal or negative temperature changes, the broader trend is a large rise in extreme heat exposure. Appendix Tables A.2 and A.3 show projected changes in our outcomes of interest.

in Ireland and 60% more than his or her counterpart in Spain (Yanatma, 2023). The optimal adaptation strategy, which equates the marginal rate of technical substitution between adaptation technologies and practices, will clearly vary across these settings.

This calculus is further complicated by potential nonlinearities in the costs and benefits of each technology. Are there returns to scale in rolling out a residential AC program in countries with very low penetration? How do these returns to scale compare to those in healthcare infrastructure within and across cities and countries? Are there extreme temperatures at which the protective effect of AC or healthcare services is no longer effective? How might these estimates vary under climate change when the frequency of these extreme events is more likely to increase? Do the co-benefits of greater comfort from AC and greater healthcare access on cooler days significantly alter these welfare calculations? Together, these questions comprise a rich agenda for future research.

## References

- (2023). Household electricity prices worldwide in december 2023, by select country. <https://www.statista.com/statistics/263492/electricity-prices-in-selected-countries/>.
- AKLIN, M., CHENG, C.-Y., URPELAINEN, J., GANESAN, K. and JAIN, A. (2016). Factors affecting household satisfaction with electricity supply in rural india. *Nature Energy*, **1** (11), 1–6.
- AMERICAN SOCIETY OF HEATING, R. and ENGINEERS, A.-C. (2017). Ashrae fundamentals handbook.
- BARRECA, A., CLAY, K., DESCHENES, O., GREENSTONE, M. and SHAPIRO, J. S. (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, **124** (1), 105–159.
- BASU, R. (2009). High ambient temperature and mortality: a review of epidemiologic studies from 2001 to 2008. *Environmental health*, **8**, 1–13.
- BELL, J., HERRING, S., JANTARASAMI, L., ADRIANOPOLI, C., BENEDICT, K., CONLON, K., ESCOBAR, V., HESS, J., LUVALL, J., GARCIA-PANDO, C., QUATTROCHI, D., RUNKLE, J. and SCHRECK, I., C.J. (2016). Ch. 4: Impacts of extreme events on human health. In *The Impacts of Climate Change on Human Health in the United States: A Scientific Assessment*, Washington, DC: U.S. Global Change Research Program, pp. 99–128.

- BLOCK, M. Á. G., MORALES, H. R., HURTADO, L. C., BALANDRÁN, A., MÉNDEZ, E., ORGANIZATION, W. H. *et al.* (2020). Mexico: health system review.
- BUNKER, A., WILDENHAIN, J., VANDENBERGH, A., HENSCHKE, N., ROCKLÖV, J., HAJAT, S. and SAUERBORN, R. (2016). Effects of air temperature on climate-sensitive mortality and morbidity outcomes in the elderly; a systematic review and meta-analysis of epidemiological evidence. *EBioMedicine*, **6**, 258–268.
- CARLETON, T. and GREENSTONE, M. (2022). A guide to updating the us government’s social cost of carbon. *Review of Environmental Economics and Policy*, **16** (2), 196–218.
- CORREIA, S., GUIMARÃES, P. and ZYLKIN, T. (2019). ppmlhdfc: Fast poisson estimation with high-dimensional fixed effects, arxiv. org.
- COUPLED MODEL INTERCOMPARISON PROJECT PHASE 6 (CMIP6) (2017). SSP5-8.5 Scenario Data. <https://esgf-node.llnl.gov/projects/cmip6/>, available at <https://esgf-node.llnl.gov/projects/cmip6/>.
- DAVIS, L., GERTLER, P., JARVIS, S. and WOLFRAM, C. (2021). Air conditioning and global inequality. *Global Environmental Change*, **69**, 102299.
- DELL, M., JONES, B. F. and OLKEN, B. A. (2014). What do we learn from the weather? the new climate-economy literature. *Journal of Economic literature*, **52** (3), 740–798.
- DESCHENES, O. (2014). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, **46**, 606–619.
- DESCHÊNES, O. and GREENSTONE, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, **3** (4), 152–185.
- DRESCHER, K. and JANZEN, B. (2024). When weather wounds workers: The impact of temperature on workplace accidents. *Journal of Public Economics*, forthcoming.
- EBI, K. L., CAPON, A., BERRY, P., BRODERICK, C., DE DEAR, R., HAVENITH, G., HONDA, Y., KOVATS, R. S., MA, W., MALIK, A. *et al.* (2021). Hot weather and heat extremes: health risks. *The lancet*, **398** (10301), 698–708.
- FANKHAUSER, S. (2017). Adaptation to climate change. *Annual Review of Resource Economics*, **9** (1), 209–230.
- FOX, M., ZUIDEMA, C., BAUMAN, B., BURKE, T. and SHEEHAN, M. (2019). Integrating public health into climate change policy and planning: state of practice update. *International journal of environmental research and public health*, **16** (18), 3232.
- FRITZ, M. (2022). Temperature and non-communicable diseases: Evidence from Indonesia’s primary health care system. *Health Economics*, **31** (11), 2445–2464.
- GEORGE, G. and RHODES, B. (2017). Is there a financial incentive to immigrate? examining of the health worker salary gap between india and popular destination countries. *Human resources for health*, **15**, 1–10.

- GOULD, C. F., HEFT-NEAL, S., HEANEY, A. K., BENDAVID, E., CALLAHAN, C. W., KIANG, M., ZIVIN, J. S. G. and BURKE, M. (2024). *Temperature Extremes Impact Mortality and Morbidity Differently*. Tech. rep., National Bureau of Economic Research.
- GOURIEROUX, C., MONFORT, A. and TROGNON, A. (1984). Pseudo maximum likelihood methods: Applications to poisson models. *Econometrica: Journal of the Econometric Society*, pp. 701–720.
- GUIDETTI, B., PEREDA, P. and SEVERNINI, E. R. (2024). *Health shocks under hospital capacity constraint: Evidence from air pollution in sao paulo, brazil*. Tech. rep., National Bureau of Economic Research.
- GUTIERREZ, E. and RUBLI, A. (2021). Shocks to hospital occupancy and mortality: Evidence from the 2009 h1n1 pandemic. *Management Science*, **67** (9), 5943–5952.
- HAARSMA, R. J., ROBERTS, M. J., VIDALE, P. L., SENIOR, C. A., BELLUCCI, A., BAO, Q., CHANG, P., CORTI, S., FUČKAR, N. S., GUEMAS, V. *et al.* (2016). High resolution model intercomparison project (highresmip v1. 0) for cmip6. *Geoscientific Model Development*, **9** (11), 4185–4208.
- HANSEN, A. L., BI, P., RYAN, P., NITSCHKE, M., PISANIELLO, D. and TUCKER, G. (2008). The effect of heat waves on hospital admissions for renal disease in a temperate city of Australia. *International Journal of Epidemiology*, **37** (6), 1359–1365.
- HARDUAR MORANO, L. and WATKINS, S. (2017). Evaluation of diagnostic codes in morbidity and mortality data sources for heat-related illness surveillance. *Public Health Reports*, **132** (3), 326–335.
- HOE, T. P. (2022). Does hospital crowding matter? evidence from trauma and orthopedics in england. *American Economic Journal: Economic Policy*, **14** (2), 231–262.
- KAHN, M. E. (2016). The climate change adaptation literature. *Review of Environmental Economics and Policy*.
- LIU, J., VARGHESE, B. M., HANSEN, A., XIANG, J., ZHANG, Y., DEAR, K., GOURLEY, M., DRISCOLL, T., MORGAN, G., CAPON, A. *et al.* (2021). Is there an association between hot weather and poor mental health outcomes? A systematic review and meta-analysis. *Environment International*, **153**, 106533.
- MICHELOZZI, P., ACCETTA, G., DE SARIO, M., D’IPPOLITI, D., MARINO, C., BACCINI, M., BIGGERI, A., ANDERSON, H. R., KATSOUYANNI, K., BALLESTER, F. *et al.* (2009). High temperature and hospitalizations for cardiovascular and respiratory causes in 12 European cities. *American Journal of Respiratory and Critical Care Medicine*, **179** (5), 383–389.
- MULLINS, J. T. and WHITE, C. (2019). Temperature and mental health: Evidence from the spectrum of mental health outcomes. *Journal of health economics*, **68**, 102240.
- MURRAY-TORTAROLO, G. N. (2021). Seven decades of climate change across Mexico. *Atmósfera*, **34** (2), 217–226.



- NHS (2024). *Summary Hospital-level Mortality Indicator (SHMI)*. Tech. rep., National Health Service, England.
- OECD (2023). *Health at a Glance 2023*.
- OWEN, G. (2020). What makes climate change adaptation effective? a systematic review of the literature. *Global Environmental Change*, **62**, 102071.
- OWID (2021). Medical doctors per 1,000 people, 2021. <https://ourworldindata.org/grapher/physicians-per-1000-people>.
- PARK, J. and PANKRATZ, N. (). Temperature, workplace safety, and labor market inequality.
- PÖRTNER, H., ROBERTS, D. C., ADAMS, H., ADELEKAN, I., ADLER, C., ADRIAN, R., ALDUNCE, P., ALI, E., BEGUM, R. A., BEDNAR-FRIEDL, B., BEZNER KERR, R., BIESBROEK, R., BIRKMANN, J., BOWEN, K., CARETTA, M. A., CARNICER, J., CASTELLANOS, E., CHEONG, T. S., CHOW, W., CISSÉ, G., CLAYTON, S., CONSTABLE, A., COOLEY, S. R., COSTELLO, M. J., CRAIG, M., CRAMER, W., DAWSON, R., DODMAN, D., EFITRE, J., GARSCHAGEN, M., GILMORE, E., GLAVOVIC, B., GUTZLER, D., HAASNoot, M., HARPER, S., HASEGAWA, T., HAYWARD, B., HICKE, J., HIRABAYASHI, Y., HUANG, C., KALABA, K., KIESSLING, W., KITO, A., LASCO, R., LAWRENCE, J., LEMOS, M. F., LEMPert, R., LENNARD, C., LEY, D., LISSNER, T., LIU, Q., LIWENGA, E., LLUCH-COTA, S., LOESCHKE, S., LUCATELLO, S., LUO, Y., MACKEY, B., MINTENBECK, K., MIRZABAEV, A., MOELLER, V., MONCASSIM VALE, M., MORECROFT, M., MORTSCH, L., MUKHERJI, A., MUSTONEN, T., MYCOO, M., NALAU, J., NEW, M., OKEM, A., OMETTO, J. P., O'NEILL, B., PANDEY, R., PARMESAN, C., PELLING, M., PINHO, P. F., PINNEGAR, J., POLOCZANSKA, E. S., PRAKASH, A., PRESTON, B., RACAULT, M.-F., RECKIEN, D., REVI, A., ROSE, S. K., SCHIPPER, E. L. F., SCHMIDT, D., SCHOEMAN, D., SHAW, R., SIMPSON, N. P., SINGH, C., SOLECKI, W., STRINGER, L., TOTIN, E., TRISOS, C., TRISURAT, Y., VAN AALST, M., VINER, D., WAIRIU, M., WARREN, R., WESTER, P., WRATHALL, D. and IBRAHIM, Z. Z. (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Tech. rep., Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.
- SAPARI, H., ISA, M. R., ISMAIL, R., MAHIYUDDIN, W. R. W. and SELAMAT, M. I. (2024). The impact of heatwaves on healthcare services in developing countries: A systematic review. *Educational Administration: Theory and Practice*, **30** (5), 4626–4645.
- SCHULTE, P. A. and CHUN, H. (2009). Climate change and occupational safety and health: establishing a preliminary framework. *Journal of occupational and environmental hygiene*, **6** (9), 542–554.
- SILVA, J. S. and TENREYRO, S. (2011). Further simulation evidence on the performance of the poisson pseudo-maximum likelihood estimator. *Economics Letters*, **112** (2), 220–222.

- SUN, Z., CHEN, C., XU, D. and LI, T. (2018). Effects of ambient temperature on myocardial infarction: A systematic review and meta-analysis. *Environmental Pollution*, **241**, 1106–1114.
- THORNTON, M., SHRESTHA, R., WEI, Y., THORNTON, P., KAO, S. and WILSON, B. (2021). Daymet: Daily surface weather data on a 1-km grid for north america, version 4.
- VAIDYANATHAN, A., SAHA, S., VICEDO-CABRERA, A. M., GASPARRINI, A., ABDUREHMAN, N., JORDAN, R., HAWKINS, M., HESS, J. and ELIXHAUSER, A. (2019). Assessment of extreme heat and hospitalizations to inform early warning systems. *Proceedings of the National Academy of Sciences*, **116** (12), 5420–5427.
- WETTSTEIN, Z. S., HOSHIKO, S., FAHIMI, J., HARRISON, R. J., CASCIO, W. E. and RAPPOLD, A. G. (2018). Cardiovascular and cerebrovascular emergency department visits associated with wildfire smoke exposure in California in 2015. *Journal of the American Heart Association*, **7** (8), e007492.
- WHITE, C. (2017). The dynamic relationship between temperature and morbidity. *Journal of the Association of Environmental and Resource Economists*, **4** (4), 1155–1198.
- YANATMA, S. (2023). Following a series of strikes across europe in 2023, nurses in the uk and italy are expected to continue taking to the streets in 2024 for better pay and work conditions. <https://www.euronews.com/next/2023/12/14/nurses-salaries-across-europe-which-countries-have-the-highest-and-lowest-wages>.
- YE, X., WOLFF, R., YU, W., VANECKOVA, P., PAN, X. and TONG, S. (2012). Ambient temperature and morbidity: a review of epidemiological evidence. *Environmental health perspectives*, **120** (1), 19–28.


# A Online Appendix

## A.1 Descriptive statistics

	Mean (1)	SD (2)	Min (3)	Max (4)	Within SD (5)
<b>Panel A: Emergency Departments</b>					
ED visits per day	39.65	40.14	0	5889	07.02
Patients sent home	29.16	31.99	0	5784	14.44
Patients hospitalized	5.36	10.38	0	3779	4.08
Patients referred to another hospital	3.81	9.51	0	659	4.13
Exit without discharge	0.20	1.20	0	99	0.48
Daily deaths in the ED	0.06	0.31	0	50	0.17
Patients sent home (%)	73.43	27.74	0	100	17.12
Patients hospitalized (%)	12.52	16.87	0	100	10.07
Patients referred to another hospital (%)	10.77	19.88	0	100	11.69
Exits without discharge	0.47	3.04	0	100	1.69
Deaths in the ED (%)	0.23	3.39	0	100	1.12
Average severity	0.02	0.02	0	0.93	0.01
<b>Panel B: Hospitals</b>					
Hospital admissions per day	10.53	14.58	0	286	4.46
Length of stay (days)	3.50	5.96	0	365	3.11
Reason for leaving: improvement	9.86	13.83	0	284	4.27
Patients transferred out	0.12	0.48	0	37	0.30
Reason for leaving: voluntary	0.02	0.13	0	1	0.11
Daily deaths	0.23	0.68	0	19	0.35
Reason for leaving: improvement (%)	92.79	16.48	0	100	13.10
Patients transferred out (%)	2.44	10.32	0	100	7.93
Reason for leaving: voluntary (%)	0.56	6.06	0	100	4.53
Deaths (%)	1.61	5.48	0	100	4.43
Average severity	0.02	0.03	0	0.98	0.02
Excess mortality	0.23	0.63	0	16.10	0.31
<b>Panel C: Deaths</b>					
Daily deaths	3.02	5.06	1	96	
Deaths in public hospitals (%)	23.17	35.17	0	100	
Deaths in private hospitals (%)	2.83	12.71	0	100	
Deaths outside of hospitals (%)	73.99	37.05	0	100	

Data from 916 EDs (A), 857 hospitals with inpatient services (B), and municipality-level mortality data in Mexico between 2012 and 2019 (C). Unit of observation: hospital-day (A and B) and municipality-day (C). Shares might not sum to 100 because some patients flee the hospital or their outcome is not recorded. Column (5) presents the within-hospital standard deviation (SD)

Table A.1: Descriptive statistics




(a) ED visits

(b) Hospitalizations

Figure A.1: Descriptive statistics: average severity

*Notes:* Data from 916 EDs and 857 hospitals with inpatient services in Mexico between 2012 and 2019. Unit of observation: hospital-day, with cohorts based on the date of admission of the patient. Observations above the 99th percentile have been excluded

## A.2 Robustness checks



(a) Heat-related admissions

(b) Total admissions excluding heat-related

(c) Heat-related share

Figure A.2: ED admissions

*Notes:* Share of heat-related admissions is expressed in percentage points. The shaded area represents 95% confidence intervals. Omitted category:  $22-24^{\circ}\text{C}$ . Lagged specification.

### A.3 Wet Bulb Temperature

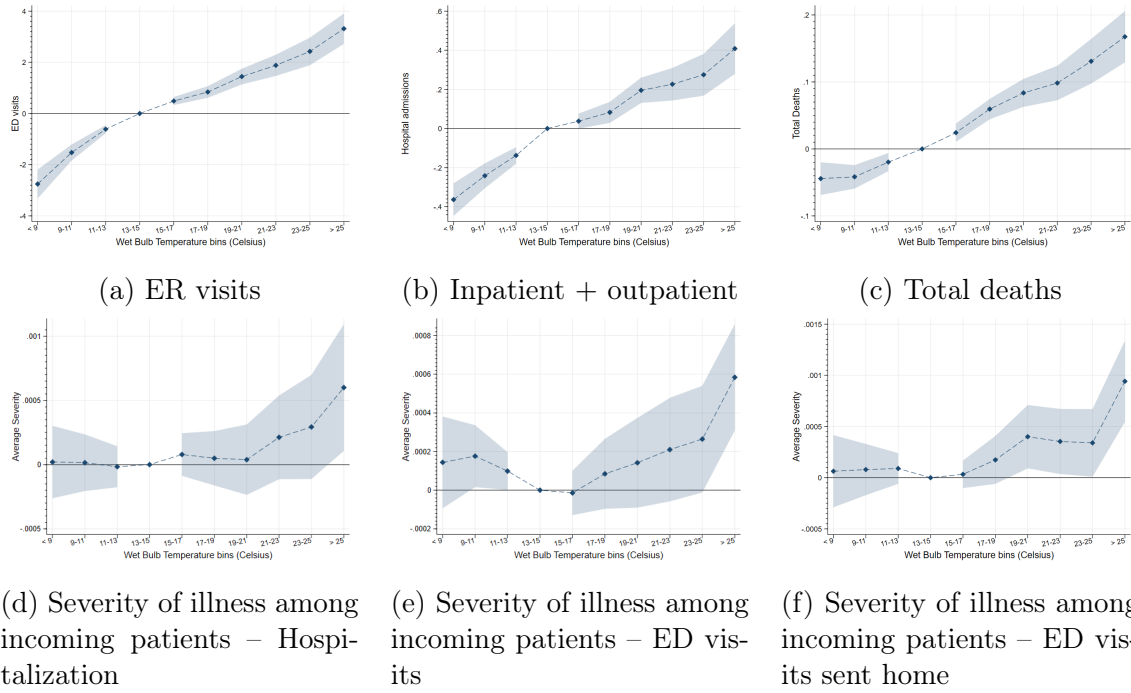



Figure A.3: Main results using wet bulb temperature


Notes: Shaded area represents 95% confidence intervals. Omitted category: 13–15 °C. Lagged specification.

## A.4 Point estimates when varying lag structure


### A.4.1 ED visits per day




(a) No lags




(b) 1 lag




(c) 2 lags




(d) 5 lags




(e) 10 lags




(f) 15 lags



(g) 20 lags



(h) 25 lags

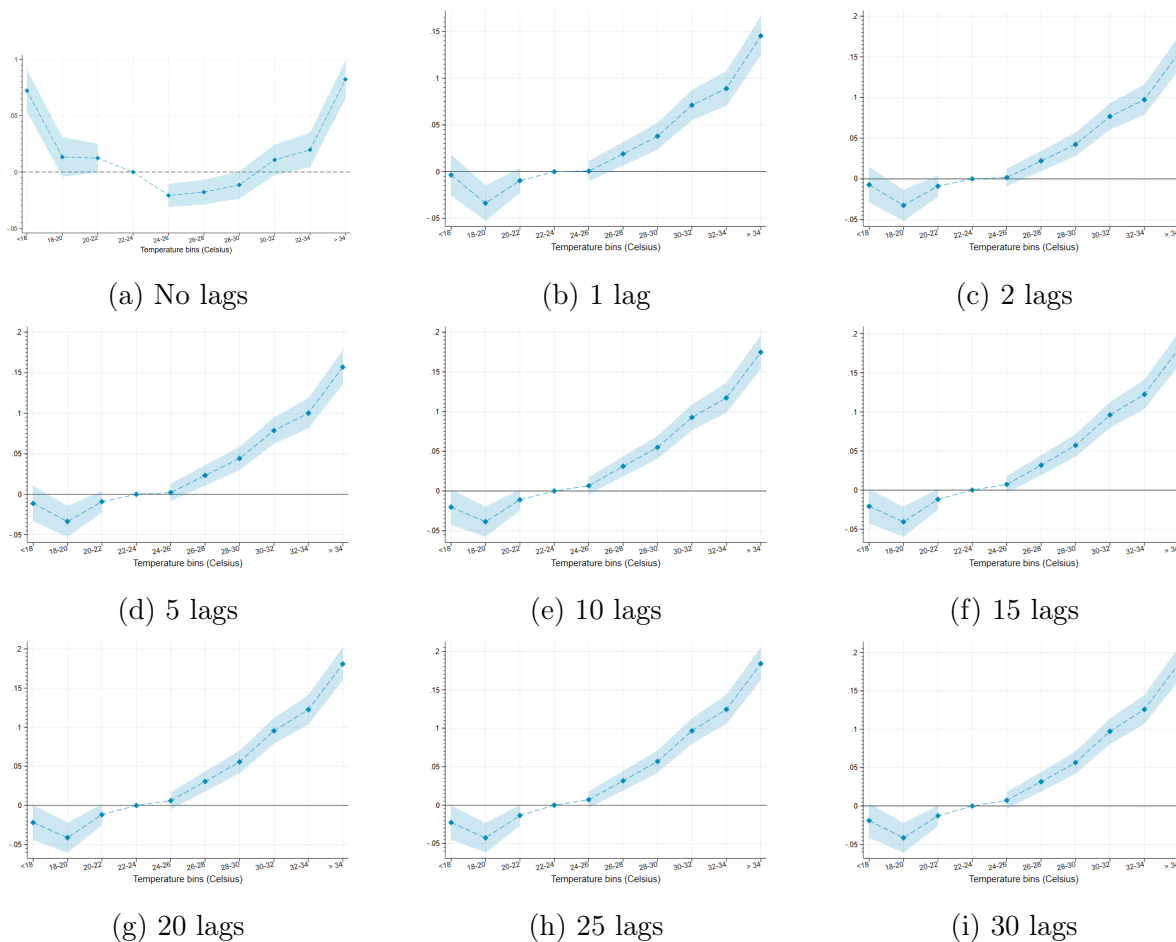


(i) 30 lags

*Notes:* Sample includes hospital data for daily ED visits from 2012 to 2019 in the universe of MoH hospitals. Estimates come from a Poisson regression using contemporaneous and 30-day lags of temperature bins. Controls include municipality-by-month, day-of-week, and year fixed effects. Standard errors are clustered at the municipality level.

Figure A.4: ED visit estimations adding lags – Poisson

### A.4.2 Total deaths (death records) per day



*Notes:* Sample includes municipality-level data for daily deaths from 2012 to 2019. The temperature data during the day are used to construct treatment bins. Estimates come from a Poisson regression using contemporaneous and 30-day lags of temperature bins. Controls include municipality-by-month, day-of-week, and year fixed effects. Standard errors are clustered at the municipality level.

Figure A.6: Total death estimations adding lags – Poisson

### A.4.3 2050 Projections

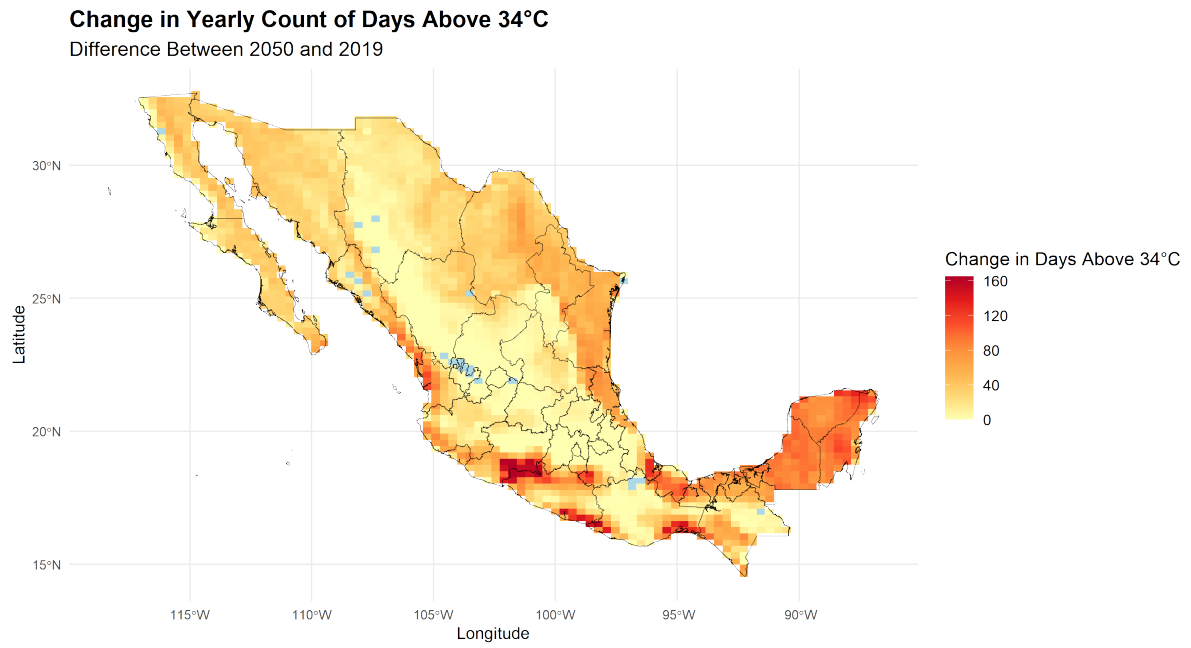


Figure A.8: Days above 34°C in 2050, change from 2019

*Source:* Own elaboration with projections from CIMP6 SSP8 (Coupled Model Intercomparison Project Phase 6 (CMIP6), 2017). Pixels in blue have negative changes. Min=-11



	ED			Hospitalizations	
	ED Visits	Share Hosp.	Share Home	Hospitalizations	Excess Mortality
2019	41.86	12.77	74.16	13.00	0.27
2050	44.30	12.54	74.45	15.41	0.31
Diff. (2050-2019)	2.43	-0.22	0.29	2.40	0.04

Table A.2: Emergency Department and Hospitalization Outcomes

*Source:* Own elaboration with projections from CIMP6 SSP8 (Coupled Model Intercomparison Project Phase 6 (CMIP6), 2017) and dose-response estimates from implementing Equation 1 using 2050 daily temperature for Mexican counties.

	Total deaths	Deaths outside the hospital system
2019	3.02	1.65
2050	8.35	3.43
Diff. (2050-2019)	5.33	1.77

Table A.3: Mortality

*Source:* Own elaboration with projections from CIMP6 SSP8 (Coupled Model Intercomparison Project Phase 6 (CMIP6), 2017) and dose-response estimates from implementing Equation 1 using 2050 daily temperature for Mexican counties.