UNVEILING THE UNSEEN ILLNESS:

PUBLIC HEALTH WARNINGS AND HEAT STROKE*

Lester Lusher,[a] Tim Ruberg[b†]

[a]University of Pittsburgh and IZA

[b]University of Hohenheim

This draft: May 31, 2024

Current version: here

Abstract

We utilize region-day variation in actual vs. forecasted wet bulb globe temperatures (i.e. forecasting errors) to investigate the effects of the first comprehensive heat-health warning system in Japan. We find that heat alerts led to an *increase* in heat stroke counts of 17%. An analysis of mechanisms utilizing several datasets suggests the effect is due to increased reporting, as opposed to potential "adverse" behaviors or substitution in health diagnoses. We further find that four times as many heat strokes are detected in low-income neighborhoods compared to high-income neighborhoods, highlighting severe environmental inequalities in health reporting behavior.

Keywords: heat stroke, climate change, warning effectiveness, avoidance behavior

JEL codes: D90, I12, I18, Q54

*We thank Mika Akesaka, Scott Carrell, Alastair Fraser, Osea Giuntella, Timothy Halliday, Yoko Ibuka, Yoshifumi Konishi, Shotaro Nakamura, Matt Naven, Yuki Nohara, Akira Sasahara, Hitoshi Shigeoka, Jeffrey Shrader, Manasvini Singh, Corey White, Mirka Zvedelikova, and seminar participants at the AASLE 2023 Conference, IAAEU Workshop on Labor Economics, University of Hohenheim, University of Tokyo, Osaka University, and Keio University for helpful comments. We thank the Fire and Disaster Management Agency (FDMA) for providing micro-level data. The views expressed in this paper are those of the authors and do not necessarily reflect the positions or policies of the organizations mentioned above. Any remaining errors are our own. Declarations of interest: none.

†Lester Lusher, Department of Economics, University of Pittsburgh, 230 South Bouquet Street, PA 15260, United States; IZA; Email: lesterlusher@pitt.edu. Tim Ruberg, University of Hohenheim, Department of Economics, Schloss Hohenheim, 70593 Stuttgart, Germany; Email: tim.ruberg@uni-hohenheim.de.

In 2023, the world experienced one of the hottest summers ever recorded in human history. July 3, 2023, was the hottest day ever recorded, a record broken the very next day (and again on July 5 and July 6) (PBS News Hour, 2023). Extreme heat is the leading cause of weather-related death in the United States (more so than hurricanes, floods, and tornadoes, for example) (National Weather Service, 2023), with the Los Angeles Times recently calling extreme heat "far deadlier than we think" (Los Angeles Times, 2021). Worldwide, extreme heat is a significant, yet under-counted, cause of overall death (Ebi et al., 2021). Climate change has led to a drastic increase in average global warming, which has caused an increase in the incidence of record high temperatures and extreme heat events measured across weather stations (Robinson et al., 2021). Historically, single extreme heatwaves have claimed thousands of lives, like in the United States in 1995 (Palecki et al., 2001), Europe in 2003 and 2022 (D'Ippoliti et al., 2010; Ballester et al., 2023), and India in 2015 (Dodla et al., 2017), raising great concern for the implications of climate change in the present and immediate future.[1]

In this study, we investigate the impacts of a public warning system for extreme heat events on behavioral and health outcomes. Despite their potential importance and growing usage, relatively little is known about the efficacy of heat warning systems.[2] The context of our study is Japan. To our knowledge, Japan is one of only several high-income countries to have adopted a heat-health warning system.[3] Experts worldwide agree that there is a global need for heat-health warning systems, yet are severely lacking, even among high-income countries (Li et al., 2022). In-progress examples in the US of heat warning systems at the state level include the HeatRisk program covering a handful of western U.S. states, currently in a piloting phase (National Weather Service, 2024), and an impending program in California, which approved a bill in September 2022 requiring the state to develop a heat wave ranking system.[4,5] Japan piloted a heat-health warning system in 2020, then fully implemented the program nationally in 2021, granting us three years of observations with a warning system in place (2020–2022).

Additional advantages from the Japanese context include its comprehensive hospitalization data and features of the warning system, which allow us to plausibly identify the causal impacts of the warning

[1] In a recent study, Gould et al. (2024) predict the mortality and morbidity effects of future climate change. See Martiello and Giacchi (2010) for a review of the relationship between high temperatures and mortality and morbidity outcomes.

[2] Small-scale and/or correlational studies from epidemiology assessing the potential health impacts from heat warnings include Weinberger et al. (2018), Benmarhnia et al. (2019), Vaidyanathan et al. (2019), and Weinberger et al. (2021).

[3] Other high-income countries to have adopted a heat-health warning system as of 2022 include France, the UK, and Spain (Kotharkar and Ghosh, 2022) as well as South Korea (Chae and Park, 2021).

[4] California Legislative Information, "AB-2238 Extreme heat: statewide extreme heat ranking system," retrieved March 1, 2023.

[5] In the United States, a 2022 National Public Radio article dissected the insufficiencies of the current heatwave warning system. For example, the primary variable for the U.S. warning system is the National Weather Service's "heat index," which underestimates the effect of extreme temperatures on the human body, thus underselling the true hazard from heat. Furthermore, only vague information is given regarding health hazards in response to heat from the U.S. warning system (National Public Radio, 2022).

system on health outcomes. First, Japan provides extensive information for heat stroke outcomes by severity diagnosis, demographic, and location, among other variables, whereas in the United States, hospitals and health care providers are not required to report heat-related illnesses to public health agencies.[6] Furthermore, in Japan, micro-level data exist for the population of ambulance dispatches, including various information on patient characteristics, health outcomes, and location. Given the time frame of our study, we are also able to utilize additional data from the Google Mobility Reports to measure behavioral changes from households (e.g., staying indoors).

For plausibly exogenous variation, Japan issues warnings across 58 spatial areas based on whether the *forecasted* wet bulb globe temperature (WBGT) exceeds a certain threshold ($33.0°C$). Combined with data on whether a warning was issued for a region and *actual* WBGT, our econometric model effectively allows us to estimate the effect of a heat warning while flexibly controlling for local daily WBGT. In other words, our data include region-days where an alert was issued but the actual WBGT was below $33.0°C$ (i.e. a false positive) *and* region-days where an alert was not issued but the actual WBGT exceeded $33.0°C$ (i.e. a false negative). Thus, our econometric model separately disentangles the effects of a heat warning from high heat itself using exogenous variation in forecasting errors.

The theoretical impacts of heat warnings on behavioral and health outcomes are ambiguous. On the one hand, a warning may make the dangers of high heat more salient, and thus encourage households to engage in avoidance behaviors (to mitigate damages from heat). On the other hand, it may be that the optimal response to extreme heat from some individuals is to leave their homes if in expectation, their "outside" option provides better protection from heat (e.g. the temperature is so hot that air conditioning at home is inferior to going elsewhere, such as a mall). An erroneous heat alert may then nudge some households into exercising this outside option when ex post they would have been better off staying at home. Further, the warning may also serve as an awareness shock, informing households to pay more attention to their own symptoms and those of others for potential heat stroke, and thus increase the probability they report a heat stroke. Finally, doctors could also be influenced by a heat warning in the sense that they are nudged to diagnose an illness with symptoms similar to a heat stroke on a day with an alert.

We precisely estimate a substantial *increase* in heat stroke cases in response to a heat-health warning being issued. For our fully specified model, which flexibly controls for same-day and prior-day local WBGT (i.e. dummies for each WBGT level, rounded to the nearest tenth) as well as prefecture and date fixed effects, a heat-health warning increases the incidence of heat stroke hospitalizations by 17%. The effect is very precisely estimated, and exists across the spectrum of severity diagnoses, age groups, and locations.

[6] Centers for Disease Control and Prevention, "Picture of America Heat-Related Illness Fact Sheet," retrieved March 1, 2023

We further show the robustness of these results to various considerations, including additional controls for region-specific time trends and COVID-19 policies, alternative difference-in-differences estimators, estimating OLS vs. log-linear vs. negative binomial regression models, different sample restrictions, commuting behavior, randomization inference, alternative fixed effects, different timing of a warning, and intertemporal models. The result appears even in a simple descriptive figure, where on average, for the exact same WBGT, region-days with a heat alert experience more heat strokes than region-days without a heat alert.

In order to shed some light onto the potential channels behind this result, we start by utilizing two datasets from Google: Google Trends and Google Mobility Reports. From Google Trends, we find that the heat alert raised general awareness, as reflected by increases in searches for the terms "heat stroke," "heat stroke alert," and "temperature." Interestingly, searches for the term "weather" decreased, suggesting searches may have substituted away from abstract weather conditions ("weather") in favor of specific characteristics ("temperature"). Searches for "air conditioner" also significantly increased. Searches for popular indoor and outdoor activities remain mostly unaffected, with a small increase in searches for "sea bathing."

We then utilize the Google Mobility Reports data, which tracked individual daily foot traffic across locations by several broad categories from 2020 to 2022, while also considering the home residence and workplace of the tracked individual. We do not find any change in people's visits in response to an alert except for a small reduction in time spent at home. Because this reduction cannot be explained by an increase in visits to any other place category in the Mobility data, it is likely this reduction is driven by people visiting the hospital to report a heat stroke.

Further, we utilize micro-data from the population of ambulance records in Japan.[7] We find that total ambulance calls increase in response to a heat alert. Thus, the main effects cannot be solely attributed to health diagnoses "substituting" away from other sudden illnesses and into heat stroke (since substitution would imply no change in ambulance pickups). We then disaggregate ambulance transports by diagnosed illnesses, and no effects on other illnesses are found, with the exception of sudden illnesses that include heat stroke. This again suggests that there is no substitution away from other sudden illnesses.

Taken together, the substantial increase in heat strokes on alert days is due to an increase in reporting of heat strokes that would otherwise have gone undetected absent an alert. Furthermore, because we also find a small increase in energy usage, searches for the term "air conditioning," and avoidance behavior from respondents of our own survey, it is likely that the *actual* number of heat strokes decreases in response to the alert. Thus, our estimated treatment effect represents a lower bound of the heat stroke cases that

[7] Discussed in further detail later, and as described in Akesaka and Shigeoka (2023), ambulance services are provided to the public for free in Japan, and thus constitute a significant share of hospital transport in Japan.

are going "unnoticed" during extreme heat. A back-of-the-envelope calculation shows that at least one in ten reported cases on high-heat days would have gone undetected without the warning system. Moreover, the warning system detects four times as many heat strokes in low-income neighborhoods vs. high-income neighborhoods (relative to their means), highlighting how lower income neighborhoods (absent the alert) are substantially less likely to visit the hospital. These results have substantial implications for the climate-health relationship and environmental injustices, as previous studies may be severely underestimating the negative health consequences of extreme heat (since many heat stroke cases are going unreported absent an alert), especially for low-income neighborhoods.

This paper provides several important contributions to the literature. To our knowledge, it is the first large-scale study to examine the causal effects of a heat warning system. It is also one of only two studies to examine the effects of *any* public alert system on direct health outcomes, with nearly the entirety of the previous literature focusing on how air quality alerts impact avoidance behavior.[8] Early work from Neidell (2009) and Zivin and Neidell (2009) examine avoidance behavior in California in response to smog alerts, finding reduced outdoor activity when alerts are issued, with reduced efficacy for alerts made on consecutive days. Noonan (2014) finds that elderly users and exercisers reduce their use of a major park in response to a smog alert, with no effect on driving behavior. Ward and Beatty (2016) find similar behavioral effects of reduced outdoor activity, particularly among the elderly. Saberian et al. (2017) find reductions in cycling in Sydney when an air quality alert is issued.[9] Finally, Ferris and Newburn (2017) find that mobile phone alerts for flash flood events in Virginia led to reductions in traffic volume and car accidents.[10]

Our study also relates to the strand of literature identifying factors that influence the climate-health relationship. Perhaps the two most directly related studies to ours come from Neidell et al. (2021) and He and Tanaka (2023), who study electricity prices and energy-saving campaigns, respectively, in the Japanese context.[11] Other prominent studies identifying adaptive behaviors to climate include Deschenes and Moretti

[8] To our knowledge, the only study examining how a public alert impacted direct health outcomes comes from Mullins and Bharadwaj (2015), who find that an "Environmental Episode" policy from Chile reduced air pollution by 20% and improved day-of elderly mortality.

[9] Beyond air quality alerts, Gutteling et al. (2018) investigate behavioral responses from 643 survey participants to a Dutch civil defense warning system, carried via cell phone messages, which warned citizens of various local natural and man-made threats (e.g. toxic fumes released with fire).

[10] A separate strand of related research investigates how private health alerts affect health outcomes (Kim et al., 2019; Iizuka et al., 2021), finding that alerts overall have little effect on behavior, but can have substantial positive effects for those who are at high risk (for diabetes).

[11] Our study also differs from these two in several noticeable ways. First, we study different treatment variables (heat-health warnings vs. energy-saving campaigns vs. electricity prices). Further, whereas the previous two studies utilize monthly-regional variation driven by the Fukushima accident, our data are more granular (daily and smaller spatial units) while our identifying variation utilizes both space-time variation and variation in whether forecasted WBGT exceeded a specific threshold. Our data also include health outcomes beyond mortality, including severity diagnoses of heat stroke. Lastly, given the recent timeframe of our study, we are able to utilize Google Mobility Reports data to measure behavioral outcomes such as staying indoors.

(2009) (migration), Zivin and Neidell (2014) (staying indoors), and Barreca et al. (2016) (air conditioning).[12] Finally, Mullins and White (2020) show that better access to health care can mitigate the negative impacts of heat. Our study contributes to this literature by being the first to look at how public alerts affect the climate-health relationship, and along with He and Tanaka (2023), is one of the first studies to evaluate any government policy specifically targeting behavioral responses to higher temperatures.[13]

1 Background

In 2020, Japan introduced a comprehensive heat warning system to raise awareness of heat-related illnesses and promote heat stroke prevention measures. In the first year of implementation, the program was tested in nine prefectures (Chiba, Gunma, Ibaraki, Kanagawa, Nagano, Saitama, Tochigi, Tokyo, and Yamanashi), and then the warning system was rolled out to all regions of Japan starting in 2021 (Ministry of Environment, 2021). Alerts are communicated via the Ministry of Environment's homepage, weather news broadcasting stations, radio, public institutions, social networking services (i.e. LINE), and smartphone applications. The way these alerts are communicated (e.g. through newspapers, TV, public billboards and loudspeakers, or TV screens in trains) makes it likely that most people receive this information.[14] The alerts encourage people to stay at home, drink a lot, take salt, avoid exercise, use air-conditioning, and report to others if they are in poor physical condition (i.e. children, elderly, and disabled). As shown in Figure 1, Google Searches for the term "heat stroke alert" in Japan began when the warning was introduced and were especially high in the summer, and as shown later econometrically, daily searches in a region increased significantly when a warning was issued.

Warnings are issued based on whether the region-specific forecasted heat index "Wet Bulb Globe Temperature" (WBGT) (Yaglou et al., 1957) exceeds a threshold of $33.0°C$. This index serves as a more comprehensive indicator of heat stress and its impact on the human body than temperature alone, because it also incorporates humidity, radiant heat, and wind speed, which in combination with temperature influence the risk of heat stroke. In order to calculate forecasted WBGT for a specific region, the Japan Meteorological Agency (JMA) collects information on temperature, humidity, amount of rain, wind strength, and solar radiation across 1,300 weather stations. All of these stations collect information on rain, while 840 stations collect further information about temperature, humidity, and wind strength/direction. Among all weather stations, only 11 stations collect information about wet bulb temperature, wbt, black bulb temperature, bbt,

[12] Additional studies broadly discuss the role of adaption on the climate-health relationship, including Deschênes and Greenstone (2011), Geruso and Spears (2018), Mullins and White (2019), and Heutel et al. (2021).

[13] Our work also complements Shrader et al. (2023) who estimate the mortality costs of inaccurate forecasts in the US.

[14] In the online appendix, we present examples of such public announcements.

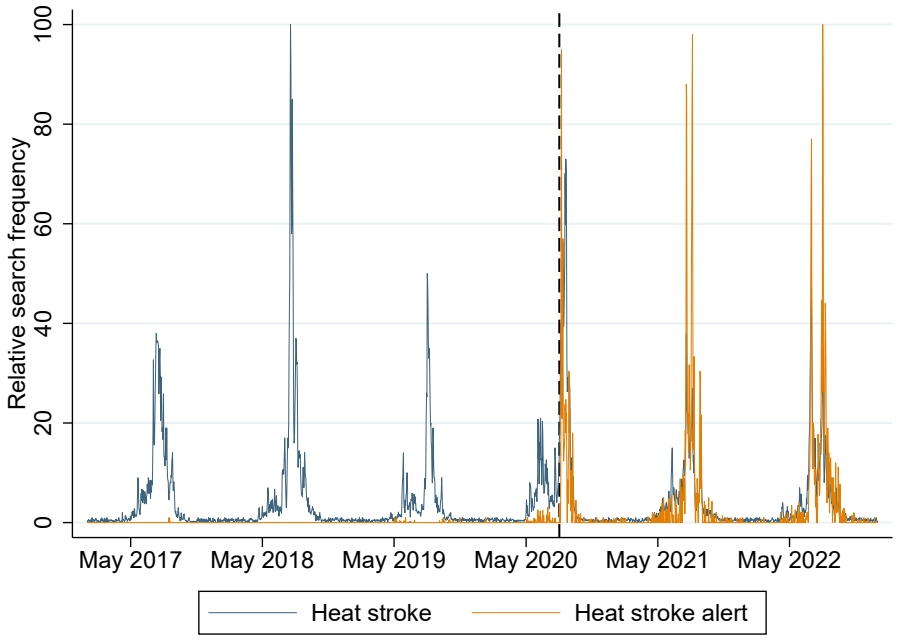


Figure 1: Google Searches for Heat Strokes and Heat Stroke Alerts

Source: Authors' presentation using data from Google Trends. *Note*: This graph shows Google Trends Searches for 熱中症 (heat stroke) and 熱中症警戒アラート (heat stroke alert), respectively. The dashed line indicates the introduction of the Heat Stroke Alert System in August 2020.

and dry bulb temperature, dbt, and calculate WBGT according to

$$\text{WBGT} = 0.7 \cdot wbt + 0.2 \cdot bbt + 0.1 \cdot dbt. \tag{1}$$

44 stations collect information about global solar radiation in order to estimate WBGT following Ono and Tonouchi (2014):

$$\begin{aligned} \text{WBGT} = {} & 0.735 \cdot Ta + 0.0374 \cdot RH + 0.00292 \cdot Ta \cdot RH + 7.619 \cdot SR \\ & - 4.557 \cdot SR^2 - 0.0572 \cdot WS - 4.064, \end{aligned} \tag{2}$$

with Ta being the temperature in $°C$, RH being the relative humidity in $\%$, SR being global solar radiation in kW/m^2, and WS being the average wind speed in m/s. 98 stations do not collect information about global solar radiation, but estimate it using the relationship between the previous 10 minutes of sunshine hours (min) and the relationship between clear-sky irradiance and total solar irradiance based on historical observations. 687 stations further do not collect information about humidity. This information is estimated using numerical forecast data or predicted values from the JMA that is reanalyzed with observed values

from surrounding observations. The remaining 460 stations only collect information about precipitation and therefore cannot estimate WBGT.

With a formula for WBGT for each of 840 weather stations, an alert is issued for a region based on whether *any* of the stations within a specific region have a forecasted WBGT greater than $33.0°C$, i.e. whether the maximum forecasted WBGT across all stations within a region exceeds $33.0°C$. Furthermore, the decision rule is based on two forecasts: 5pm of the previous day or 5am of the same day. For example, if the 5pm forecast for tomorrow's WBGT is $33.1°C$ for a single station in the Osaka region, then an alert is issued for Osaka (and will be broadcast across Osaka as a warning for tomorrow). If tomorrow, the same-day 5am forecast is at most $32.9°C$ across all stations in Osaka, then the alert still applies. In some cases, the 5pm forecast is below $33.0°C$ while the 5am forecast is above $33.0°C$, in which case an alert is issued in the morning, even though no warning was issued the evening before. In our main analysis, for simplicity, we consider whether a warning was ever issued, regardless of the consistency between 5pm the previous day and 5am the same day, while later, for robustness, we separately consider the warnings issued based on the two forecasts.

2 Data and Summary Statistics

2.1 Heat Alerts

The Ministry of Environment provides information about the issuance of heat alerts for each of 58 regions on a given day since the implementation of the program in 2020 (Ministry of Environment, 2024). As described in the previous section, these alerts are either issued at 5pm on the previous day or 5am on the same day. Appendix Figure A1 shows the fraction of dates between May 1 and September 30, 2022, that had a heat stroke warning by region. There exists substantial variation across regions in how often alerts were issued, with some regions issuing zero alerts (Aomori, Akita, Iwate, and Miyagi) and others issuing alerts for as many as 26% (Oita) of the summer days in 2022.

2.2 Wet Bulb Globe Temperature (WBGT)

Although heat alerts are issued on the regional level based on the maximum forecasted WBGT among all weather stations within a region, data from the Ministry of Environment only includes daily measurements for the maximum *realized* WBGT for each of 840 weather stations, rounded to the first decimal. The data

cover the dates between May 1 and September 30 for the years 2017 to 2022.[15] To reflect the decision variable, we calculate the maximum realized WBGT across all stations within a region on a specific date as our primary control variable. Appendix Figure A2 plots the full distribution of observed daily WBGT across region-days for the full sample. While most region-days have WBGTs between $25°C$ and $33°C$, there are also extreme WBGTs of more than $35°C$, which are extremely dangerous and can be fatal.

Importantly, because we do not have information about the 5pm and 5am forecasted WBGTs, i.e. the "running variable" used to determine whether an alert was issued for a specific region, our identification strategy will effectively leverage variation in whether an alert was issued controlling for actual WBGT (described further in Section 3). In particular, alerts do not perfectly match actual WBGT because of forecasting errors. Thus, although we do not observe the actual decision variable for determining an alert, we are still able to flexibly control for the sole concerning confounder of the heat alert (actual WBGT), and are thus able to disentangle the effect of the alert from high heat itself.

2.3 Behavioral Outcomes

2.3.1 Google Trends

To analyze potential mechanisms, we utilize data from Google Trends to examine how searches for various terms on Google's search engine changed in response to heat alerts. Google Trends provides measurements of the relative search frequency of a queried search item (e.g. heat stroke alert) on Google Search for a selected geography (e.g. Tokyo) and time period (e.g. May 1 to September 30, 2021), indexed to a range of 0 and 100, where 100 represents the date within the specified geography that had the greatest search volume for the queried search item. Importantly, these indices are generated from relative search volume, not total searches on Google. Google Trends takes a random sample of total searches for the queried geography and time frame, then generates the indices based on the share of total searches that the queried item constituted. We consider searches for the terms heat strokes, heat stroke alert, weather, temperature, air conditioner, outdoor, sea bathing, park, indoor, cinema, and karaoke for each region for May 1 to September 30 for 2020 to 2022, respectively. These terms describe the awareness of an alert and the risk of heat stroke as well as the intention of potential avoidance behavior. Appendix Table A2 provides summary statistics for these terms.[16]

[15] Detailed data can be downloaded following the description in the manual: https://www.wbgt.env.go.jp/man15NH/R04_wbgt_data_service_manual.pdf, retrieved October 10, 2022.

[16] We downloaded the data for each prefecture and year (May 1 to September 30) separately, such that numbers are relative to the searches within a specific prefecture in a year. Alternatively, one could download all regions on a given date and repeat this for all required dates. Since the data are indices related to a specific geographic and temporal area, the numbers can change. Nevertheless, our results are robust to downloading the data for all regions of Japan at different points in time (not reported).

2.3.2 Google COVID-19 Community Mobility Reports

In response to the COVID-19 pandemic, Google provided aggregated, anonymized "reports" which tracked daily movement by location. The data used to generate the reports are the same used in Google Maps, where users can track how busy a specific location is. In particular, the Google Mobility Reports include information on shares of people who left their residence, and which locations individuals visited across different categories: retail and recreation, groceries and pharmacies, parks, transit stations, workplaces, and residential. These reports were provided from January 2020, with daily updates ceasing on October 15, 2022. Visitations are reported as percentage changes in daily visits relative to visitations for that weekday pre-Covid. Naturally, these data have been used to investigate questions related to traveling behaviors in response to COVID-19 and corresponding policies (e.g. Chetty et al., 2020; Fernández-Villaverde and Jones, 2020; Sulyok and Walker, 2020; Karaivanov et al., 2021; Mendolia et al., 2021). For our study, we utilize mobility reports at the region-day level in order to investigate behavioral responses to the heat alerts.[17] Since these data only start from 2020, analyses using the Community Mobility Reports do not include any pre-treatment observations.[18] Appendix Table A3 provides summary statistics for these reports.

2.4 Health Outcomes

2.4.1 Heat Stroke

The primary health outcome for our study is heat stroke. Heat stroke is the most serious heat-related illness (vs. heat exhaustion, rhabdomyolysis, heat syncope, heat cramps, and heat rash) (Center for Disease Control and Prevention, 2023). With a heat stroke, the body typically loses the ability to control its own temperature, which can lead to permanent disability or death without emergency treatment. Symptoms include very high body temperature, potential loss of consciousness, seizures, dry skin, or profuse sweating. Treatment typically includes cooled IV fluids, a cooling blanket, an ice bath, medication for seizures, or supplemental oxygen. Broadly speaking, there are two types of heat strokes: exertional heat stroke results from physical overexertion in hot, humid conditions, while non-exertional, or "classic," heat stroke is the more common type, and is due to passive exposure to extreme heat and is most common in older people and those in poor health (Bouchama et al., 2022).

The Fire and Disaster Management Agency (FDMA) provides administrative data on the number of heat strokes by age group, place of incidence, and severity diagnosis, for each of 47 prefectures on a daily

[17] To our knowledge, our study is the first to use these data to investigate mobility responses outside of Covid-related policies.
[18] More information on these reports can be found at https://www.google.com/covid19/mobility/.

basis in the summer months since 2008.[19] These data include all individuals who were transported to a hospital via an ambulance and diagnosed with heat stroke. As Akesaka and Shigeoka (2023) describe, ambulance services are provided to the public for free, and thus, utilized similarly by individuals across socioeconomic status. Ambulance use is extremely prevalent among patients in Japan, and in general, Japanese policymakers express significant concern of overuse of ambulance services. For example, Kadooka et al. (2017) estimate that nearly 50% of ambulance rides were for minor conditions where patients could have taken a taxi instead of an ambulance. In response to the summer heat of 2023, Tokyo is on pace to set a new record number for ambulance dispatches (The Japan Times, 2023). Thus, though we do not observe visits made to hospitals via methods other than an ambulance, the data likely cover most visits to hospitals. Furthermore, it is important to note that the data consist of *diagnosed* and not actual heat strokes, which allows for potential misdiagnoses, especially in terms of severity.[20] Thus, the planned duration of the hospital stay may differ from the actual duration depending on the success of the patient's treatment and constraints on the part of the patient or the hospital.

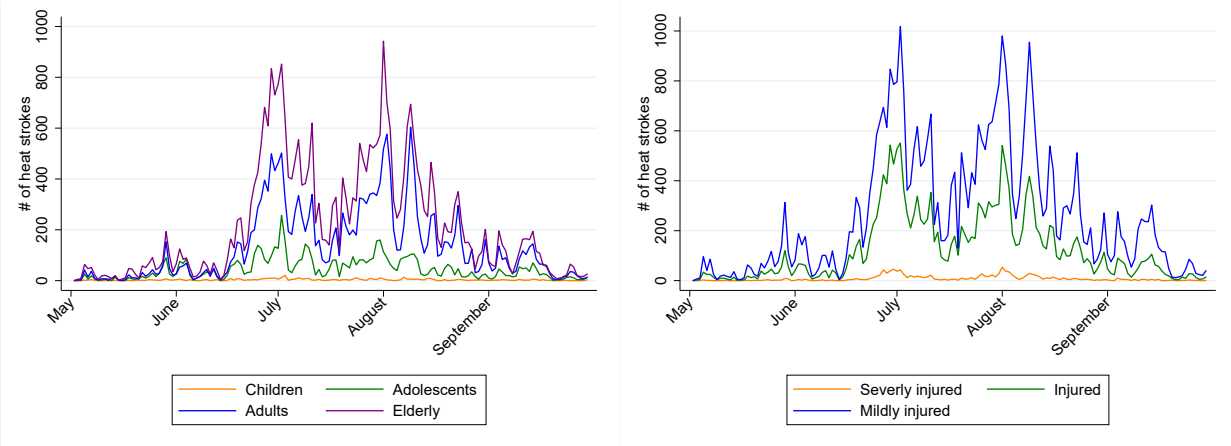
In Figure 2, we plot counts of daily heat strokes by age group and severity diagnosis for the entirety of Japan in 2022. Consistent with the prior health literature, elderly are at the most risk for heat stroke, and heat stroke incidence is strongest in July and August. Although most cases are mild, there are a few severe diagnoses (which require multiple weeks of hospitalization).

We match these prefecture-daily data to the Ministry of Environment's region-daily data. Three prefectures (Hokkaido, Kagoshima, and Okinawa) are separated into 8, 2, and 4 regions, respectively, whereas the remaining 44 prefectures perfectly map into the corresponding 44 regions. For simplicity, we drop these three prefectures.[21] Since the Ministry of Environment's data start from 2017, our matched sample begins in 2017, granting us an additional three years of observations pre-warning system. In Appendix Figure A3, we overlap counts of heat strokes with heat alerts for all of Japan for 2021 and 2022. Unsurprisingly, a positive correlation arises since heat alerts are issued on hotter days, and heat stroke incidence is higher on hotter days.

[19] Website: https://www.fdma.go.jp/disaster/, retrieved October 10, 2022.

[20] The diagnosis about the degree of injury or illness is made during the initial examination at the hospital to which the patient was transported.

[21] Beside these data reasons, at least Hokkaido and Okinawa are unique in their own right and thus should be excluded from the analysis. Hokkaido is the northernmost island relatively close to Siberia and as such has a different climate than the rest of Japan. Similarly, Okinawa, a cluster of islands to the south of Kyushu, is historically and geographically different from the rest of Japan.



(a) Age

(b) Severity diagnoses

Figure 2: Heat Strokes by Age and Severity Diagnoses

Source: Authors' presentation using data from Fire and Disaster Management Agency (FDMA). *Note*: This graph shows the number of heat strokes by age group (a) and severity diagnoses (b) in 2022 for Japan. "Mildly injured" are any heat stroke incidents that do not require hospitalization. "Injured" are incidents that require no more than three weeks of hospitalization, while "severely injured" are those that require at least three weeks of hospitalization.

2.4.2 Ambulance Transports

The FDMA also provides administrative data on the population of ambulance transportations to hospitals in Japan. The data cover all transports from 2017 to 2021 for each of the 47 prefectures on a daily basis. Appendix Table A1 provides summary statistics for the number of transports by severity diagnosis and type of illness. There are on average 130 ambulance transports per million people per day. The majority of these transports are for mild cases (92%), while 7% are for severe cases, and 1% result in death. Similar to the heat stroke severity categorization, "severe" cases suggest multiple nights of hospitalization, while "mild" cases require no hospitalization.

We utilize these data for multiple purposes. First, assuming the majority of cases that result in death in Japan utilize an ambulance, these records allow us to observe the effect of the heat alerts on region-day-level mortality. Second, discussed in further detail later, these data allow us to investigate potential substitution between heat stroke cases and other sudden illnesses. In particular, we can distinguish between injuries resulting from natural disasters, fire, or water (0.1%), traffic accidents (6.9%), work- or sports-related accidents (2.0%), sudden and general illnesses (81.7%), injuries due to violence and self-harm (1.0%), and transfers from one hospital to another (8.2%). Since heat stroke falls under the "sudden and general illnesses" category, the other categories also grant us useful placebo tests. Third, because the data allows

analyses on the fire station level,[22] we merge municipality-level data on local income to conduct heterogeneity analyses within regions.

2.5 Sample Restrictions

Our heat stroke data consist of 44 regions with information about maximum realized WBGT and heat alerts for each day between May 1 and September 30 for 2017 to 2022. However, because there are almost no heat alerts for region-days with a low WBGT, our primary analyses restrict the data to days with a WBGT of at least $28°C$. Further, there are a few region-days with an extremely high number of heat strokes; to lessen the effect that potential outlier events drive our results, we drop region-days with heat strokes above the 99^{th} percentile. Robustness exercises test the sensitivity of our main results to these two sample selections. Our final primary sample consists of 20,612 observations.

2.6 Summary Statistics for Heat Stroke

Table 1 provides summary statistics for the months of May through September for the years 2017 to 2022 at the region-day level. There are approximately 5.4 heat strokes every day per million people. The majority of heat strokes are diagnosed as mild injury, which are cases that do not require overnight hospitalization. Fewer cases are diagnosed to suggest hospitalization of up to three weeks whereas only a few cases suggest hospitalization of more than three weeks. While only very few heat strokes are experienced by children, over half of all heat strokes are experienced by those aged 65 years and older. Interestingly, the most common location for a heat stroke to occur is in one's home. This, coupled with the fact that elderly are at greater risk of heat stroke, highlights how the majority of heat strokes occur due to age and underlying health, and are not necessarily due to one overexerting themselves outdoors. Heat alerts were issued on approximately six percent of region-days. This number increases to almost 17% when only considering the years 2021 and 2022, the period the alert system was fully implemented. Finally, the average WBGT in the summer months was around $31°C$ (after restricting the data to WBGT $\geq 28°C$).

3 Empirical Strategy

A heat stroke warning was issued for region r on day t when the maximum forecasted WBGT from either the 5pm forecast on day $t-1$ or the 5am forecast on day t exceeded $33.0°C$ at any weather station

[22] There are 752 fire stations, each responsible for on average 2.6 municipalities.

Table 1: Summary Statistics

	(1) Obs.	(2) Mean	(3) SD	(4) Min	(5) Max
Panel A: Outcome variables					
Heat strokes: total	20612	5.379	6.10	0.00	70.19
Heat stroke diagnosis: mild injury	20612	3.411	3.92	0.00	53.90
Heat stroke diagnosis: injury	20612	1.814	2.38	0.00	33.99
Heat stroke diagnosis: severe injury	20612	0.123	0.33	0.00	8.24
Heat strokes: 28 days–6 years	20612	0.040	0.14	0.00	5.40
Heat strokes: 7–17 years	20612	0.610	1.12	0.00	36.35
Heat strokes: 18–64 years	20612	1.877	2.32	0.00	29.76
Heat strokes: $\geq$65 years	20612	2.851	3.29	0.00	34.22
Heat stroke loc.: home	20612	2.167	2.71	0.00	34.38
Heat stroke loc.: workplace (construction)	20612	0.603	0.99	0.00	14.09
Heat stroke loc.: workplace (field)	20612	0.132	0.36	0.00	7.36
Heat stroke loc.: education	20612	0.298	0.64	0.00	27.29
Heat stroke loc.: public (inside)	20612	0.409	0.64	0.00	20.02
Heat stroke loc.: public (outside)	20612	0.651	1.13	0.00	36.35
Heat stroke loc.: street	20612	0.821	1.05	0.00	17.99
Heat stroke loc.: other	20612	0.298	0.58	0.00	13.56
Panel B: Treatment variable					
$= 1$ if heat alert	20612	0.063	0.24	0.00	1.00
Panel C: Control variable					
WBGT index	20612	31.212	1.86	28.00	37.60

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table presents descriptive statistics of all key variables used in the analysis. Summary statistics for heat stroke are in terms of counts per million people. Mild injuries do not require overnight hospitalization. Injuries require up to three weeks of inpatient care. Severe injuries require more than three weeks of inpatient care. Home: all locations on the property. Workplace (construction): road construction sites, factories, manufacturing, etc. Workplace (field): fields, forests, oceans, rivers, etc. (only if agricultural, livestock, or fishery operations are conducted). Education: kindergartens, nursery schools, elementary schools, junior high schools, high schools, vocational schools, universities, etc. Public (inside): indoor areas where people enter and exit (theaters, concert halls, restaurants, department stores, hospitals, public bathhouses, train stations (underground platforms), etc.). Public (outside): outdoor portions of places where people enter and exit (stadiums, outdoor parking lots, outdoor concert venues, train stations (outdoor platforms), etc.). Street: general roads, sidewalks, toll roads, highways, etc. Summary statistics for heat stroke are weighted by a region's population in 2019. WBGT $\geq$ 28.

within region r. Conversely, a region issued no warning if $\mathrm{WBGT}_{rt} < 33$ for both the evening-before (5pm) and the morning-of (5am) forecasts across all stations within region r.

If one observed the forecasted WBGT, one could compare region-days with a forecasted WBGT just marginally below the 33.0°C threshold against region-days with a forecasted WBGT just marginally above the threshold in a Regression Discontinuity Design (RDD). This identification strategy would then compare region-days with very similar forecasts that differed within some tight range of forecasted WBGTs, and assume that region-days on opposite sides of the cutoff are otherwise similar to each other. However, as discussed in the previous section, our data do not include station-level forecasts of WBGT, and thus, we do not observe the "running variable" necessary for an RDD.[23]

Our data include realized WBGT and indicators for whether a heat stroke warning was issued due to the evening-before or morning-of forecasts, or both, i.e. whether the forecasted WBGT crossed the threshold or not. With these data, we are able to regress our outcome variables y_{rt} on a dummy variable $Alert_{rt}$, indicating a heat stroke warning in region r at time t, while flexibly controlling for actual WBGT:

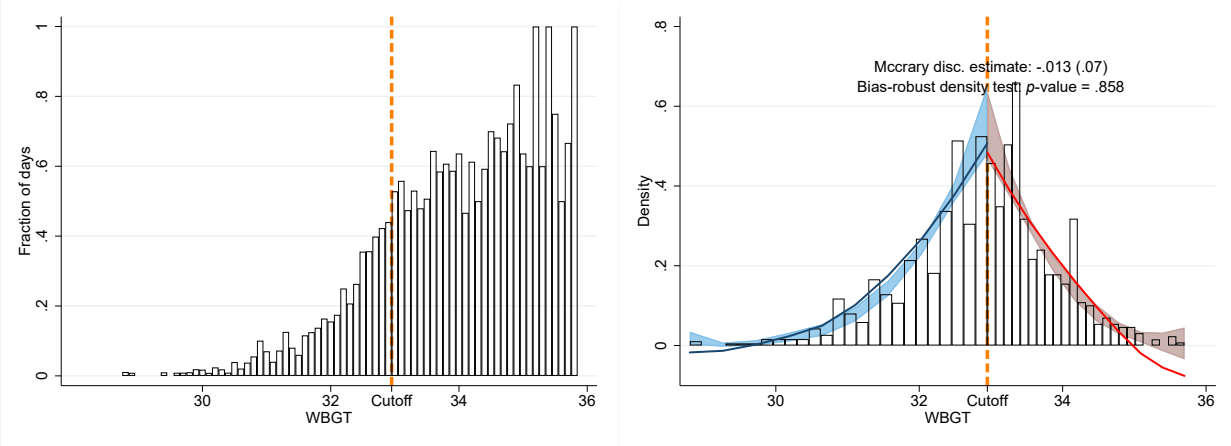
$$y_{rt} = \alpha + \beta Alert_{rt} + \sum_{k=28.0°C}^{K} \gamma_k \mathbb{1}\left(\text{WBGT}_{rt} = k\right) + \delta_r + \lambda_t + \zeta_r year + \epsilon_{rt}, \tag{3}$$

where $\mathbb{1}(\cdot)$ is an indicator function, δ_r are region fixed effects, λ_t are day fixed effects, and ζ_r are region-specific linear time trends. ϵ_{rt} is an idiosyncratic error term. To control for WBGT, we follow the decision rule for the alerts and assign each region-day rt the maximum observed WBGT across stations within that region-day (calculated as WBGT_{rt}), and include dummies for each possible WBGT_{rt}, rounded to the nearest tenth.[24]

We are able to separately estimate β from γ_k since $Alert_{rt}$ does not perfectly correlate with WBGT_{rt}. In other words, we are able to separately disentangle the effect of the heat warning from heat itself on outcomes since (unobserved) forecasted WBGT does not perfectly predict actual (observed) WBGT. In essence, our identification strategy relies on variations in forecasted and actual WBGT (i.e. "forecasting errors"), and whether forecasted WBGT and actual WBGT laid on the opposite sides of the 33.0°C threshold. Because of these forecasting errors, we observe region-days where an alert was issued but actual WBGT was below 33.0°C (i.e. false positive), and region-days where no alert was issued but actual WBGT exceeded 33.0°C (i.e. false negative). This identification strategy relies on the assumption that "forecasting errors" are random and therefore independent from unobservable features of region-days that could potentially affect heat

[23] Moreover, the most concerning confounder for identifying the effect of an alert on heat stroke would be actual WBGT, not forecasted WBGT, since actual weather is what directly impacts individuals. Still, it is possible that individuals make decisions (e.g. plans for tomorrow) based on forecasted WBGT, independent of the alert, and fail to adjust their plans when actual WBGT differs from forecasted WBGT. So long as the forecasting errors are distributed around zero, any such possibility should merely attenuate estimated treatment effects of an alert on outcomes. Also note that forecasted WBGT should theoretically be perfectly collinear with the alert, and thus it is impossible to disentangle the effect of the alert from the effect of forecasted WBGT.

[24] Unreported in equation (3) for simplicity, we similarly flexibly control for WBGT_{rt-1}, the region's previous day's WBGT.



(a) Fraction of region-days with alert

(b) Distribution of false alerts

Figure 3: Alerts and WBGT

Source: Authors' presentation using data from Ministry of Environment. *Note*: Figure (a) shows the share of region-days with an alert among all region-days for all values of WBGT. Figure (b) shows the distribution of "false alerts" for different values of WBGT. Local polynomial approximations on both sides of the cutoff are based on the triangular kernel and local quadratic regressions. Bandwidths for the local polynomial regressions on both sides are 2.35. 95%-confidence intervals are shown as shaded area. $N = 1125$. Data for 2020–2022. In 2020, the sample is restricted to the nine regions that had the heat warning system implemented. WBGT $= [28, 36]$. Results are obtained from the Stata command `rddensity` of Cattaneo et al. (2018).

stroke. Region fixed effects further control for all time-invariant differences across regions, such as average temperature and average heat stroke incidence. Date fixed effects control for all unobserved nationwide trends, including potential changes in average weather or heat stroke incidence. Since the heat stroke warnings are only based on forecasted WBGT, which is a function of actual WBGT and an idiosyncratic error term, controlling for actual WBGT in this panel regression specification allows us to identify the causal effect of the heat warning on our outcomes y_{rt}.

Our identification strategy relies on the fact that regions issue alerts based on the forecasted WBGT across all stations, yet due to forecasting errors, the same region may experience different days with identical WBGTs (e.g. 33.1°C) but some with a warning issued and others without. To illustrate the relationship between WBGT and alerts, in panel (a) of Figure 3, we plot the fraction of region-days that issued an alert by the WBGT for that region-day. We see that around the 33.0°C cutoff, the probability an alert is issued increases from around 45% to around 55%, and slowly increases in WBGT. Still, due to forecasting errors, substantial variation remains: some region-days with maximum WBGT as low as 30.0°C issued alerts, while some region-days with maximum WBGT above 34.0°C failed to issue an alert. If forecasting errors did not exist, then the entire mass of observations would lie above 33.0°C.

The core assumption is that regions do not systematically issue "false" alerts with respect to the 33.0°C threshold. More specifically, we assume that regions are not more or less likely to issue an alert when

WBGT is below 33.0°C relative to failing to issue an alert when WBGT exceeds 33.0°C. This is equivalent to assuming that forecasting errors are random. In panel (b) of Figure 3, we present the distribution of false alerts across actual WBGT.[25] If forecasting errors are independent from actual WBGT, the distribution of false alerts should be smooth around the threshold. Conversely, if regions manipulate the forecasts or the decision rule such that more alerts are issued on either side of the 33.0°C threshold, then we would expect to see a higher share of false alerts just above or below the threshold. To formally test for such a discontinuous jump in the distribution, we perform the manipulation test of Cattaneo et al. (2020) and fail to reject the null hypothesis of no manipulation. Moreover, unsurprisingly, there is a greater mass of false alerts around the 33.0°C cutoff, since this is where even a small forecasting error can result in a false alert. Thus, though we do not observe forecasted WBGT, this evidence suggests that it is highly unlikely that forecasting errors are systematically biased away from zero, or that regions implementing the alerts systematically deviated from the 33.0°C decision rule.

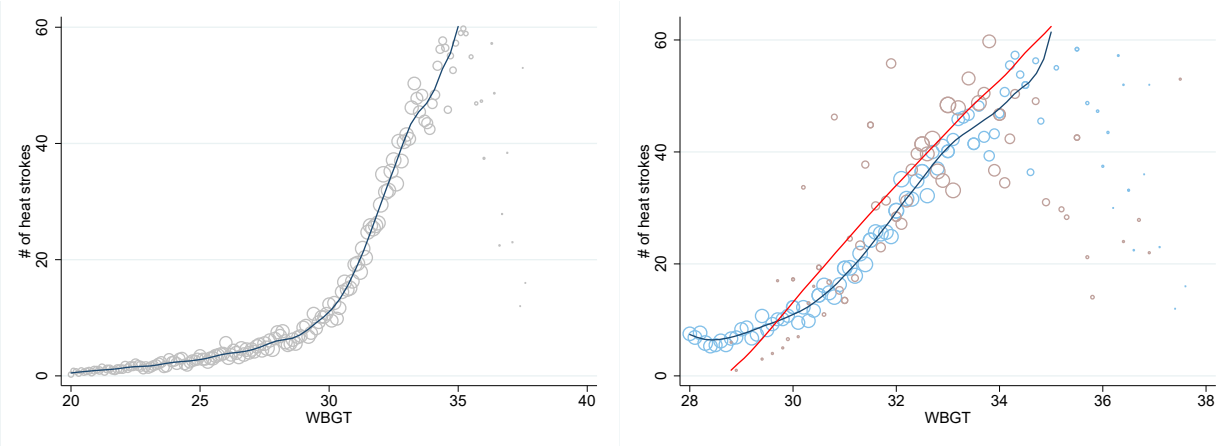
4 Main Results

4.1 Descriptive Evidence

We begin by investigating the graphical relationship between temperature, heat stroke, and heat alerts. In panel (a) of Figure 4, we plot the relationship between WBGT and the number of heat strokes across the full bandwidth of possible WBGTs. The size of the circles indicate the number of region-days that experienced that specific WBGT; for example, WBGTs above 35°C were relatively rare and thus have relatively small circles. Unsurprisingly, the relationship between WBGT and heat stroke is positive, and interestingly, appears to be nonlinear, where the number of heat strokes increases at an increasing rate for higher WBGT.

Then, in panel (b), we plot the relationship between WBGT and heat stroke separately for region-days without an alert (in blue) and with an alert (in red). Immediately, in just the "raw" data, we see that region-days when a heat alert was issued systematically have *more* heat stroke hospitalizations across the full range of WBGTs. This difference is similar for smaller and higher values of WBGT with more region-days experiencing an alert. Thus, it appears that an alert lead to more heat stroke counts in a region. In the next section, we test whether this descriptive result still manifests in our econometric models.

[25] A region-day issued a "false" alert if either (a) an alert was issued but actual WBGT was below 33.0°C ($N = 591$, 46% of all region-days with alert), or (b) an alert was not issued but actual WBGT was greater than 33.0°C ($N = 534$, 9% of all region-days without alert).



(a) WBGT and heat stroke

(b) WBGT and heat stroke by heat alert days

Figure 4: Relationship Between Heat Stroke and Wet Bulb Globe Temperature

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: These graphs show the relationship between the WBGT index and heat stroke. Panel (b) separates the relationship by region-days with a heat alert (in red) vs. without a heat alert (in blue). Local linear regressions are based on a triangular kernel and the rule-of-thumb bandwidth. Size of the circles indicates the number of observations for this point. Data is weighted by a region's population in 2019. WBGT $\geq$ 28.

4.2 Econometric Results

In the next step, we estimate Equation (3) for our full sample in Table 2, where we regress the incidence of heat strokes across region-days in response to a heat alert being issued, also controlling nonparametrically for the maximum region-day WBGT and region-previous-day WBGT. Across columns, we consider combinations of measurements of the outcome variable, regression models, and the sensitivity of the results to region-specific time trends. We cluster standard errors at the region level. Starting with the first two columns, which take the log of region-day heat strokes per million people, we estimate that the incidence of heat strokes increased by nearly 18% for region-days that issued a heat alert. In columns (3) and (4), we estimate the count of heat strokes per million people as the outcome variable to find similarly large effects (i.e. 1.08 additional heat strokes per million people). Since our primary outcome variable is count data (that includes zeros and is overdispersed), perhaps the most appropriate regression model is the negative binomial regression model, which we consider in columns (5) and (6).[26] For our fully specified model in column (6), the expected log count of heat strokes on heat alert days is 0.17 higher than the expected log count for region-days without an alert. Across all specifications, our estimates are statistically significant at the 1%-level. In the remainder of our analyses, we will estimate our full model using negative binomial regressions.

[26]Results do not change if we use a Poisson regression model instead.

Table 2: Effect of Heat Alerts on Heat Strokes

	Log (# of heat strokes +1 per million people)		# of heat strokes per million people		# of heat strokes	
	OLS				Neg. binomial regression	
	(1)	(2)	(3)	(4)	(5)	(6)
= 1 if heat alert	0.190***	0.180***	1.040***	1.078***	0.174***	0.172***
	(0.018)	(0.019)	(0.166)	(0.164)	(0.016)	(0.015)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓
Mean heat strokes	5.071	5.071	5.071	5.071	5.071	5.071
Observations	20612	20612	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from OLS and negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). Regressions in columns (1) to (4) are weighted by a region's population in 2019. WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

To present the heat-health relationship, in Appendix Figure A4, we plot the coefficients from the same-day and previous-day WBGT dummies. Unsurprisingly, higher WBGTs produce more heat strokes (conditional on the previous day's WBGT and on whether a heat alert was issued). We also see some evidence that previous-day WBGT impacts heat stroke, which is consistent with "classical" heat stroke, where the consequences of heat exposure can accumulate over several days. Still, the same-day impact of WBGT on heat stroke is substantially stronger.

4.3 Heterogeneity in Heat Strokes

In Table 3, we consider our full specification while splitting heat stroke counts by different severity diagnosis of the heat stroke. We estimate positive coefficients across the spectrum of severity diagnoses. Strong positive effects are found for mild injuries that do not require hospitalizations and injuries that require some hospitalization. Interestingly, we precisely estimate a large effect for "severe injury" as well, which is associated with at least three weeks of hospitalization. It is important to note, however, that these diagnoses are made on arrival at the hospital and do not necessarily reflect the actual length of hospitalization. In addition, Japanese doctors are very careful with their diagnoses and often keep patients in hospital longer

Table 3: Effect of Heat Alerts on Heat Strokes by Severity Diagnosis

	# of heat strokes		
	(1) Mild injury	(2) Injury	(3) Severe injury
$= 1$ if heat alert	0.167***	0.161***	0.166***
	(0.016)	(0.023)	(0.059)
WBGT	✓	✓	✓
Lagged WBGT	✓	✓	✓
Region FE	✓	✓	✓
Day FE	✓	✓	✓
Region-spec. trend	✓	✓	✓
Mean heat strokes	3.231	1.699	.112
Observations	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes by severity diagnosis in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

than necessary, especially the elderly (Katayama et al., 2021; OECD, 2024). It is therefore not clear whether these "severe cases" actually required such extended care.

In Table 4, we estimate our full specification while splitting the counts of heat strokes across four age groups: children (aged less than seven years), adolescents (aged seven to 17), adults (aged 18 to 64), and elderly (aged 65 and above). We observe large and statistically significant effects for all age groups except for children. Although adolescents are the most resilient to heat, an alert increases heat strokes for this age group as well.[27] However, the largest increase in heat stroke incidence is experienced by adults and the elderly in particular.[28]

In Appendix Table A4, we consider heterogeneity by the location of the heat stroke incident. We find that the positive effects are consistent across all locations, both indoors and outdoors. Though one may expect stronger effects in outdoor locations, these location outcomes do not trace movement (e.g. going from outdoors to indoors) but rather the place of ambulance pick up and are net of any potential avoidance or adverse behavioral effects (i.e. people with varying degrees of risk of heat stroke changing where they spend their day in response to the alert). It is therefore not possible to trace where the actual heat exposure took place.

[27] Because we cannot distinguish exertional and "classic" heat stroke in our data, it is possible that heat strokes are exertional for young people due to strenuous exercise in extreme heat, whereas "classic" heat stroke is more common for older people.

[28] In the online appendix, we present various robustness checks for levels of severity diagnosis and age group.

Table 4: Effect of Heat Alerts on Heat Strokes by Age Group

	# of heat strokes			
	(1) Child	(2) Adolesc.	(3) Adult	(4) Elderly
$= 1$ if heat alert	0.001	0.121***	0.166***	0.181***
	(0.097)	(0.033)	(0.020)	(0.016)
WBGT	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓
Mean heat strokes	.040	.593	1.765	2.672
Observations	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes for different age groups in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

4.4 Intertemporal Effects

Next, we consider several models to test for potential intertemporal effects of the heat alerts. First, given heat strokes can develop over several days, it may be that any gains from avoidance behaviors are not realized until several days in the future. That is, if an alert encouraged avoidance behavior on day t, then perhaps we would see a reduction in heat strokes on day $t+1$ or $t+2$. In these same models, we also consider whether future heat alerts (on day $t+1$ and $t+2$) had an effect on heat strokes on day t. This largely serves as a placebo exercise, though given an alert for $t+1$ could be generated on day t based on the 5pm forecast, it is possible that an alert made for $t+1$ could lead to an "increased reporting" effect on day t. The results from this exercise are presented in Table 5. By and large, we find no evidence of an intertemporal effect of alerts on heat stroke: both past and future heat alerts have almost no effect on heat alerts on day t, with the entirety of the main effects still occurring in response to an alert on day t. These results suggest that any increases in avoidance behavior did not capitalize into a reduction of heat strokes in the future.

We next consider whether there were any dynamic effects based on the issuance of alerts across consecutive days. Zivin and Neidell (2009) evaluate avoidance behavior in response to smog alerts, and find that there was little effect of a smog alert being issued on the second day of consecutive smog alert days.

Table 5: Intertemporal Effect of Heat Alerts on Heat Strokes

	# of heat strokes			
	(1)	(2)	(3)	(4)
$= 1$ if heat alert$_{t-2}$			-0.007	-0.007
			(0.022)	(0.021)
$= 1$ if heat alert$_{t-1}$	0.014	0.014	0.017	0.018
	(0.016)	(0.015)	(0.015)	(0.015)
$= 1$ if heat alert$_{t}$	0.163***	0.162***	0.159***	0.159***
	(0.015)	(0.014)	(0.015)	(0.015)
$= 1$ if heat alert$_{t+1}$	0.011	0.011	0.003	0.004
	(0.017)	(0.017)	(0.017)	(0.017)
$= 1$ if heat alert$_{t+2}$			0.028*	0.029*
			(0.016)	(0.016)
WBGT$_{t-2}$			✓	✓
WBGT$_{t-1}$, WBGT$_{t}$, WBGT$_{t+1}$	✓	✓	✓	✓
WBGT$_{t+2}$			✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend		✓		✓
Mean heat strokes	5.075	5.075	5.075	5.075
Observations	20600	20600	20600	20600

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

In Appendix Table A5, in column (1) we first consider our baseline model with a regressor for an alert on days t and $t-1$ separately. In column (2), we then interact the dummies for the two alert days to find no additional interactive effect for alerts being issued on consecutive days. For completeness, the remaining columns consider further interactions for consecutive-day alerts, including two day gaps, and still we find no evidence of intertemporal effects.

5 Potential Mechanisms

In this section, we consider several exercises in order to unpack the potential mechanisms of our main results. Before moving to the analyses, we first hypothesize several channels of potential responses resulting from the heat warning. First, as is the primary intention of the program, we may expect some type of "avoidance" behavior - in response to an alert, individuals engage in behaviors to avoid or mitigate the

consequences of exposure to high heat. Generally speaking, this would involve the movement of individuals from hotter to cooler areas, or greater usage of appliances to reduce ones temperature (e.g. air conditioning). Of course, this channel cannot explain our main results in Table 2, since such avoidance behavior should mitigate exposure to heat and thus reduce heat stroke. Still, avoidance behavior could exist, yet be outweighed by other opposing factors.

Thus, we further posit a secondary set of behavioral responses that may emerge due to a heat alert and could plausibly explain our findings. The first we call an "increased reporting" effect. Given the alert is explicitly about heat stroke health, it may be that the alert raises one's awareness of the potential presence of heat stroke. That is, in the absence of an alert, some people who experience a heat stroke may not realize it or may fail to consider the seriousness of their symptoms - subsequently, their heat stroke case goes unreported and uncounted. This would be consistent with a health literature that suggests that the damages from high heat are significantly under-counted and unreported (Ebi et al., 2021) and would have serious implications for the climate-health relationship estimated in other studies (Mullins and White, 2020; Neidell et al., 2021; He and Tanaka, 2023). Then, in the presence of the alert, individuals may now visit the hospital and report their heat stroke or pay more attention to other people's symptoms.

The second channel we hypothesize is a "substitution" effect in how health cases are diagnosed. Heat stroke is generally brought upon by high heat, while high heat also substantially increases the risk of other sudden respiratory and cardiovascular diseases (e.g. Michelozzi et al., 2009). In the United States, the EPA estimates that approximately one quarter of heat-related deaths are tied to cardiovascular disease interacting with high heat (United States Environmental Protection Agency, 2023). Given these important interactions, we posit that heat stroke diagnoses can potentially be conflated with other sudden illnesses when patients are transported to a hospital. Then, in the presence of the heat alert, the identification of heat-related symptoms by doctors and hospitals lead to "substitutions" away from other related sudden illnesses and into heat stroke. Furthermore, because doctors fill out a form providing information about the ambulance transport (i.e. cause or reason for the transport, severity, location of incidence, ...) upon arrival at the hospital, these on-the-spot decisions could also lead to substitution away from other categories to sudden illnesses.[29]

Note that with both of these channels ("increased reporting" and "substitution"), the *actual* incidence of heat stroke is unchanged (all else equal), while *reported* heat stroke increases. Thus, it is possible that even in the presence of avoidance behavior (which subsequently reduces actual heat stroke), the total count of reported heat strokes increases in response to the alert, if the "increased reporting" and "substitution" effects

[29] Anecdotal evidence shows that these decisions may also reflect the attitudes of doctors. For example, in the case of work-related accidents, managers and companies are held responsible for the victim's compensation, while in the case of a sudden illness, the medical expenses need to be covered by the victim itself.

outweigh reductions in actual heat stroke from any "avoidance" behaviors.

Lastly, we consider the possibility of an "adverse" behavioral response - the heat alert may be associated with other changing factors which then induce individuals to engage in behaviors that place them at a greater risk of heat stroke. For example, it may be that on days when a warning was issued, workplaces closed more, and perhaps individuals are better protected from high heat at their workplace relative to their home (perhaps due to better air conditioning).[30] If people were less likely to go into work, then they may have engaged in other adverse behaviors with their newfound free time as well. Moreover, He and Tanaka (2023) found increased mortality in high heat months after the Fukushima accident, which they attribute to a behavioral response from households saving energy (where electricity prices increased). In our context, it may be that households associate energy savings and electricity prices with the heat alert, and thus perhaps individuals engage in adverse behaviors in order to save energy on (what they perceive to be) a high heat day. In general, forecasting errors in our setting - when actual WBGT does not coincide with forecasted WBGT (and thus, the heat warning) - may lead people to engage in behaviors placing them at a greater risk of heat stroke (and as highlighted by Shrader et al. (2023), inaccurate forecasts contribute to a significant share of annual deaths in the United States).

In an effort to identify these potential mechanisms, our next analyses consider (a) searches on Google Trends, (b) foot traffic outcomes as measured through Google Mobility Reports, (c) other health and mortality outcomes using the population of ambulance records, (d) energy consumption behavior, and (e) evidence from our own survey.

5.1 Google Trends

We start with Google Trends to see whether region-day level searches for specific terms responded to an alert. In Table 6, using the 2020 to 2022 sample, we consider searches for the terms (Japanese equivalent) "heat stroke," "heat stroke alert," "weather," "temperature," and "air conditioner." We find large increases in searches for "heat stroke" and "heat stroke alert," of around 28% and 44% relative to the mean, highlighting people's increased awareness of the warning and for the primary outcome of interest. Interestingly, we see a significant increase in searches for the word "temperature" and a significant decrease in searches for "weather," suggesting perhaps a substitution away from a generic term ("weather") in favor of becoming more temperature-conscious on days when an alert was issued.[31] Further, we observe a strong increase in

[30] Other interactive effects could occur too: for example, an alert could induce people to go into work to receive the air conditioning ("avoidance" behavior), but then perhaps an increased exposure to peers increases the likelihood that one reports another as suffering from a heat stroke ("increased reporting").

[31] Anecdotal evidence also suggests that it is common to search for "weather" when people are interested in precipitation. This could further explain the strong decrease in search frequency if an alert has been issued and rain is unlikely.

Table 6: Google Trends – Effect of Heat Alerts on Searches Related to Awareness

	(1) Heat stroke	(2) Heat stroke alert	(3) Weather	(4) Temperature	(5) Air conditioner
$= 1$ if heat alert	5.340***	3.705***	−4.385***	3.834***	2.168***
	(0.992)	(1.319)	(0.654)	(0.914)	(0.771)
WBGT	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓	✓
Mean outcome	19.100	8.364	53.125	23.392	37.311
Observations	10647	10647	10647	10647	10647

Source: Authors' calculations using data from Google Trends and Ministry of Environment. *Note*: This table shows results from OLS regressions of Google Trends Searches for 熱中症 (heat stroke), 熱中症警戒アラート (heat stroke alert), 天気 (weather), 気温 (atmospheric temperature), and エアコン (air conditioner), respectively, in a region on the presence of a heat alert and a set of control variables. The mean outcome is for regions without heat alerts (control group). All regressions are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2020–2022. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

searches for the term "air conditioner," highlighting again people's awareness of heat on a day with an alert.

Next, in Table 7, we consider whether searches for activities related to potential avoidance behavior responded to a heat alert, since searches could be indicative of intentions for that day's activities (e.g. a reduction in searches for outdoor activities could indicate intentions to stay at home more). We see a small increase in searches for the outdoor activity "sea bathing" in response to an alert being issued, suggesting perhaps people were more likely to combat high heat by going to the sea on (what they perceive to be) a high heat day. Still, this result should be interpreted cautiously since an alert could just be reflective of a change in the composition of Google Search users for that region-day, and searches do not necessarily reflect actual actions (e.g. an alert piques an individual's curiosity as to whether a nearby beach is closed for the day). This action may therefore constitute "avoidance behavior" in that options such as "sea bathing" may better protect some individuals from heat relative to their own homes. For other search terms, however, we find no statistically significant coefficient. These activities, whether indoor or outdoor, seem unaffected by the alerts. In Appendix Figure A5, we plot the coefficients on the dummies for maximum daily WBGT for each Google Search term. The relationship is positive for heat stroke, temperature, and air conditioner, suggesting that these searches are responsive to heat itself.[32]

[32]It should be noted that older people, who are at most risk of heat stroke, do usually not use search engines. Thus, we selected activities that are most frequently used by young people or adults. Nevertheless, this exercise is indicative of people's behavioral response to an alert in general.

Table 7: Google Trends – Effect of Heat Alerts on Searches Related to Avoidance Behavior

	Outdoor			Indoor		
	(1) Outdoor	(2) Sea bathing	(3) Park	(4) Indoor	(5) Cinema	(6) Karaoke
$= 1$ if heat alert	0.668	2.793*	0.691	1.407	0.100	−0.059
	(0.659)	(1.517)	(0.501)	(1.440)	(0.869)	(0.656)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓	✓	✓
Mean outcome	37.334	16.212	33.776	21.114	40.102	47.789
Observations	10647	10647	10647	10647	10647	10647

Source: Authors' calculations using data from Google Trends and Ministry of Environment. *Note*: This table shows results from OLS regressions of Google Trends Searches for 屋外 (outdoor), 海水浴 (sea bathing), 公園 (park), 屋内 (indoor), 映画館 (cinema), and カラオケ (karaoke), respectively, in a region on the presence of a heat alert and a set of control variables. The mean outcome is for regions without heat alerts (control group). All regressions are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2020–2022. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

5.2 Google COVID-19 Community Mobility Reports

Next, we turn to the Google Mobility Reports data to investigate how household behavior responded to heat alerts. The results are presented in Table 8. Recall that the outcome is measured as the percentage change in visits to a specific place (e.g. "Retail & recreation") relative to pre-Covid;[33] for example, looking at the summary statistics of the outcome variables in Appendix Table A3, people spent more time at home and at grocery stores and pharmacies and less time at recreational places, parks, transit stations, and workplaces relative to pre-Covid.

We do not find any change in people's visits in response to an alert except for a reduction in time spent at home. Because this reduction in time at home cannot be explained by an increase in visits to any other place that we observe, it is likely that this reduction is purely mechanical due to people visiting the hospital because of heat stroke. In Appendix Figure A6, we plot the coefficients from the dummies for daily WBGT for each location category - interestingly, we find that people are more likely to leave their home residence in response to a higher WBGT (conditional on an alert being issued), which again could be due to having a heat stroke and thus going to the hospital.

[33] In particular, the reports show percentage changes relative to a reference day, which is the median of the corresponding weekday in the five weeks preceding the COVID-19 pandemic (i.e. January 3 to February 6).

Table 8: Google Mobility Reports – Effect of Heat Alerts on Mobility

	Percentage change in visits					
	(1) Retail & recreation	(2) Grocery & pharmacy	(3) Parks	(4) Transit stations	(5) Workplaces	(6) Residential
$= 1$ if heat alert	0.122	0.028	0.220	0.499	0.258	-0.245^{***}
	(0.313)	(0.129)	(0.717)	(0.531)	(0.160)	(0.066)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓	✓	✓
Mean outcome	−8.128	3.905	−4.211	−16.434	−12.211	6.223
Observations	10647	10647	10293	10636	10647	10647

Source: Authors' calculations using data from Google COVID-19 Community Mobility Reports and Ministry of Environment. *Note*: This table shows results from OLS regressions of the percentage changes in number of visits to specific places in a region on the presence of a heat alert and a set of control variables. The mean outcome is for regions without heat alerts (control group). All regressions are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2020–2022. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

5.3 Ambulance Records – Mortality and Substitution

We next turn to the ambulance transports data, which include information on the population of ambulance transportations made in Japan from 2017 to 2021. One possible alternative explanation of our main results is that there's a "substitution" effect in how health cases are diagnosed. Heat stroke is generally brought upon by high heat, while high heat also substantially increases the risk of other sudden respiratory and cardiovascular diseases (e.g. Michelozzi et al., 2009). In the United States, the EPA estimates that approximately one quarter of heat-related deaths are tied to cardiovascular disease interacting with high heat (United States Environmental Protection Agency, 2023). Given these important interactions, we posit that heat stroke diagnoses can potentially be conflated with other sudden illnesses when patients are transported to a hospital. In particular, the diagnosis of exertional heat stroke is often complicated by overlapping symptoms such as hyponatraemia, hypotension, and hypoglycaemia, or inaccurately measured body temperature due to immediate treatment (Bouchama et al., 2022).[34] Then, in the presence of the heat alert, the identification of heat-related symptoms by doctors and hospitals lead to "substitutions" away from other related sudden illnesses and into heat stroke.

[34]Diagnosis of "classic" heat stroke, however, is usually straightforward.

Table 9: Effect of Heat Alerts on Ambulance Transports

	# of ambulance transports							
	All		Mild		Severe		Deaths	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$= 1$ if heat alert	0.015***	0.014***	0.015***	0.013***	0.019*	0.023**	0.005	0.015
	(0.004)	(0.004)	(0.004)	(0.005)	(0.011)	(0.012)	(0.033)	(0.030)
WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓		✓
Mean # of transports	130.301	130.301	119.806	119.806	9.228	9.228	1.267	1.267
Observations	16879	16879	16879	16879	16879	16879	16879	16879

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of ambulance transports by severity diagnosis in a region on the presence of a heat alert and a set of control variables. The mean number of transports per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Data for 2017–2021. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

In Table 9 we estimate our Equation (3) for all ambulance transportations using negative binomial regressions in columns (1) and (2), then split by severity diagnosis in the remaining columns. Similar to our main results, we estimate that heat alerts led to increases in ambulance usage, primarily due to mild and severe cases. These results suggest that the main effects cannot be driven solely by a "substitution" effect, where the increase in heat strokes comes from a reduction in ambulance usage stemming from other sudden illnesses.

In the final two columns of Table 9, we find that the heat alert did not have a statistically significant effect on mortality (involving an ambulance call). Again this suggests that if there were any avoidance behaviors, the benefits were not capitalized into reduced mortality. They also show that the increase in heat strokes did not result in a statistically significant increase in mortality - concern remains however on the size of the effect, which is positive and comparable in magnitude to the effect sizes for mild and severe cases. We conclude that at a minimum, this is no evidence that the heat alerts led to reductions in mortality.[35]

We can further utilize the ambulance records to identify the effect of heat alerts on patient outcomes by the "cause" or reason for the ambulance transport. We consider six categories: hospital transfer (an ambulance transferring a patient from one hospital to another), violence and self-harm, sudden and general illness

[35]Because these fatalities are ambulance transports, an increase could still be indicative of increased reporting of cases that would have occurred without the use of an ambulance.



(a) Percentage change

(b) Absolute change

Figure 5: Effect of Heat Alerts on Ambulance Transports

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: Panel (a) shows results from negative binomial regressions of the number of ambulance transports by severity diagnosis in a region on the presence of a heat alert and a set of control variables. Panel (b) shows these results multiplied by the daily mean number of ambulance transports per million people for regions without heat alerts (control group). WBGT $\geq$ 28. Data for 2017–2021. 95%-confidence intervals are based on robust standard errors clustered at the regional level.

(which includes heat strokes), work and sport, traffic accident, and fire and water and natural disaster.[36] In Figure 5, we plot the coefficients from estimating our full specification for each of these categories, while also considering subgroups by severity diagnosis (mild, severe, death). In panel (a), we plot the coefficients and their 95%-confidence intervals directly from our main equation, while panel (b) scales the coefficients by the daily mean number of ambulance transports per million people. Reassuringly, we find that the positive effect of heat alerts on ambulance usage is driven solely by sudden and general illnesses (including heat stroke), whereas the remaining five categories (which effectively act as placebo tests) all display statistically insignificant effects. When split by the three categories of severity diagnosis, only one of the 15 placebo estimates (mortality due to fire and water and natural disasters) displays statistical significance, which can likely be attributed to general randomness and is very small in size. This implies that substitution within the sudden illnesses category is unlikely and cannot explain the results. We also do not find reductions in other accidents and illnesses that could be substituted for sudden illnesses. This is also further evidence against the "adverse" behavior hypothesis, as a change in people's behavior would be reflected in a change in hospitalizations, for instance related to drowning while swimming in the sea.

[36] The raw FDMA data provide finer categories than these six groupings. For example, sudden illness is separately reported from general illness, and accidents due to fire, water, and natural disasters are separated from each other. We group sudden illnesses with general illnesses since heat strokes could fall into either of these categories, and in these data we otherwise cannot distinguish whether a specific ambulance transport was directly related to a heat stroke. We group fire, water, and natural disasters together due to the nature of their causes and to increase their overall counts.

In the next step, we analyze heterogeneous responses based on income levels. The ambulance records data comes at the emergency response unit, which covers several municipalities.[37] By merging data on taxable income per capita in 2020 on the municipality level, we can distinguish low-income and high-income emergency response units and thus perform the same analysis as before at this level instead. In Appendix Tables A6 and A7 we find that the effect of a heat alert is three times as large for low-income as for high-income neighborhoods. It even seems that the effect on severe diagnoses is solely driven by low-income neighborhoods with effect sizes of up to 6.4%. This finding has serious implications for the health consequences of climate change. Low-income households may be less informed about the risk of heat stroke and therefore react stronger to the information provided by the alert, while high-income households are well informed and know how to protect themselves from high heat.

5.4 Compliance Behavior

The question remains whether people follow the policy's recommendation and engage in any avoidance behavior that reduces the risk of heat stroke. Although we found a strong increase in Google searches for air conditioning, it is unclear whether people actually turned up air conditioning in response to an alert. To provide some evidence of the actual energy consumption response that is most likely due to increased air conditioning (which is part of the policy's recommendation), we utilize data from the Japan Renewable Energy Foundation. These data contain energy consumption from various sources (e.g. nuclear energy, fossil fuels, solar energy, wind energy). Though these data are provided at the hourly level, they unfortunately are aggregated to the energy-grid level, which comprises nine regions, each covering several prefectures (i.e. our level of analysis).[38] Thus, this analysis suffers from having relatively few clusters, and the data are not well suited to unmask potential heterogeneities in energy consumption by, for example, household type. Still, examining aggregate consumption could reveal potential avoidance responses across Japan. We assign the consumption level for each region-day to the energy-grid level and adjust our regressions to ensure that regions covering more prefectures are appropriately weighted. The results are presented in Appendix Table A8, and in general we find a weak overall increase in energy consumption of around 3% on days when an alert was issued (conditional on same day and prior day WBGT), suggesting evidence of avoidance behavior. Because air conditioning is the largest share of households' overall energy consumption (around 30%) (Reuters, 2023), especially during the summer, it is likely that this increase is due to people spending more time at home and using more air conditioning as response to a heat alert. This is also reflected by the

[37] Our sample contains 752 emergency response units covering 1,916 municipalities, and thus on average 2.5 municipalities per emergency response unit.

[38] The regions are Hokkaido, Tohoku, Kanto, Hokuriku, Chubu, Kansai, Chugoku, Shikoku, and Kyushu & Okinawa.

stronger response during daytime compared to nighttime, when heat is a bigger problem.

5.5 Survey Evidence

To further complement our analyses into potential mechanisms, we conducted a survey in which we asked more than 200 adults between the ages of 18 and 66 from all parts of Japan and from different socioeconomic backgrounds about their knowledge of the alert system, experience of heat stroke, and behavioral responses to an alert. After dropping invalid responses and conducting validity checks, we are left with 192 observations. Individual background characteristics are further described in the online appendix.

Using Google Trends data, we found that alerts have increased the awareness of the risk of heat stroke and as such "increased reporting" serves as the most likely explanation for the increase in heat stroke counts in response to an alert. According to results from our survey presented in Figure 6, 39% have at least heard of the heat index WBGT, 27% have suffered from a heat stroke, 53% know someone who has suffered from a heat stroke, and 63% have experienced a heat stroke alert. These results show not only that the majority of people are aware of the risk of heat stroke from their own experience or acquaintances, but also that they are familiar with the warning system and have experienced a heat stroke alert themselves. As such, our awareness hypothesis seems to be confirmed through the survey, even though we do not observe the counterfactual in absence of the warning system.

Because most people have experienced a heat stroke alert in their region, we further asked about their behavioral response to an alert. Results are presented in Figure 7. Most people stayed inside ($n = 76$), turned up the air conditioning ($n = 40$), and drank more water than usual ($n = 86$), which is evidence of avoidance behavior that should reduce the risk of heat stroke and is in line with the policy's recommendation and previous results (see Section 5.4). Interestingly, 45 people responded to watch out for their own and other people's physical condition, further supporting our awareness and "increased reporting" hypothesis. 11 people even went to the hospital after having symptoms and 6 people encouraged others to go the hospital. Finally, any "adverse" behavior that might explain our results can also be dismissed by these responses, as only a small fraction of people went outside.

6 Welfare Implications

Our results show a strong increase in heat stroke counts on days with a heat stroke alert. Because we find no evidence of "adverse" behavioral responses or substitution away from other related sudden illnesses, this increase can only be explained by an increased awareness and thus reporting of previously unidentified heat

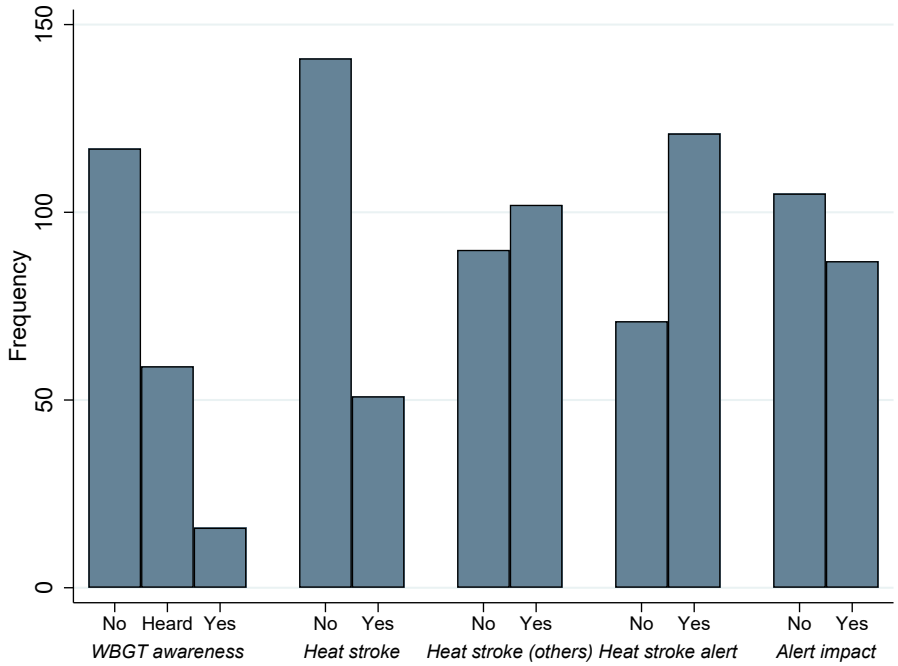


Figure 6: Survey Response: Awareness & Heat Stroke Experiences

Source: Own survey data. *Note*: The questions (in Japanese) are: WBGT awareness: "Do you know what WBGT means?" Heat Stroke: "Have you ever suffered from a heat stroke before?" Heat stroke (others): "Do you know someone who has ever suffered from a heat stroke before?" Heat stroke alert: "Have you ever experienced a *Heat Stroke Alert* in your region?" Alert impact: "Did the *Heat Stroke Alerts* affect your behavior in any way?" $n = 192$.

strokes. However, there are still people who are complying with the policy and as such potentially reducing the risk of heat stroke. As the *actual* number of heat strokes potentially decreases due to partial compliance with the policy's recommendations, this means that the increased reporting effect is a lower limit to the number of cases that would go unidentified in absence of the warning system.

To obtain an estimate of the number of heat strokes that would have gone undetected in the absence of the warning system, we perform a back-of-the-envelope calculation for region-days with a WBGT above $33°C$. First, we divide the actual number of heat strokes on a specific day by 1 + (0.172 $\times$ share of regions with a heat alert).[39] We then compare this counterfactual estimate with the actual number of heat strokes in Figure 8. We find that the warning system detected more than 2,550 additional heat strokes in 2021 and 2022, or 10.2% of all heat strokes recorded on these high-heat days. In other words, more than one in ten cases of heat stroke would have gone unidentified in absence of the warning system. Because only 57% of the region-days with a WBGT above $33°C$ issued an alert due to "forecasting errors," the number of detected heat strokes could be even higher if the warning system was based on more accurate forecasts (Shrader et al., 2023). If we perform this back-of-the-envelope calculation for low-income and high-income

[39]The estimate 0.172 is based on column (6) in Table 2.

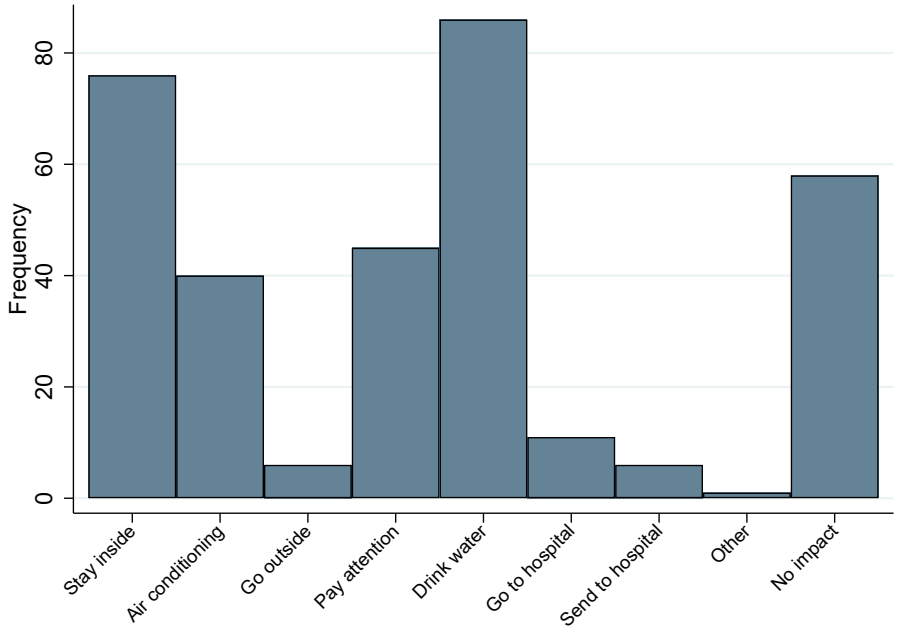


Figure 7: Survey Response: Behavioral Response to a Heat Stroke Alert

Source: Own survey data. *Note*: The question (in Japanese) is: "If yes [Did the *Heat Stroke Alerts* affect your behavior in any way?], how did you react? Please choose the options that describe all your possible actions best." $n = 192$.

neighborhoods separately using the ambulance data for 2021 instead and scale it by the population size, we find that around four times as many heat strokes were detected in low-income compared to high-income neighborhoods. While in low-income neighborhoods 16.6 heat strokes per million people were detected in 2021, the figure was only 4.3 for high-income regions. This again has serious implications for environmental justice concerns (Bakkensen et al., 2024; Colmer et al., 2024) because low-income neighborhoods seem to be less-informed and potentially more vulnerable to the consequences of climate change.

If heat stroke is not treated properly, within only a few hours, serious health problems can arise, which can even lead to multiple organ damage or death (Bouchama et al., 2022). Therefore, if the increase in the number of heat strokes after a warning is only due to reporting, we should see an improvement in overall health outcomes a few days after the warning. On the other hand, overuse of ambulance for mild cases that do not require treatment could crowd out other illnesses and bind important resources, especially during the COVID-19 pandemic when hospital were running at their capacity limit (Kokudo and Sugiyama, 2021). It is therefore unclear whether the treatment of (mild) unidentified cases has positive welfare implications for the health of the population or whether the negative effects of an overcrowded health care system outweigh these benefits.

To test for any intertemporal effect, we conduct an event-study type approach in which we regress

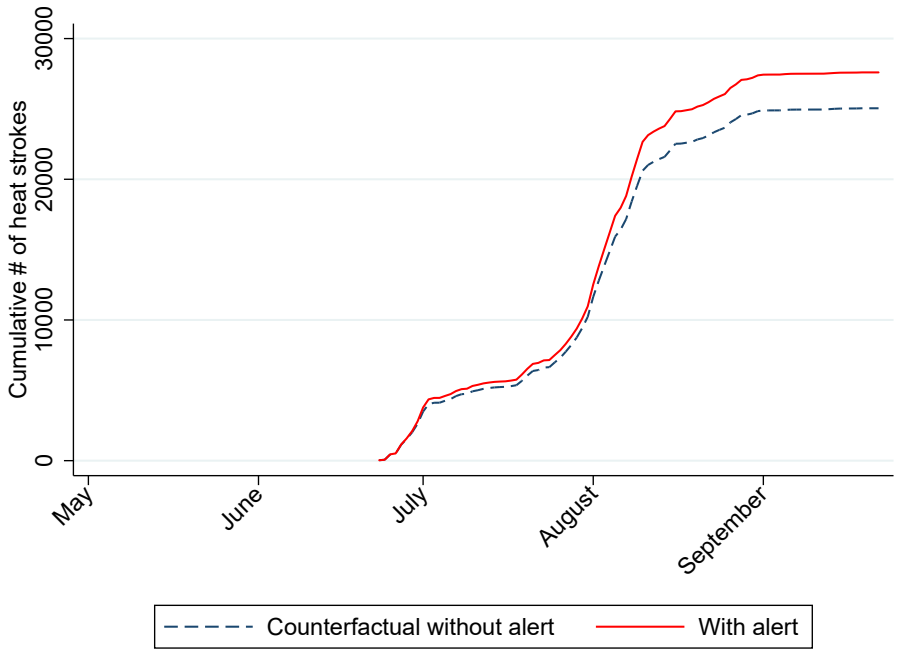


Figure 8: Counterfactual Evolution of Heat Stroke for WBGT $\geq 33°C$

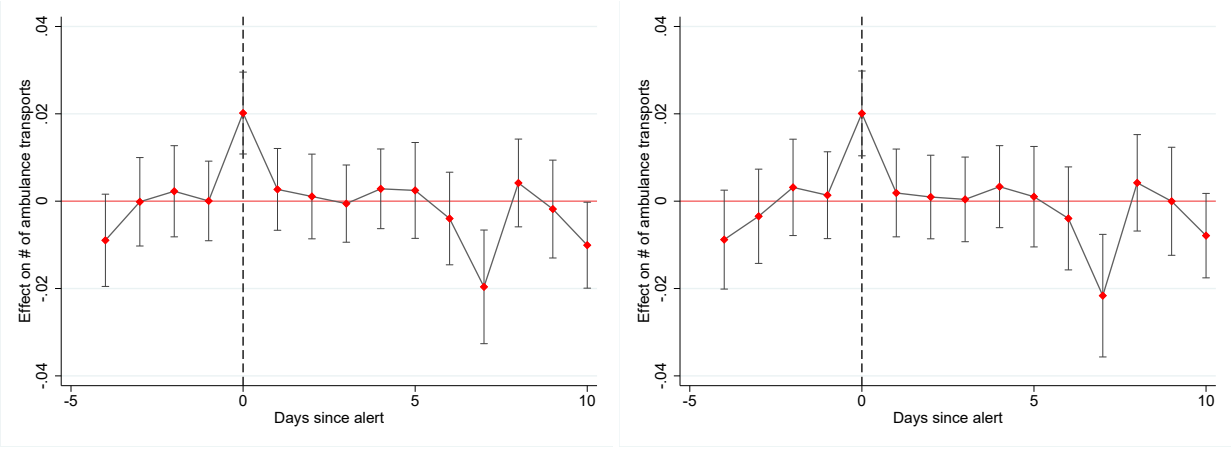
Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This graph shows the cumulative number of heat strokes in 2021 and 2022 (in red) and the counterfactual evolution of this time series in absence of the alert system. The counterfactual estimate is obtained by dividing the actual number of heat strokes on a specific day by $1 + (0.172 \times \text{share of regions with a heat alert})$. Only region-days with a WBGT above $33°C$ are considered.

ambulance transports on leads and lags of the heat alert. While leads mostly serve as placebo check, we expect a spike in ambulance transports on the day of an alert and a reduction in ambulance transports for some days after the alert if the positive effects of treatment outweigh the potential negative effects of the crowding out. Results from Figure 9 clearly show an increase in ambulance transports on the day of the alert. While the effect remains stable at zero for most days after the alert, there is a strong decrease on the 7th day and a declining trend around the 10th day. This pattern is mostly due to mild cases, although the tendency remains for severe cases as well.[40] All in all, it seems that a heat alert does not lead to more ambulance transports a few days later, but rather fewer, because people receive proper treatment when heat stroke is detected.[41]

If the increase in heat strokes is solely due to reporting, the question arises as to whether patients diagnosed with more than three weeks of hospitalization (i.e. severe cases) would really have gone undetected in absence of the warning system. Japan is known for keeping their patients for extremely long periods in

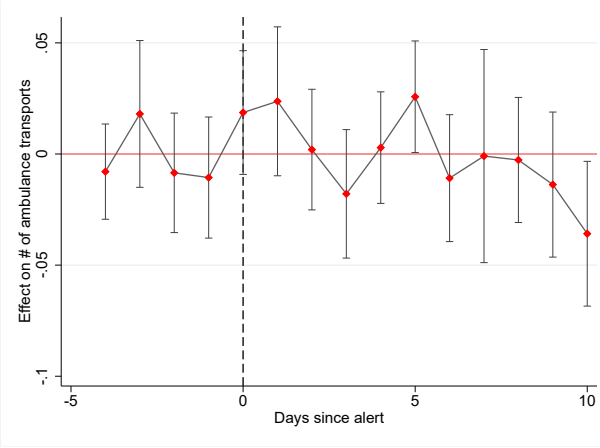
[40] The jump in severe cases on the fifth day after an alert may be explained by the corresponding decrease in deaths, which again indicates an improvement in general health.

[41] If we separate all ambulance transports by "cause" or reason, we do not find any increase that can be attributed to a potential crowding out effect (not reported).

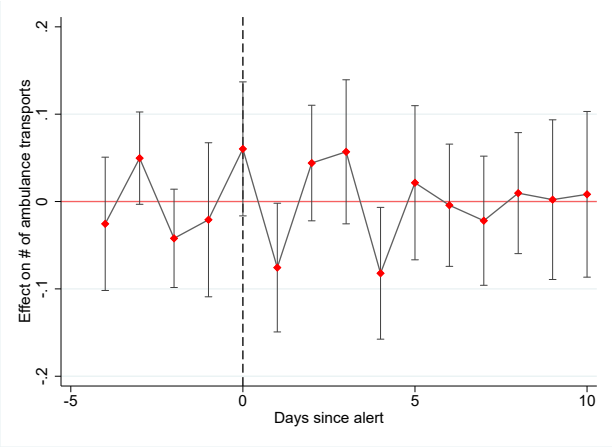


(a) All

(b) Mild



(c) Severe



(d) Death

Figure 9: Intertemporal Effects of Heat Alerts on Ambulance Transports

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: These graphs show results from negative binomial regressions of the number of ambulance transports by severity diagnosis in a region on the presence of a heat alert, its leads and lags, and a set of control variables. WBGT $\geq$ 28. Data for 2017–2021. 95%-confidence intervals are based on robust standard errors clustered at the regional level.

the hospital, even after initial treatment has been completed. According to the OECD, the average length of stay in Japan is 16 days, compared to the OECD average of around 6 days (OECD, 2024). It is likely that doctors are playing it safe and "overdiagnose" patients for various reasons. First, Japan is a risk-averse country in the sense that uncertainties are avoided through more thorough medical examination. Second, since around 70% of hospitals are privately owned, they have a financial interest in keeping patients longer than necessary.[42] Finally, when considering Japan's COVID-19 policies, the online appendix shows that the

[42]There is also suggestive evidence that mothers are kept in hospital for around two weeks after birth. Further, there is a lot of pressure from families to keep older patients in the hospital longer to do all kinds of examinations. Finally, a large proportion of a hospital's profits come from long-term patients.

effect of an alert during the state of emergency completely vanishes. This suggests that doctors diagnose severe cases as mild or regular cases only in times of hospital capacity constraints.[43] Further, we find no evidence of negative intertemporal effects of these changed diagnoses (not reported), i.e. no intertemporal interaction effect of the alert and state of emergency, suggesting that patients receive proper treatment, even if they only stay for a short period.

7 Robustness Checks and Placebo Exercises

In this section, we consider a series of robustness checks and placebo exercises. First, due to the scarcity of alerts being issued on days with low temperatures, we restricted our main sample to region-days with a WBGT above 28.0°C. We test the robustness of our results in Appendix Table A9 by estimating our main specification across different sample cutoffs in WBGT. Across all sample considerations, the estimated coefficients remain consistent and statistically significant at the 1%-level. Furthermore, results remain unchanged if we do not restrict the sample to region-days with heat stroke counts below the 99^{th} percentile (see Appendix Table A10).

As an alternative econometric robustness check, particularly given a relatively low number of region clusters ($n = 44$), we conduct a randomization inference test (Bertrand et al., 2004). To do so, we randomly reassign each region's treatment pattern (heat alerts) across regions, and estimate our main specification. We repeat this procedure 1,000 times, then plot the resulting placebo distribution of treatment effects alongside the actual treatment effect estimated in column (6) from Table 2. The results from this exercise are displayed in Appendix Figure A7, and once again we see that the primary treatment effect is highly statistically significant.

Because people can move across regional borders, an alert in Tokyo, for instance, could have a very different impact on a person living in Chiba and commuting to Tokyo than on a person living and working in Tokyo. To test for potential commuting behavior driving our results, we exclude the three major commuting zones in Japan from our analysis and present results in Appendix Table A12. Whether we drop the Greater Tokyo area with Saitama, Chiba, Tokyo, and Kanagawa, the Chukyo metropolitan area with Gifu, Aichi, and Mie, or the Keihanshin metropolitan area with Kyoto, Osaka, and Kobe, the estimates remain unchanged and thus we conclude that commuting behavior does not affect our results. This is also confirmed by results from Appendix Table A13 in which we run regressions separately for weekdays and weekends, not finding any significant difference.

[43] Guidetti et al. (2021) emphasize the importance of considering hospital capacity constraints when estimating the impacts of natural hazards on health outcomes.

Next, we consider whether the source of the alert (previous-day 5pm forecast vs. same-day 5am forecast) had any differential bearing on heat strokes. In Appendix Table A11, we estimate our main specification separately for alerts issued on the same day (5am) and the previous day (5pm). Note that an alert from 5pm the previous day was never "retracted" if the 5am same day WBGT forecast was below $33.0°C$. Therefore, the model for the same-day alert corresponds to our main model with an alert at any time. Meanwhile, only 4.6% of all region-days had an alert issued at 5am even though there was no alert from 5pm the previous day. In columns (1) to (4), we find that there is virtually no difference if we code the alert as stemming from 5am or 5pm. If we allow for interactions of both alerts, we find that around two thirds of the effect is driven by the same-day alert and only one third stems from the previous-day's alert that is confirmed on the same day.

Because the introduction of the heat warning system coincided with the COVID-19 pandemic, it is possible that some of our results are due to policies affecting people's daily lives that coincided with heat alerts (e.g. movement restrictions). As a measure to prevent the spread of the virus, Japan's Prime Minister Shinzo Abe declared a state of emergency for Tokyo and six other regions on April 7, 2020. Although this measure was regional and temporary, it was extended to the entire country on April 16 and lasted until May 25. In the following two years, the Japanese government proclaimed four more states of emergency for various regions.[44] Although this measure did not restrict people's movement directly, people were asked to stay at home, and businesses as well as public facilities such as schools reduced operating hours or shut down completely. To test whether those measures coincided with heat alerts and thus drive our results, in Appendix Table A14, we estimate model (3) and include a dummy of whether a region-day was in a state of emergency. The results show that this policy did neither affect heat stroke nor the effect of a heat alert on heat stroke.[45]

As an additional robustness check, we consider alternative fixed effects as used in our main specification. In Appendix Table A15, we present results from models that control for region-month-year fixed effects or region-WBGT fixed effects in addition to day fixed effects, where WBGT was rounded to the whole number. Again, the results remain virtually unchanged.

In essence, our main specification corresponds to a two-way fixed effects regression in which units switch on treatment at different points in time for exactly one period. A recent literature has shown that in such staggered Difference-in-Differences (DiD) settings standard two-way fixed effects regressions are problematic in presence of treatment effect heterogeneity over time (Borusyak and Jaravel, 2018; De Chaise-

[44] In the online appendix, we present the timeline of these policies across prefectures.

[45] In the online appendix, we also show that the alerts did not have any effect on the number of Covid infections or deaths and that the effect of an alert on heat stroke is not driven by the number of Covid infections or deaths.

martin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). To address the issues arising from the negative weighting problem, in Appendix Table A16, we estimate our main specification using the alternative estimator proposed by De Chaisemartin and d'Haultfoeuille (2020), which is robust to heterogeneous treatment effects. Coefficients are remarkably similar to our main results (Table 2, columns (1) to (4)) and highly statistically significant, leaving us confident that treatment effect heterogeneity is not a concern in our setting.

Finally, in Appendix Figure A8, we allow the effect to vary with subsamples of different WBGT cutoffs and for subsamples centered around specific WBGT values, respectively. Neither of these figures exhibit any heterogeneity. The results are rather stable around 0.15 to 0.18 and highly statistically significant. In contrast to the descriptive evidence from Figure 4, this suggests that the effect of a heat alert is the same for different temperatures. In particular, there is no differential effect for region-days with actual temperatures below and above the threshold of $33°C$.

8 Conclusion

This study investigates the behavioral and health impacts of the first comprehensive heat-health warning system. The context of our study is Japan, which in 2020, introduced a heat warning system to raise awareness of heat-related illnesses and promote heat stroke prevention measures. Our identification strategy utilizes variation across region-days in whether an alert was issued while flexibly controlling for same-day and previous-day wet bulb globe temperature (WBGT). We document substantial variation, with no manipulation, in whether an alert was issued for a region-day, controlling for observed WBGT. The lack of perfect collinearity comes from unobserved "forecasting errors," where alerts are determined solely by (inaccurate) forecasts.

Our findings suggest an increase in hospitalizations for heat stroke in response to a heat alert being issued. The effect exists across the spectrum of severity diagnoses and location, and is most prevalent for adults and elderly. The effects are large, precisely estimated, and robust to a vast array of model considerations. Unsurprisingly, we further document that same-day and previous-day WBGT adversely impact heat stroke as well.

To unpack the potential mechanisms driving these results, we hypothesize that heat alerts can have various impacts on people's behavior. First, alerts may raise one's awareness of the potential presence and risk of heat stroke (possibly without changing the actual incidence of heat stroke), leading to increased reporting. Second, we hypothesize that heat strokes may conflate with other sudden illnesses, such as cardiovascular disease and stroke, and when an alert is issued, there is effectively a "substitution" in how

diseases are recorded, away from other illnesses and into heat stroke. Finally, an alert may be associated with other changing factors which then induce people to engage in behaviors that place them at a greater risk of heat stroke (e.g. leaving their home).

To identify these potential behavioral responses, we utilize additional data from Google Trends, Google Mobility Reports, the population of ambulance records, energy consumption behavior, and our own survey. In total, we do not find any evidence of "adverse" behavior or "substitution" away from other sudden illnesses. Instead, our evidence suggests increased awareness and people reporting heat strokes that would otherwise have gone undetected. Since we also find evidence of avoidance behavior, the actual number of heat strokes probably decreased due to the warning system, while our estimates represent a lower bound of reported heat strokes that would otherwise have gone undetected. This finding has important implications for the climate-health relationship, as previous studies may have severely underestimated the negative health consequences of extreme heat in absence of a warning system.

Finally, a back-of-the-envelope calculation shows that the warning system detected more than 2,550 additional heat strokes on high-heat days in 2021 and 2022, or one in ten heat strokes would have gone unidentified in absence of the warning system. This effect is strongest for low-income neighborhoods which report four times as many heat strokes (relative to their mean) as high-income neighborhoods in response to an alert. Further, we find no evidence of a crowding out of other illnesses. Instead, we find a reduction in overall ambulance transports around seven days after an alert, most likely due to proper treatment of otherwise unidentified cases.

Though not impacting our study, in response to the hottest summer ever recorded in Japan and a rising number of heat stroke deaths in 2023, the Ministry of Environment announced a special warning system that will apply from October 23, 2024 in addition to the current heat stroke alert system (Japan Today, 2024). In contrast to the current system which only issues behavioral recommendations, the new system will call for "canceling or postponing sports and other events that cannot take sufficient measures to prevent heat strokes" in addition to providing cooling shelters (i.e. libraries and community centers) to the public. Such a warning will be issued when the WBGT exceeds $35.0°C$ in the previous-day 2pm forecast. We welcome further research on the mortality effects of this new policy, since it is at a potentially fatal threshold and can have very different impacts than the current warning system.

References

AKESAKA, M. AND H. SHIGEOKA (2023): ""Invisible killer": Seasonal allergies and accidents," NBER Working Paper 31593.

BAKKENSEN, L. A., L. MA, L. MUEHLENBACHS, AND L. BENITEZ (2024): "Cumulative impacts in environmental justice: Insights from economics and policy," *Regional Science and Urban Economics*, 103993.

BALLESTER, J., M. QUIJAL-ZAMORANO, R. F. MÉNDEZ TURRUBIATES, F. PEGENAUTE, F. R. HERRMANN, J. M. ROBINE, X. BASAGAÑA, C. TONNE, J. M. ANTÓ, AND H. ACHEBAK (2023): "Heat-related mortality in Europe during the summer of 2022," *Nature Medicine*, 1–10.

BARRECA, A., K. CLAY, O. DESCHENES, M. GREENSTONE, AND J. S. SHAPIRO (2016): "Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century," *Journal of Political Economy*, 124, 105–159.

BENMARHNIA, T., L. SCHWARZ, A. NORI-SARMA, AND M. L. BELL (2019): "Quantifying the impact of changing the threshold of New York City heat emergency plan in reducing heat-related illnesses," *Environmental Research Letters*, 14, 114006.

BERTRAND, M., E. DUFLO, AND S. MULLAINATHAN (2004): "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics*, 119, 249–275.

BORUSYAK, K. AND X. JARAVEL (2018): *Revisiting event study designs*, SSRN Scholarly Paper ID 2826228, Social Science Research Network, Rochester, NY.

BOUCHAMA, A., B. ABUYASSIN, C. LEHE, O. LAITANO, O. JAY, F. G. O'CONNOR, AND L. R. LEON (2022): "Classic and exertional heatstroke," *Nature Reviews Disease Primers*, 8, 8.

CATTANEO, M. D., M. JANSSON, AND X. MA (2018): "Manipulation testing based on density discontinuity," *The Stata Journal*, 18, 234–261.

——— (2020): "Simple local polynomial density estimators," *Journal of the American Statistical Association*, 115, 1449–1455.

CENTER FOR DISEASE CONTROL AND PREVENTION (2023): "Heat stress – Heat related illness," https://www.cdc.gov/niosh/topics/heatstress/heatrelillness.html, retrieved July 13, 2023.

CHAE, Y. AND J. PARK (2021): "Analysis on effectiveness of impact based heatwave warning considering severity and likelihood of health impacts in Seoul, Korea," *International Journal of Environmental Research and Public Health*, 18, 2380.

CHETTY, R., J. N. FRIEDMAN, N. HENDREN, M. STEPNER, AND THE OPPORTUNITY INSIGHTS TEAM (2020): "Real-time economics: A new platform to track the impacts of COVID-19 on people, businesses, and communities using private sector data," NBER Working Paper 27431.

COLMER, J. M., S. QIN, J. L. VOORHEIS, AND R. WALKER (2024): "The changing nature of pollution, income, and environmental inequality in the United States," in *AEA Papers and Proceedings*, vol. 114, 41–46.

DE CHAISEMARTIN, C., X. D'HAULTFOEUILLE, AND Y. GUYONVARCH (2019): "DID_MULTIPLEGT: Stata module to estimate sharp Difference-in-Difference designs with multiple groups and periods," Tech. rep., Boston College Department of Economics.

DE CHAISEMARTIN, C. AND X. D'HAULTFOEUILLE (2020): "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 110, 2964–2996.

DESCHÊNES, O. AND M. GREENSTONE (2011): "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US," *American Economic Journal: Applied Economics*, 3, 152–185.

DESCHENES, O. AND E. MORETTI (2009): "Extreme weather events, mortality, and migration," *Review of Economics and Statistics*, 91, 659–681.

D'IPPOLITI, D., P. MICHELOZZI, C. MARINO, F. DE'DONATO, B. MENNE, K. KATSOUYANNI, U. KIRCHMAYER, A. ANALITIS, M. MEDINA-RAMÓN, A. PALDY, ET AL. (2010): "The impact of heat waves on mortality in 9 European cities: Results from the EuroHEAT project," *Environmental Health*, 9, 1–9.

DODLA, V. B., G. C. SATYANARAYANA, AND S. DESAMSETTI (2017): "Analysis and prediction of a catastrophic Indian coastal heat wave of 2015," *Natural Hazards*, 87, 395–414.

EBI, K. L., A. CAPON, P. BERRY, C. BRODERICK, R. DE DEAR, G. HAVENITH, Y. HONDA, R. S. KOVATS, W. MA, A. MALIK, ET AL. (2021): "Hot weather and heat extremes: Health risks," *The Lancet*, 398, 698–708.

FERNÁNDEZ-VILLAVERDE, J. AND C. I. JONES (2020): "Macroeconomic outcomes and COVID-19: A progress report," Tech. rep., National Bureau of Economic Research.

FERRIS, J. S. AND D. A. NEWBURN (2017): "Wireless alerts for extreme weather and the impact on hazard mitigating behavior," *Journal of Environmental Economics and Management*, 82, 239–255.

GERUSO, M. AND D. SPEARS (2018): "Heat, humidity, and infant mortality in the developing world," NBER Working Paper 24870.

GOODMAN-BACON, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277.

GOULD, C. F., S. HEFT-NEAL, A. K. HEANEY, E. BENDAVID, C. W. CALLAHAN, M. KIANG, J. S. GRAFF ZIVIN, AND M. BURKE (2024): "Temperature extremes impact mortality and morbidity differently," NBER Working Paper 32195.

GUIDETTI, B., P. PEREDA, AND E. SEVERNINI (2021): ""Placebo tests" for the impacts of air pollution on health: The challenge of limited health care infrastructure," in *AEA Papers and Proceedings*, vol. 111, 371–375.

GUTTELING, J. M., T. TERPSTRA, AND J. H. KERSTHOLT (2018): "Citizens' adaptive or avoiding behavioral response to an emergency message on their mobile phone," *Journal of Risk Research*, 21, 1579–1591.

HE, G. AND T. TANAKA (2023): "Energy saving may kill: Evidence from the Fukushima nuclear accident," *American Economic Journal: Applied Economics*, 15, 377–414.

HEUTEL, G., N. H. MILLER, AND D. MOLITOR (2021): "Adaptation and the mortality effects of temperature across US climate regions," *Review of Economics and Statistics*, 103, 740–753.

IIZUKA, T., K. NISHIYAMA, B. CHEN, AND K. EGGLESTON (2021): "False alarm? Estimating the marginal value of health signals," *Journal of Public Economics*, 195, 104368.

JAPAN TODAY (2024): "Japan launches new alert system as heat stroke deaths rise," https://japantoday.com/category/national/japan-launches-new-alert-system-as-heat-stroke-deaths-rise, retrieved April 25, 2024.

KADOOKA, Y., A. ASAI, A. ENZO, AND T. OKITA (2017): "Misuse of emergent healthcare in contemporary Japan," *BMC Emergency Medicine*, 17, 1–6.

KARAIVANOV, A., S. E. LU, H. SHIGEOKA, C. CHEN, AND S. PAMPLONA (2021): "Face masks, public policies and slowing the spread of COVID-19: Evidence from Canada," *Journal of Health Economics*, 78, 102475.

KATAYAMA, Y., T. KITAMURA, J. TANAKA, S. NAKAO, M. NITTA, S. FUJIMI, Y. KUWAGATA, T. SHIMAZU, AND T. MATSUOKA (2021): "Factors associated with prolonged hospitalization among patients transported by emergency medical services: A population-based study in Osaka, Japan," *Medicine*, 100, e27862.

KIM, H. B., S. A. LEE, AND W. LIM (2019): "Knowing is not half the battle: Impacts of information from the National Health Screening Program in Korea," *Journal of Health Economics*, 65, 1–14.

KOKUDO, N. AND H. SUGIYAMA (2021): "Hospital capacity during the COVID-19 pandemic," *Global Health & Medicine*, 3, 56–59.

KOTHARKAR, R. AND A. GHOSH (2022): "Progress in extreme heat management and warning systems: A systematic review of heat-health action plans (1995–2020)," *Sustainable Cities and Society*, 76, 103487.

LI, T., C. CHEN, AND W. CAI (2022): "The global need for smart heat-health warning systems," *The Lancet*, 400, 1511–1512.

LOS ANGELES TIMES (2021): "The L.A. Times investigation into extreme heat's deadly toll," https://www.latimes.com/environment/story/2021-10-07/la-times-investigation-extreme-heat, retrieved May 14, 2024.

MARTIELLO, M. A. AND M. V. GIACCHI (2010): "High temperatures and health outcomes: A review of the literature," *Scandinavian Journal of Public Health*, 38, 826–837.

MENDOLIA, S., O. STAVRUNOVA, AND O. YEROKHIN (2021): "Determinants of the community mobility during the COVID-19 epidemic: The role of government regulations and information," *Journal of Economic Behavior & Organization*, 184, 199–231.

MICHELOZZI, P., G. ACCETTA, M. DE SARIO, D. D'IPPOLITI, C. MARINO, M. BACCINI, A. BIGGERI, H. R. ANDERSON, K. KATSOUYANNI, F. BALLESTER, 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, 383–389.

MINISTRY OF ENVIRONMENT (2021): ""Heat Stroke Alert" starts throughout Japan," https://www.env.go.jp/en/headline/2512.html, retrieved May 14, 2024.

——— (2024): "Heat illness prevention information," https://www.wbgt.env.go.jp/en/, retrieved May 14, 2024.

MULLINS, J. AND P. BHARADWAJ (2015): "Effects of short-term measures to Curb Air Pollution: Evidence from Santiago, Chile," *American Journal of Agricultural Economics*, 97, 1107–1134.

MULLINS, J. T. AND C. WHITE (2019): "Temperature and mental health: Evidence from the spectrum of mental health outcomes," *Journal of Health Economics*, 68, 102240.

——— (2020): "Can access to health care mitigate the effects of temperature on mortality?" *Journal of Public Economics*, 191, 104259.

NATIONAL PUBLIC RADIO (2022): "NWS HeatRisk," https://www.npr.org/2022/09/13/1122491546/why-heat-wave-warnings-are-falling-short-in-the-u-s, retrieved May 14, 2024.

NATIONAL WEATHER SERVICE (2023): "Weather related fatality and injury statistics," https://www.weather.gov/hazstat/, retrieved February 28, 2023.

——— (2024): "NWS HeatRisk," https://www.wpc.ncep.noaa.gov/heatrisk/, retrieved May 14, 2024.

NEIDELL, M. (2009): "Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations," *Journal of Human Resources*, 44, 450–478.

NEIDELL, M., S. UCHIDA, AND M. VERONESI (2021): "The unintended effects from halting nuclear power production: Evidence from Fukushima Daiichi accident," *Journal of Health Economics*, 79, 102507.

NOONAN, D. S. (2014): "Smoggy with a chance of altruism: The effects of ozone alerts on outdoor recreation and driving in Atlanta," *Policy Studies Journal*, 42, 122–145.

OECD (2024): "Length of hospital stay (indicator)," https://data.oecd.org/healthcare/length-of-hospital-stay.htm, retrieved May 2, 2024.

ONO, M. AND M. TONOUCHI (2014): "Estimation of wet-bulb globe temperature using generally measured meterological indices," *Japanese Journal of Biometeorology*, 50, 147–157.

PALECKI, M. A., S. A. CHANGNON, AND K. E. KUNKEL (2001): "The nature and impacts of the July 1999 heat wave in the Midwestern United States: Learning from the lessons of 1995," *Bulletin of the American Meteorological Society*, 82, 1353–1368.

PBS NEWS HOUR (2023): "Summer of record-breaking heat paints story of a warming world, scientists say," https://www.pbs.org/newshour/science/summer-of-record-breaking-heat-paints-story-of-a-warming-world-scientists-say, retrieved July 25, 2023.

REUTERS (2023): "Emissions watchers wary of Japan's power saving pleas," https://www.reuters.com/markets/commodities/emissions-watchers-wary-japans-power-saving-pleas-2023-06-12/, retrieved October 26, 2023.

ROBINSON, A., J. LEHMANN, D. BARRIOPEDRO, S. RAHMSTORF, AND D. COUMOU (2021): "Increasing heat and rainfall extremes now far outside the historical climate," *npj Climate and Atmospheric Science*, 4, 45.

SABERIAN, S., A. HEYES, AND N. RIVERS (2017): "Alerts work! Air quality warnings and cycling," *Resource and Energy Economics*, 49, 165–185.

SHRADER, J. G., L. BAKKENSEN, AND D. LEMOINE (2023): "Fatal errors: The mortality value of accurate weather forecasts," NBER Working Paper 31361.

SULYOK, M. AND M. WALKER (2020): "Community movement and COVID-19: A global study using Google's Community Mobility Reports," *Epidemiology & Infection*, 148, e284.

THE JAPAN TIMES (2023): "Ambulance dispatches in Tokyo increasing at record pace as summer heat arrives," https://www.japantimes.co.jp/news/2023/07/01/national/tokyo-ambulance-calls-rise/, retrieved September 7, 2023.

UNITED STATES ENVIRONMENTAL PROTECTION AGENCY (2023): "Climate change indicators: Heat-related deaths," https://www.epa.gov/climate-indicators/climate-change-indicators-heat-related-deaths, retrieved October 16, 2023.

VAIDYANATHAN, A., S. SAHA, A. M. VICEDO-CABRERA, A. GASPARRINI, N. ABDUREHMAN, R. JORDAN, M. HAWKINS, J. HESS, AND A. ELIXHAUSER (2019): "Assessment of extreme heat and hospitalizations to inform early warning systems," *Proceedings of the National Academy of Sciences*, 116, 5420–5427.

WARD, A. L. S. AND T. K. BEATTY (2016): "Who responds to air quality alerts?" *Environmental and Resource Economics*, 65, 487–511.

WEINBERGER, K. R., X. WU, S. SUN, K. R. SPANGLER, A. NORI-SARMA, J. SCHWARTZ, W. REQUIA, B. M. SABATH, D. BRAUN, A. ZANOBETTI, ET AL. (2021): "Heat warnings, mortality, and hospital admissions among older adults in the United States," *Environment International*, 157, 106834.

WEINBERGER, K. R., A. ZANOBETTI, J. SCHWARTZ, AND G. A. WELLENIUS (2018): "Effectiveness of National Weather Service heat alerts in preventing mortality in 20 US cities," *Environment International*, 116, 30–38.

YAGLOU, C., D. MINAED, ET AL. (1957): "Control of heat casualties at military training centers." *Arch. Indust. Health*, 16, 302–16.

ZIVIN, J. G. AND M. NEIDELL (2009): "Days of haze: Environmental information disclosure and intertemporal avoidance behavior," *Journal of Environmental Economics and Management*, 58, 119–128.

——— (2014): "Temperature and the allocation of time: Implications for climate change," *Journal of Labor Economics*, 32, 1–26.

A Appendix Tables and Figures

Table A1: FDMA Data – Summary Statistics

	(1) Obs.	(2) Mean	(3) SD	(4) Min	(5) Max
Panel A: Overall					
# of transports	16879	130.473	21.87	59.90	251.56
# of mild cases	16879	119.950	22.67	46.88	245.66
# of severe cases	16879	9.253	4.09	0.00	40.11
# of deaths	16879	1.270	0.88	0.00	10.59
Panel B: Cause for ambulance transport					
# of transports (Cause 1)	16879	0.185	0.81	0.00	41.09
# of mild cases (Cause 1)	16879	0.135	0.74	0.00	39.62
# of severe cases (Cause 1)	16879	0.031	0.14	0.00	7.25
# of deaths (Cause 1)	16879	0.019	0.10	0.00	4.45
# of transports (Cause 2)	16879	8.959	3.14	0.00	39.06
# of mild cases (Cause 2)	16879	8.589	3.04	0.00	37.76
# of severe cases (Cause 2)	16879	0.338	0.44	0.00	7.19
# of deaths (Cause 2)	16879	0.032	0.13	0.00	3.01
# of transports (Cause 4)	16879	106.574	20.24	44.17	216.94
# of mild cases (Cause 4)	16879	98.844	20.85	34.36	211.72
# of severe cases (Cause 4)	16879	6.632	2.90	0.00	32.09
# of deaths (Cause 4)	16879	1.097	0.80	0.00	9.78
# of transports (Cause 5)	16879	1.363	0.74	0.00	10.03
# of mild cases (Cause 5)	16879	1.133	0.68	0.00	8.60
# of severe cases (Cause 5)	16879	0.147	0.23	0.00	3.60
# of deaths (Cause 5)	16879	0.083	0.19	0.00	3.60
# of transports (Cause 6)	16879	10.735	4.66	0.00	43.71
# of mild cases (Cause 6)	16879	8.747	3.85	0.00	31.65
# of severe cases (Cause 6)	16879	1.968	1.53	0.00	17.19
# of deaths (Cause 6)	16879	0.020	0.11	0.00	2.87

Source: Authors' calculations using data from Fire and Disaster Management Agency (FDMA). *Note*: This table presents descriptive statistics of all key variables used in the analysis. Summary statistics are in terms of counts per million people. Cause 1: Fire & Water & Natural disaster. Cause 2: Traffic accident. Cause 3: Work & Sports. Cause 4: Sudden & General illness (incl. heat stroke). Cause 5: Violence & Self-harm. Cause 6: Hospital transfers. Summary statistics are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2017–2021.

Table A2: Google Trends – Summary Statistics

	(1) Obs.	(2) Mean	(3) SD	(4) Min	(5) Max
Heat strokes	10647	21.585	21.57	0.00	100.00
Heat stroke alert	10647	9.589	19.30	0.00	100.00
Weather	10647	52.095	19.55	11.00	100.00
Temperature	10647	25.214	18.65	0.00	100.00
Air conditioner	10647	38.394	20.25	0.00	100.00
Outdoor	10647	37.152	26.79	0.00	100.00
Sea bathing	10647	17.186	23.30	0.00	100.00
Park	10647	33.950	13.77	0.00	100.00
Indoor	10647	21.539	22.82	0.00	100.00
Cinema	10647	40.981	23.40	0.00	100.00
Karaoke	10647	48.014	19.30	0.00	100.00

Source: Authors' calculations using data from Google Trends. *Note*: This table presents descriptive statistics of all key variables used in the analysis. Google Trends Searches for 熱中症 (heat stroke), 熱中症警戒アラート (heat stroke alert), 天気 (weather), 気温 (atmospheric temperature), エアコン (air conditioner), 屋外 (outdoor), 海水浴 (sea bathing), 公園 (park), 屋内 (indoor), 映画館 (cinema), and カラオケ (karaoke). Summary statistics are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2020–2022.

Table A3: Google Mobility Reports – Summary Statistics

	(1) Obs.	(2) Mean	(3) SD	(4) Min	(5) Max
Retail & recreation	10647	−7.934	8.79	−89.00	56.00
Grocery & pharmacy	10647	4.153	6.29	−82.00	45.00
Parks	10293	−3.949	22.94	−80.00	280.00
Transit stations	10636	−16.554	17.27	−85.00	105.00
Workplaces	10647	−12.656	12.77	−78.00	8.00
Residential	10647	6.333	3.82	−2.00	34.00

Source: Authors' calculations using data from Google COVID-19 Community Mobility Reports. *Note*: This table presents descriptive statistics of all key variables used in the analysis. Summary statistics are weighted by a region's population in 2019. WBGT $\geq$ 28. Data for 2020–2022.

Table A4: Effect of Heat Alerts on Heat Strokes by Place of Incidence

	# of heat strokes							
	(1) Home	(2) Workpl. (constr.)	(3) Workpl. (field)	(4) Educ.	(5) Public (ins.)	(6) Public (outs.)	(7) Street	(8) Other
$= 1$ if heat alert	0.154^{***}	0.204^{***}	0.145^{**}	0.208^{***}	0.127^{***}	0.195^{***}	0.164^{***}	0.167^{***}
	(0.020)	(0.034)	(0.066)	(0.051)	(0.046)	(0.033)	(0.030)	(0.045)
WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓	✓	✓	✓	✓
Mean heat strokes	2.040	.555	.125	.291	.392	.619	.765	.283
Observations	20612	20612	20612	20612	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes by place of incidence in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A5: Effect of Consecutive Heat Alerts on Heat Strokes

	# of heat strokes			
	(1)	(2)	(3)	(4)
Heat strokes: total				
$= 1$ if heat alert_{t-2}			-0.004	-0.004
			(0.030)	(0.030)
$= 1$ if heat alert_{t-1}	0.016	0.004	0.005	0.005
	(0.015)	(0.021)	(0.021)	(0.021)
$= 1$ if heat alert_{t}	0.166^{***}	0.156^{***}	0.158^{***}	0.166^{***}
	(0.015)	(0.018)	(0.019)	(0.021)
$\text{Alert}_t \times \text{alert}_{t-2}$			-0.006	-0.051
			(0.041)	(0.057)
$\text{Alert}_t \times \text{alert}_{t-1}$		0.027	0.030	0.009
		(0.033)	(0.033)	(0.037)
$\text{Alert}_t \times \text{alert}_{t-1} \times \text{alert}_{t-2}$				0.064
				(0.051)
WBGT	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓
Mean heat strokes	5.075	5.075	5.075	5.075
Observations	20600	20600	20600	20600

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA).
Note: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT ≥ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A6: Effect of Heat Alerts on Ambulance Transports (Low Income)

	# of ambulance transports							
	All		Mild		Severe		Deaths	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$= 1$ if heat alert	0.023***	0.027***	0.019**	0.022**	0.057***	0.064***	0.023	0.006
	(0.008)	(0.008)	(0.009)	(0.009)	(0.019)	(0.021)	(0.041)	(0.041)
WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓		✓
Mean # of transports	133.382	133.382	115.717	115.717	15.465	15.465	2.198	2.198
Observations	110141	110141	110141	110141	110141	110141	110141	110141

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of ambulance transports by severity diagnosis in a region on the presence of a heat alert and a set of control variables. The mean number of transports per million people is for regions without heat alerts (control group). Analysis at the emergency response unit level. WBGT $\geq$ 28. Data for 2017–2021. Only regions with taxable per capita income below the median are considered. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A7: Effect of Heat Alerts on Ambulance Transports (High Income)

	# of ambulance transports							
	All		Mild		Severe		Deaths	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$= 1$ if heat alert	0.011**	0.009*	0.013**	0.011*	−0.010	−0.006	−0.024	−0.005
	(0.005)	(0.005)	(0.005)	(0.006)	(0.011)	(0.012)	(0.045)	(0.043)
WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓		✓
Mean # of transports	126.771	126.771	115.481	115.481	9.852	9.852	1.439	1.439
Observations	125906	125906	125906	125906	125906	125906	125906	125906

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of ambulance transports by severity diagnosis in a region on the presence of a heat alert and a set of control variables. The mean number of transports per million people is for regions without heat alerts (control group). Analysis at the emergency response unit level. WBGT $\geq$ 28. Data for 2017–2021. Only regions with taxable per capita income above the median are considered. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A8: Effect of Heat Alert on Energy Consumption

	Log (energy consumption (in GW))					
	All		Night		Day	
	(1)	(2)	(3)	(4)	(5)	(6)
$= 1$ if heat alert	0.047**	0.029**	0.037*	0.018	0.051**	0.033***
	(0.021)	(0.011)	(0.020)	(0.011)	(0.021)	(0.012)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓
Mean energy use (in GW)	56.326	56.326	21.958	21.958	34.368	34.368
Observations	20816	20816	20816	20816	20816	20816

Source: Authors' calculations using data from Japan Renewable Energy Foundation and Ministry of Environment. *Note*: This table shows results from OLS regressions of the energy consumption in a region on the presence of a heat alert and a set of control variables. Night begins at 7pm and ends at 6am. Day begins at 7am and ends at 6pm. The mean outcome is for regions without heat alerts (control group). All regressions are weighted by the number of prefectures included in a region. WBGT ≥ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A9: Robustness Check – Different Cutoffs for WBGT

	# of heat strokes					
	(1) WBGT $= [20, max]$	(2) WBGT $= [25, max]$	(3) WBGT $= [28, max]$	(4) WBGT $= [20, 35]$	(5) WBGT $= [25, 35]$	(6) WBGT $= [28, 35]$
$= 1$ if heat alert	0.164***	0.164***	0.172***	0.165***	0.165***	0.173***
	(0.016)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend	✓	✓	✓	✓	✓	✓
Mean heat strokes	3.177	3.839	5.071	3.094	3.740	4.940
Observations	37212	29708	20612	36919	29415	20319

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A10: Robustness Check – Unrestricted Sample

	Log (# of heat strokes +1 per million people)		# of heat strokes per million people		# of heat strokes	
	OLS				Neg. binomial regression	
	(1)	(2)	(3)	(4)	(5)	(6)
$= 1$ if heat alert	0.230***	0.223***	1.214***	1.267***	0.191***	0.190***
	(0.026)	(0.024)	(0.205)	(0.213)	(0.021)	(0.020)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓
Mean heat strokes	5.559	5.559	5.559	5.559	5.559	5.559
Observations	20816	20816	20816	20816	20816	20816

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from OLS and negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. Data is *not* restricted to the 99^{th} percentile of the number of heat strokes. The mean number of heat strokes per million people is for regions without heat alerts (control group). Regressions in columns (1) to (4) are weighted by a region's population in 2019. WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A11: Robustness Check – Alert Due to Same-Day 5am vs. Previous-Day 5pm Forecasts

	# of heat strokes					
	Morning		Night		Morning & night	
	(1)	(2)	(3)	(4)	(5)	(6)
$= 1$ if heat alert at 5pm			0.171***	0.167***	0.061***	0.061***
			(0.017)	(0.015)	(0.020)	(0.018)
$= 1$ if heat alert at 5am	0.174***	0.172***			0.132***	0.130***
	(0.016)	(0.015)			(0.020)	(0.020)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓
Mean heat strokes	5.071	5.071	5.071	5.071	5.071	5.071
Observations	20612	20612	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A12: Robustness Check – Without Large Commuting Zones

	# of heat strokes					
	W/o Greater Tokyo area		W/o Chukyo metropolitan area		W/o Keihanshin metropolitan area	
	(1)	(2)	(3)	(4)	(5)	(6)
$= 1$ if heat alert	0.208^{***}	0.171^{***}	0.186^{***}	0.174^{***}	0.190^{***}	0.173^{***}
	(0.019)	(0.017)	(0.018)	(0.014)	(0.020)	(0.016)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
Region-spec. trend		✓		✓		✓
Mean heat strokes	6.164	6.164	5.031	5.031	5.006	5.006
Observations	18672	18672	19161	19161	19230	19230

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA).
Note: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. Models in columns (1) and (2) drop regions in the Greater Tokyo area (i.e. Saitama, Chiba, Tokyo, and Kanagawa). Models in columns (3) and (4) drop regions in the Chukyo metropolitan area (i.e. Gifu, Aichi, and Mie). Models in columns (5) and (6) drop regions in the Keihanshin metropolitan area (i.e. Kyoto, Osaka, and Hyogo). The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT ≥ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A13: Robustness Check – Effect of Heat Alerts on Heat Strokes on Weekdays vs. Weekends

	# of heat strokes			
	Weekdays		Weekends	
	(1)	(2)	(3)	(4)
$= 1$ if heat alert	0.171^{***}	0.163^{***}	0.180^{***}	0.183^{***}
	(0.018)	(0.017)	(0.029)	(0.027)
Mean WBGT	✓	✓	✓	✓
Lagged Mean WBGT	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend		✓		✓
Mean heat strokes	4.971	4.971	5.320	5.320
Observations	14791	14791	5821	5821

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA).
Note: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT ≥ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A14: Robustness Check – State of Emergency (COVID-19)

	# of heat strokes			
	(1)	(2)	(3)	(4)
$= 1$ if heat alert	0.175***	0.173***	0.177***	0.176***
	(0.016)	(0.015)	(0.017)	(0.016)
$= 1$ if state of emerg.	0.018	0.041	0.023	0.046
	(0.049)	(0.056)	(0.050)	(0.056)
Alert $\times$ state of emerg.			−0.025	−0.027
			(0.043)	(0.044)
WBGT	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend		✓		✓
Mean heat strokes	5.071	5.071	5.071	5.071
Observations	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert, presence of state of emergency, and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A15: Robustness Check – Alternative Fixed Effects

	Log (# of heat strokes +1 per million people)		# of heat strokes per million people		# of heat strokes	
	OLS				Neg. binomial regression	
	(1)	(2)	(3)	(4)	(5)	(6)
$= 1$ if heat alert	0.175***	0.182***	0.989***	1.189***	0.192***	0.168***
	(0.022)	(0.016)	(0.265)	(0.140)	(0.019)	(0.015)
WBGT	✓	✓	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓	✓	✓
Region $\times$ Month $\times$ Year FE	✓		✓		✓	
Region $\times$ WBGT (rounded) FE		✓		✓		✓
Day FE		✓		✓		✓
Mean heat strokes	5.071	5.071	5.071	5.071	5.071	5.071
Observations	20612	20612	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from OLS and negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The mean number of heat strokes per million people is for regions without heat alerts (control group). Regressions in columns (1) to (4) are weighted by a region's population in 2019. WBGT $\geq$ 28. Robust standard errors clustered at the regional level are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively.

Table A16: Robustness Check – Estimator of De Chaisemartin and d'Haultfoeuille (2020)

	Log(# of heat strokes +1 per million people)		# of heat strokes per million people	
	(1)	(2)	(3)	(4)
= 1 if heat alert	0.154***	0.154***	1.052***	1.052***
	(0.025)	(0.025)	(0.215)	(0.215)
WBGT	✓	✓	✓	✓
Lagged WBGT	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓
Region-spec. trend		✓		✓
Mean heat strokes	5.071	5.071	5.071	5.071
Observations	20612	20612	20612	20612

Source: Authors' calculations using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This table shows results from OLS regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables, using the estimator of De Chaisemartin and d'Haultfoeuille (2020). WBGT and lagged WBGT is controlled for linearly. The mean number of heat strokes per million people is for regions without heat alerts (control group). All regressions are weighted by a region's population in 2019. WBGT $\geq$ 28. Robust standard errors clustered at the regional level are obtained from 199 bootstrap replications and are given in parentheses. *, **, and *** denote significance at the 10%-, 5%-, and 1%-level, respectively. Results are obtained from the Stata command `did_multiplegt` of De Chaisemartin et al. (2019).

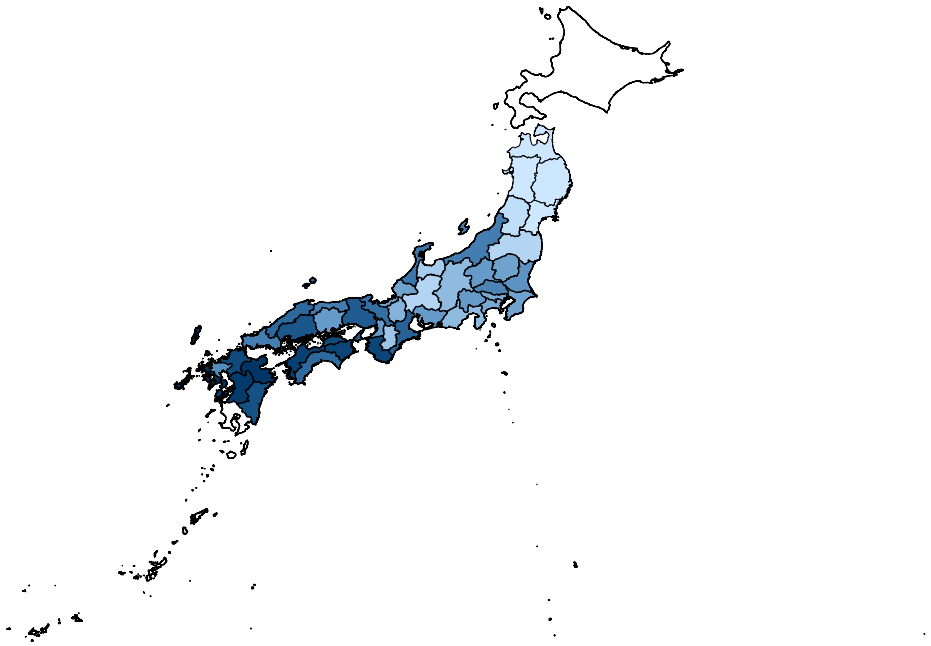


Figure A1: Share of Dates with Heat Alert in 2022

Source: Authors' presentation using data from Ministry of Environment. *Note*: This graph shows the share of dates with a heat alert between May 1 and September 30, 2022, across Japan. Regions in white are dropped due to imperfect mapping of regions to prefectures (described further in text). Regions in bright blue had zero alerts. Darker color indicates a higher share.

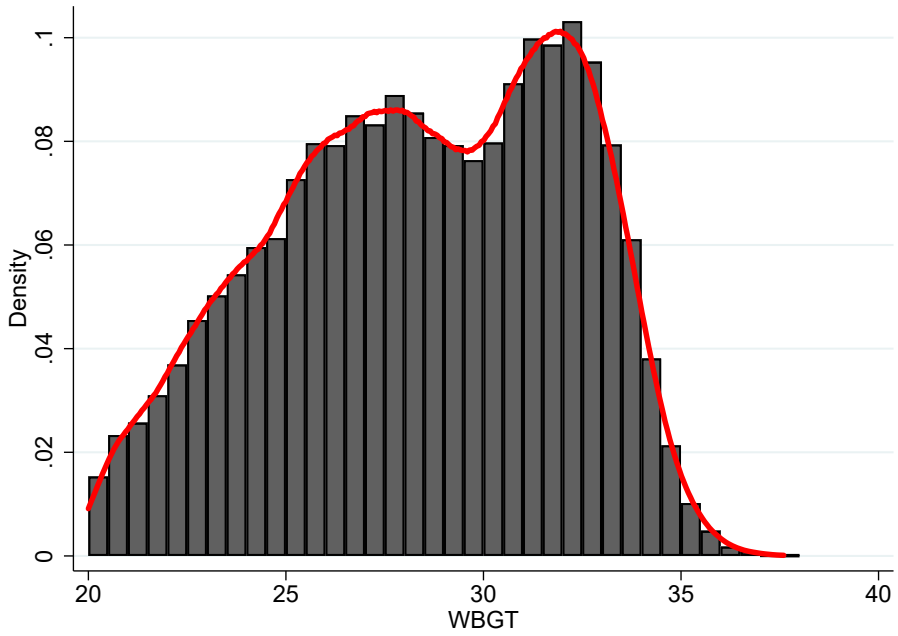
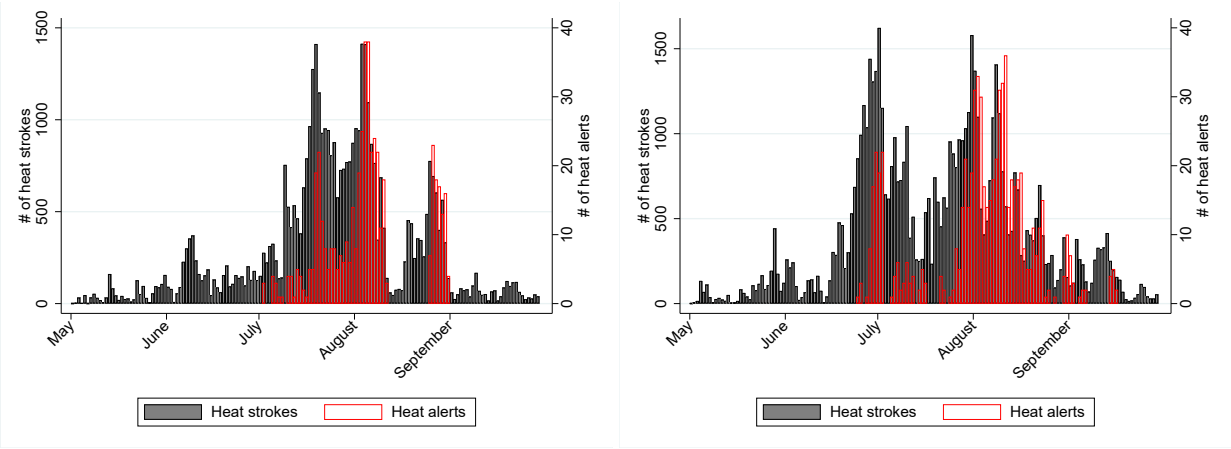


Figure A2: Distribution of Realized WBGT

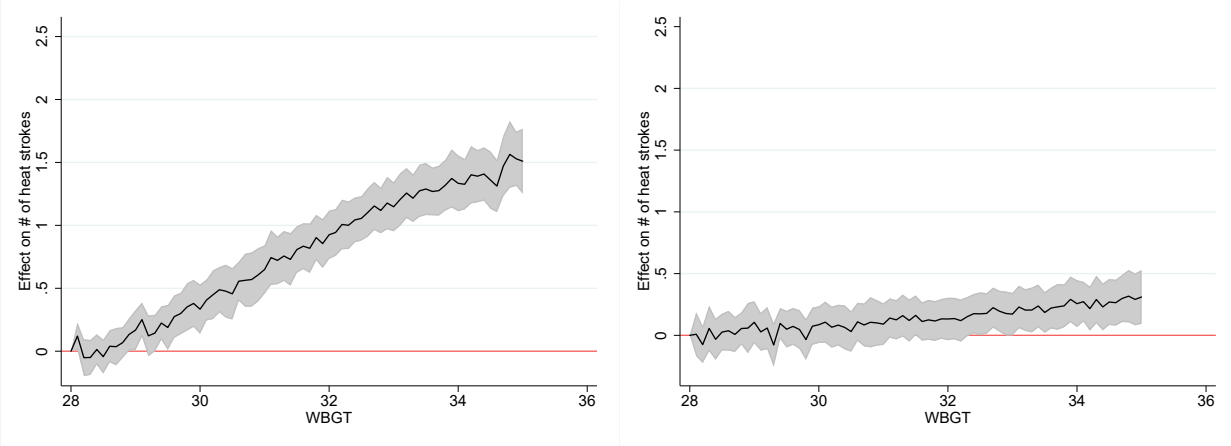
Source: Authors' presentation using data from Ministry of Environment. *Note*: This graph shows the distribution of maximum realized WBGT in a region for the dates between May 1 and September 30 for the years 2017 to 2022. A kernel density estimate based on the epanechnikov kernel and the rule-of-thumb bandwidth is shown in red.



(a) 2021 (b) 2022

Figure A3: Heat Strokes and Heat Alerts Over Time

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: These graphs show the distribution of the number of heat strokes and heat alerts for 2021 (a) and 2022 (b).



(a) Actual WBGT

(b) Lagged WBGT

Figure A4: Relationship between WBGT and Heat Strokes

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: These graphs show the relationship between the WBGT index (a) and the lagged WBGT index (b), respectively, and heat strokes according to our baseline specification presented in column (6) of Table 2. The estimated coefficients are presented for WBGT $= [28, 35]$, whereas WBGT $= 28.0$ serves as baseline. All regressions are weighted by a region's population in 2019. The grey area denotes 95%-confidence intervals based on robust standard errors clustered at the regional level.

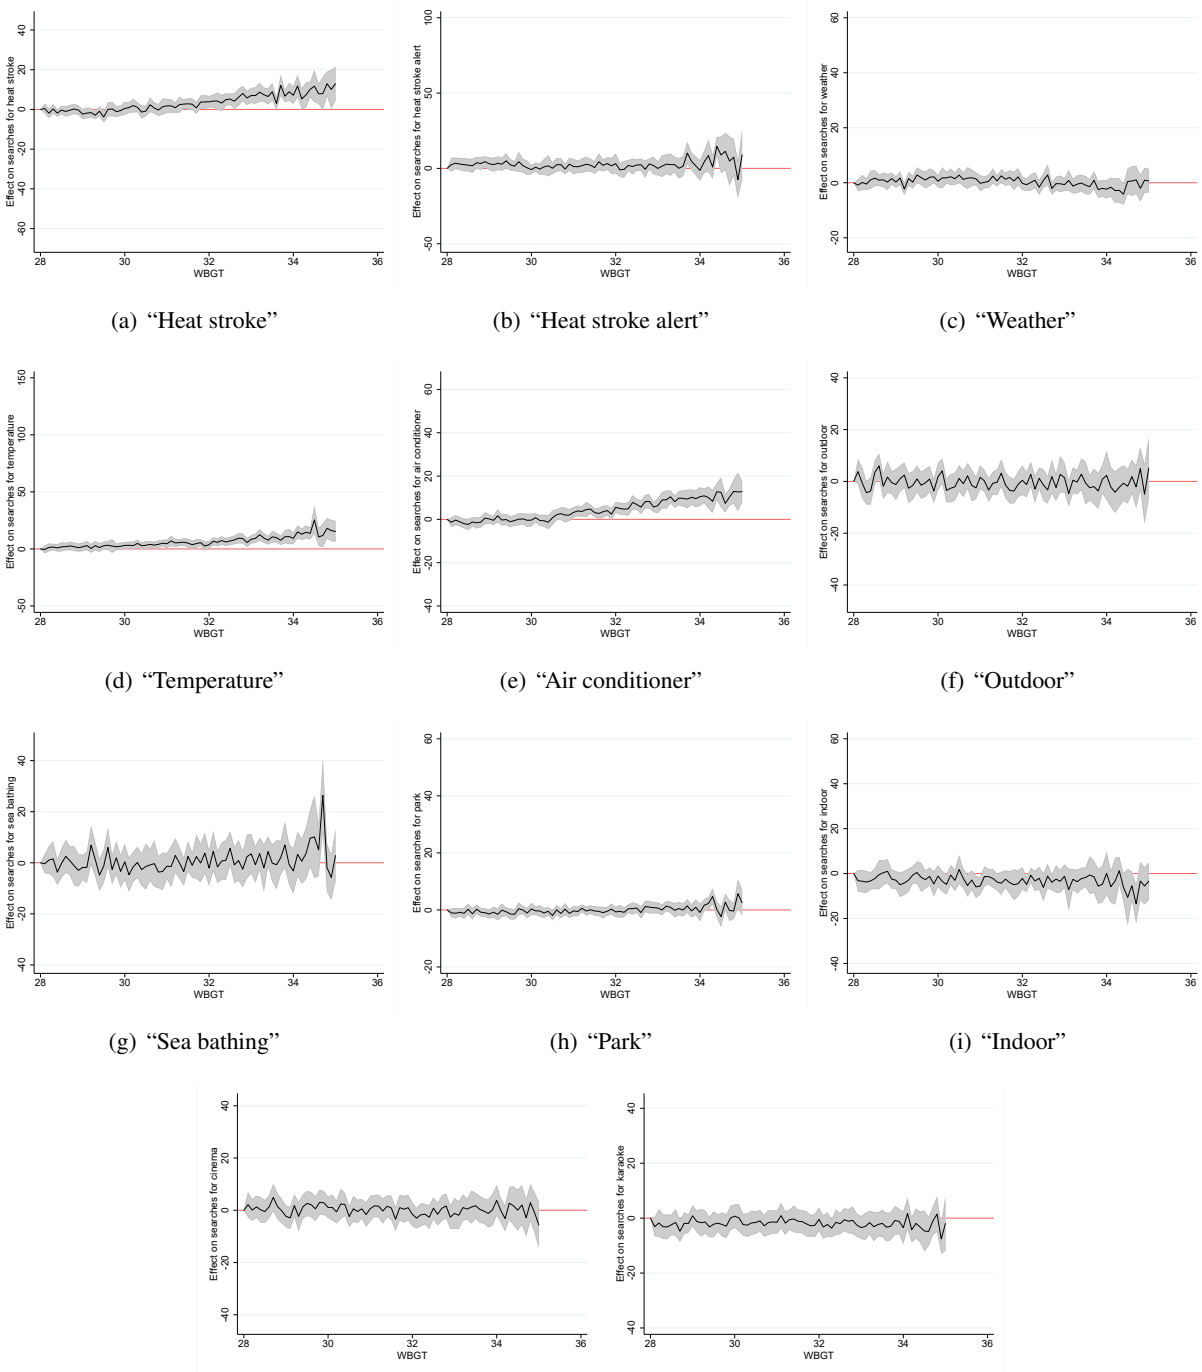
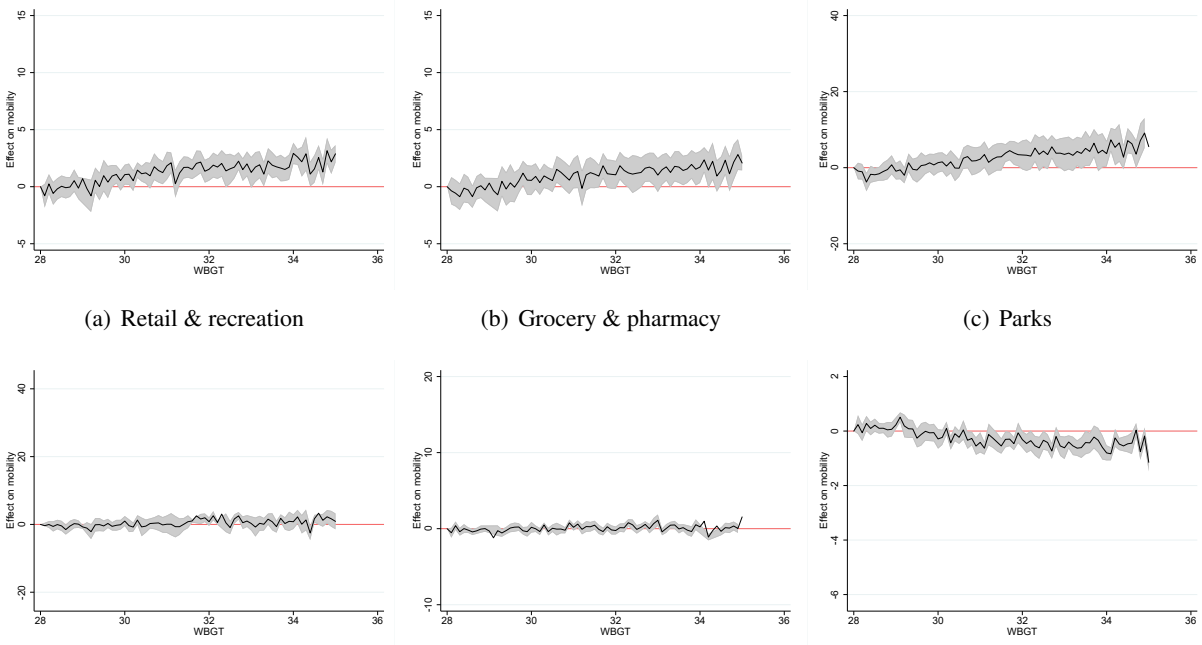


Figure A5: Relationship between WBGT and Google Trend Searches

Source: Authors' presentation using data from Google Trends and Ministry of Environment. *Note*: These graphs show the relationship between the WBGT index and Google Trend Searches for 熱中症 (heat stroke), 熱中症警戒アラート (heat stroke alert), 天気 (weather), 気温 (atmospheric temperature), エアコン (air conditioner), 屋外 (outdoor), 海水浴 (sea bathing), 公園 (park), 屋内 (indoor), 映画館 (cinema), and カラオケ (karaoke), respectively, according to our baseline specification presented in Table 6 and 7. The effects are presented for WBGT $= [28, 35]$, whereas WBGT $= 28.0$ serves as baseline. All regressions are weighted by a region's population in 2019. Data for 2020–2022. The grey area denotes 95%-confidence intervals based on robust standard errors clustered at the regional level.



(a) Retail & recreation (b) Grocery & pharmacy (c) Parks

(d) Transit stations (e) Workplaces (f) Residential

Figure A6: Relationship between WBGT and Mobility

Source: Authors' presentation using data from Google COVID-19 Community Mobility Reports and Ministry of Environment. *Note*: These graphs show the relationship between the WBGT index and the percentage changes in number of visits to specific places in a region, according to our baseline specification presented in Table 8. The effects are presented for WBGT $= [28, 35]$, whereas WBGT $= 28.0$ serves as baseline. All regressions are weighted by a region's population in 2019. Data for 2020–2022. The grey area denotes 95%-confidence intervals based on robust standard errors clustered at the regional level.

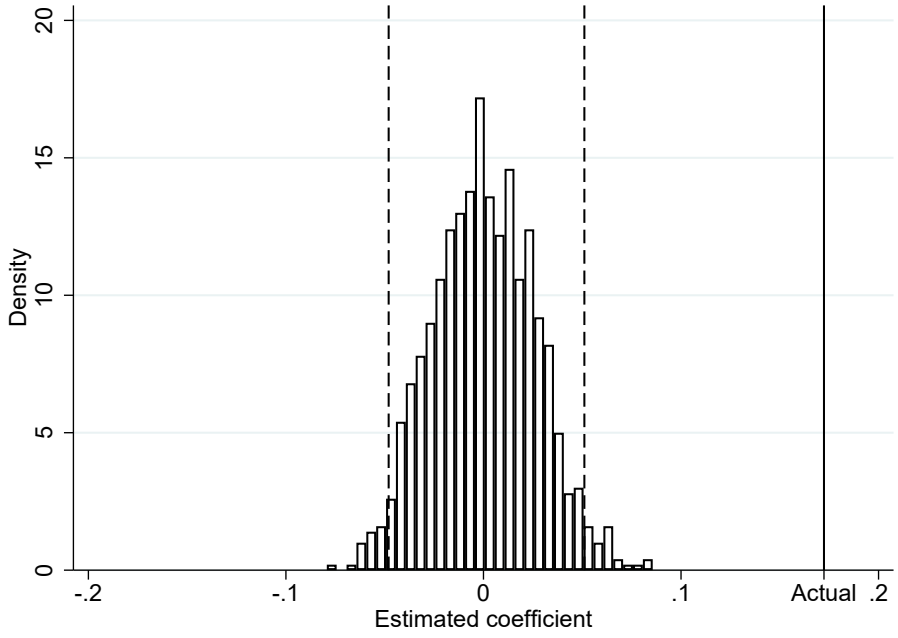
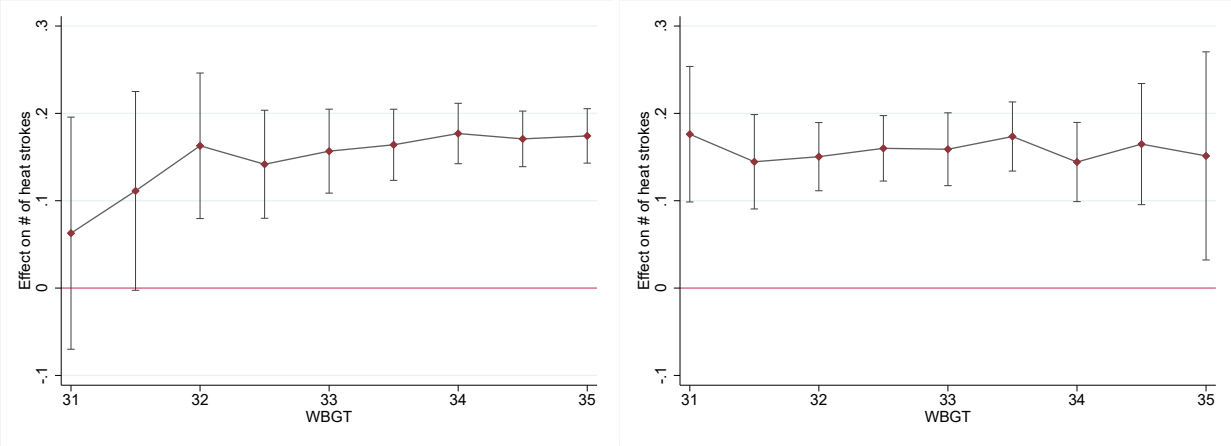


Figure A7: Placebo Distribution of Heat Alert Effects

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: This graph shows the distribution of placebo negative binomial regressions of the number of heat strokes in a region on the presence of a heat alert and a set of control variables. The actual estimate based on column (6) of Table 3 is shown as solid line. Heat alert patterns over time are randomly assigned to regions (without replacement) and the negative binomial regression as described before is run. This procedure is repeated 1,000 times. This approach is based on Bertrand et al. (2004). Dashed lines indicate empirical 95%-confidence intervals.



(a) Increasing sample size

(b) Moving average

Figure A8: Effect of Heat Stroke Alerts on Heat Strokes by WBGT Intervals

Source: Authors' presentation using data from Ministry of Environment and Fire and Disaster Management Agency (FDMA). *Note*: These graphs show the effect of heat stroke alerts on heat strokes according to our baseline specification, while increasing the sample size in 0.5-intervals of the WBGT index (a) and using subsamples for specific WBGT intervals, respectively. In panel (a), presented estimates are the effect of an alert for the sample with temperatures just including the respective WBGT. In panel (b), presented estimates are the effect of an alert for a sample with temperatures centered around the respective WBGT $\pm$ 1. WBGT $\geq$ 28. 95%-confidence intervals are based on robust standard errors clustered at the regional level.