Discussion Paper Series

IZA DP No. 18601

April 2026

The Sun is for Everyone, the Heat for Some: Heatwaves and Mortality within Cities

Vinicius Peçanha
Institute for Health Policy Studies

Rudi Rocha
São Paulo School of Business Administration, FGV, Institute for Health Policy Studies and IZA@LISER

Dimitri Szerman
Amazon

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

IZA DP No. 18601

April 2026

The Sun is for Everyone, the Heat for Some: Heatwaves and Mortality within Cities*

Abstract

If much of the variation in climate exposure occurs across short distances, then so too might the health consequences of heatwaves and the potential for place-based adaptation. We test this by combining high-resolution satellite data and administrative death records from Rio de Janeiro to estimate neighborhood-level heat effects. Nearly 60% of excess elderly mortality is driven by localized exposure differences. Yet as temperatures rise, spatial variation declines and city-wide shocks become more dominant. Preventive care and proximity to emergency services attenuate mortality, but only emergency access remains protective under localized exposure. Intervention points may thus lie hidden within city-level averages.

JEL classification

I14, I15, Q54

Keywords

heat waves, mortality, mitigation policies, healthcare

Corresponding author

Rudi Rocha
rudi.rocha@fgv.br

* We thank Claudio Ferraz, Thomas Lemieux, Siwan Anderson, Rafael Lalive, Sébastien Houde, Nathan Nunn, Jamie McCasland, Sumeet Gulati, Bruna Borges, Julia Guerra and seminar participants at XXI RIDGE Forum for the Health Economics Network, Insper, UBC Development Lunch, UBC Environmental seminar, 41st Brazilian Econometric Meeting, Global Open Series in Environmental Economics, and the CIREQ meeting.

1 Introduction

Recent evidence from climate science and environmental studies reveals that temperature exposure varies sharply within cities. Satellite-based thermal data show that neighboring urban areas can differ in surface temperatures by several degrees on the same day, driven by microclimates, built environments, and vegetation cover (e.g. Yin et al., 2023). These patterns suggest that climate vulnerability is not only a national or regional phenomenon, but deeply local.

If much of the variation in climate exposure occurs across short distances, then so too might the health consequences of heat waves and the potential for place-based adaptation. This premise is now central to urban climate policy. Cities worldwide are investing in spatially targeted strategies such as neighborhood greening, localized cooling infrastructure, and community outreach programs aimed at protecting vulnerable residents during extreme heat (Burke et al., 2025). These investments implicitly assume that climate-related health risks are also localized. Yet the standard empirical studies rely on spatially aggregated measures of heat and mortality, leaving open a fundamental question: how important are localized differences in climate vulnerability and health impacts? Without evidence at finer spatial scales, we cannot assess whether the most consequential impacts of extreme heat, or the most effective points of intervention, lie hidden within city-level averages.

This paper provides the first test of whether heat-related mortality operates at a local scale, and whether localized policy responses can be effective. We study Rio de Janeiro, one of the world's largest cities, combining satellite-derived temperature data at fine spatial resolution with individually geocoded administrative mortality records to construct a neighborhood-by-month panel. These data allow us to decompose heat-related mortality into two components: deaths attributable to localized differences in exposure across neighborhoods, and deaths attributable to city-wide temperature shocks. We then test whether neighborhood-level health policies can mitigate each component. The answer, we show, depends fundamentally on which type of heat stress dominates.

Our empirical strategy draws on fixed effects specifications used to identify causal impacts of temperature on health outcomes (e.g. Deschênes and Moretti, 2009; Deschênes and Greenstone, 2011; Barreca et al., 2016). By exploring the residual variation in temperature leveraged from different combinations of time and neighborhood fixed effects, we decompose heat-related mortality into components driven by localized exposure and by common city-wide shocks. This approach also allows us to test whether the effectiveness of neighborhood-level health policies differs across these two sources of heat stress.

Implementing this design requires measuring both temperature and mortality at fine spatial scales. For temperature, we use land surface temperature (LST) derived from satellite imagery, which captures thermal radiation emitted by the land surface at a nominal resolution of 1 km^2. Unlike weather stations, whose spatial coverage within cities is typically sparse, LST data provide daily temperature for each of Rio's 1,524 pixels, which we aggregate to neighborhoods using census-tract population weights. For mortality, we draw on administrative microdata containing the universe of deaths in Rio, including exact neighborhood of residence, date, and cause of death. We focus on mortality from chronic diseases among individuals aged 60 and over — the population most physiologically vulnerable to heat stress, which can trigger cardiovascular, respiratory, and endocrine complications (Hajat et al., 2014; Achebak et al., 2018; Kephart et al., 2022b; Ballester et al., 2023).

Rio de Janeiro provides a relevant empirical setting. First, the city has more than 6.3 million people scattered across diverse neighborhoods regarding geographical elements and socioeconomic characteristics. Forested massifs reaching heights superior to 1,000 meters affect patterns of winds and temperature, as they shape the penetration of the Atlantic sea breeze into the hinterland and provide shade, contributing to the formation of micro-climates and significant variation in temperature within the city (Neiva et al., 2017). In that sense, Rio's thermal landscape provides us with a prime testing ground to detect whether localized patterns in the heat-mortality relationship exist, a question that would not be as salient in more thermally homogeneous settings.

Second, Rio also provides unique variation in access to health services. We exploit the restructuring of the city's primary healthcare system, mainly driven by the staggered expansion of health units under the Family Health Program (FHP) strategy. This program aimed to expand primary care in their respective catchment areas. Population coverage started from 3.7% in 2008, reaching more than 50% in 2015, portraying significant variation over time both within and across neighborhoods. In parallel, the network of emergency room (ER) facilities improved with the implementation of more than 20 non-hospital emergency care units. We rely on precise geocoding of the FHP health units' catchment areas and ER facilities and exploit idiosyncratic variation in differential access to services across time and neighborhoods to examine whether and how the design of local health systems can mitigate impacts.

Our analysis yields three key findings. First, intra-city variation in high-temperature exposure is large enough for the temperature–mortality relationship to emerge. In our less saturated specification, conditional on year and neighborhood × calendar month fixed effects, we find that one additional hot day per month increases mortality by 1.016 deaths

per 100,000 individuals aged 60 and over, which corresponds to roughly 0.56% of the average monthly mortality rate from chronic diseases in this age group. This result is robust across alternative definitions of heat stress and to potential harvesting effects.

Second, the estimated effect declines by roughly 40% when we introduce time fixed effects, thus absorbing city-wide temperature shocks and identifying effects solely from cross-neighborhood variation. This implies that approximately 60% of the total heat-related mortality effect stems from differences in localized exposure. In line with this result, we show that temperature variance across neighborhoods narrows as city-wide temperatures rise, making extreme heat events more spatially uniform. In other words, as average temperatures rise and exposure converges across neighborhoods, the share of mortality attributable to localized differences falls, while the role of common shocks becomes more pronounced. These findings underscore the importance of distinguishing between general and localized heat impacts when evaluating both risks and responses.

Third, we find that both proximity to emergency services and to localized preventive care are effective in reducing mortality under city-wide shocks. However, when neighborhood-specific differences in exposure dominate, the effectiveness of localized mitigation narrows while proximity to emergency services remains important. In other words, while the mitigating effect of preventive care becomes statistically weaker once time fixed effects absorb city-wide shocks, the point estimates for ER-proximity interactions remain stable, indicating that ER access continues to buffer mortality. This pattern indicates that local preventive strategies, while important, may be insufficient on their own, emphasizing the need for coordinated public health interventions to protect vulnerable populations under more extreme climate conditions.

This paper departs from the existing literature on heat-related mortality in a fundamental way. Prior studies, including influential work using subnational fixed effects models, identify heat-mortality effects from variation in temperature across locations, measured from a small number of weather stations near cities or in city centers as a proxy for heat stress, and treating within-location exposure as homogeneous (Deschênes and Moretti, 2009; Deschênes and Greenstone, 2011; Barreca et al., 2016; Burgess et al., 2017; Cohen and Dechezleprêtre, 2022). This reflects a design that, by construction, cannot distinguish localized from aggregate heat impacts or assess whether place-based policies can mitigate them. Our contribution is to answer these questions directly. By measuring exposure at the neighborhood level and exploiting variation within city, we isolate localized effects and assess whether they can be addressed through localized responses. We find that they can, but only partially and depending on the nature of heat stress.

Finally, our focus on an urban environment addresses a first-order policy concern. While

the need for adaptation is widespread worldwide, urban areas will be on the front lines of climate warming. Cities are projected to host nearly 6.7 billion people, or 70% of the world's population by 2050, with rapid expansion concentrated in tropical regions already facing extreme heat (UN, 2018). Our results are therefore informative for the design of localized mitigation policies in an increasingly urbanized and heat-exposed world. This focus is also important given the mixed evidence on whether the heat–mortality relationship differs between rural and urban areas. The existing evidence ranges from predominantly rural effects (Burgess et al., 2017; Aragón et al., 2021) to negligible differences (Cohen and Dechezleprêtre, 2022). In many developing countries, rural heat exposure is often confounded by effects on agricultural productivity, food intake, or income (Banerjee and Maharaj, 2020). By contrast, our setting isolates the direct physiological pathway from heat to mortality in a purely urban context, providing evidence directly applicable to the adaptation challenges facing the world's growing cities.

The remainder of this paper is organized as follows. Section 2 describes our empirical setting and the temperature data, while Section 3 describes the health data and the auxiliary variables used in our analysis. Section 4 presents and discusses the empirical strategy. In Section 5, we present the main results on the heat-mortality relationship in the city of Rio and the main robustness checks. In Section 6, we test whether and how access to health services can mitigate health impacts. Section 7 concludes.

2 Background and Temperature Data

2.1 Rio de Janeiro: Geography and Climate

The city of Rio de Janeiro spreads over 1,225 km^2 on the tropical zone of the planet, under intense insolation all year round. Rio experiences higher temperatures between December to March, during the summer, and relatively warm winters. Rio is well known for its natural landscape, with dramatic geographical features scattered across the city that determine temperature variation, along with proximity to the ocean and general and secondary atmospheric circulation (Neiva et al., 2017). Forested massifs reaching heights greater than 1,000 meters affect wind and temperature patterns, as they provide shade to some areas and shape the penetration of the Atlantic sea breeze into the hinterland, contributing to the formation of microclimates. Sites behind the massifs receive winds warmer than those in the windward position (Neiva et al., 2017; Rio de Janeiro, 2016; Serra and Ratisbonna, 1941). Neighborhoods in the North and West Zones are usually the warmest, in contrast to areas in the South Zone, where the Atlantic sea breeze cools the air.

The left map of Figure 1 shows the main geographical elements of Rio and its division

into Administrative Regions (North, West, South, and Central Zones). The city is home to more than 6.3 million people and to significant socioeconomic heterogeneity across neighborhoods. Poverty rates are nearly twofold in neighborhoods in the North and West Zones in comparison to those in the South Zone. Human Development Index (HDI) ranges from 0.970 in the neighborhood of Gávea to 0.732 in the contiguous Rocinha, both in the South Zone, to 0.700 in Complexo do Alemão, where income per capita is only a tenth in comparison to the first (IETS, 2015; IBGE, 2010).[1]

2.2 Air Temperature and Weather Data from Ground Stations

Obtaining high-quality temperature data remains a major challenge in developing countries due to limited spatial and temporal coverage of ground weather stations (Auffhammer et al., 2013). The city of Rio de Janeiro is no exception. Two parallel meteorological systems operate in the city, totaling only eight ground weather stations. Of these, merely two have consistent temperature records dating back to the early 2000s.

The first system, *Alerta Rio*, maintains seven stations across the city. However, only one of these stations offers uninterrupted weather data for the full 2003–2016 period. The second system is operated by the National Institute of Meteorology (INMET), which maintains a single station in Rio de Janeiro. Both systems collect a range of meteorological variables – such as temperature, humidity, wind speed, and piche evaporation – but their limited number and coverage severely constrain fine-grained temperature measurement over time and space.

To address the limited spatial coverage of ground stations, many studies rely on gridded weather datasets that interpolate temperature over space and time using available station data (e.g., Matsuura and Willmott, 2018). However, in regions outside the United States, the spatial resolution of such products is often coarse, typically on the order of several kilometers (Hooker et al., 2018). Given that only two stations in our study area have continuous data for the full period, relying solely on station-based measurements would result in minimal spatial variation in temperature.

2.3 Land Surface Temperature

To overcome this challenge, we take an alternative approach and use data products derived from satellite imagery that collect *land surface temperature* (LST), as opposed to *air temperature* (Tair), which is what ground weather stations measure. LST is constructed from spectrum bands that measure radiation emitted by the land surface observed by

[1] The values cited refer to IBGE (2010). The year 2010 is the mid-point in our analysis period and refers to the latest Population Census available.

satellites (Wan, 1999). We use Version 6 of the MYD11A1 data product derived from the MODIS (Moderate Resolution Imaging Spectroradiometer) instrument aboard NASA's Aqua satellite (Wan et al., 2015), which has provided a high-quality global LST product (Lian et al., 2017).[2] This data product provides daily per-pixel land surface temperature at a nominal pixel spatial resolution of 1km. Images have been captured at approximately 1pm local time every day since July 2002.

The use of thermal remote sensing has been an emergent trend in the natural sciences and environmental epidemiology, and applications of satellite-derived LST are now widely reported to characterize urban heat islands (Azevedo et al., 2016; Hu and Brunsell, 2013; De Ridder et al., 2012; Dousset et al., 2011; Rajasekar and Weng, 2009), to study areas of higher relative temperature within cities (Yin et al., 2023; Neiva et al., 2017; White-Newsome et al., 2013; Johnson and Wilson, 2009), and to estimate the near-surface temperature in the absence of meteorological stations (Chen et al., 2015; Good, 2015; Kilibarda et al., 2014).[3]

Conceptually, LST is the physical temperature of the top few micrometers underlying the Earth's surface. At the same time, Tair is the thermodynamic temperature of the air at the height of approximately 2 meters above the surface (Lian et al., 2017). Despite these fundamental differences, a growing stream of studies has shown that LST and Tair are strongly related (Good et al., 2017). Both worldwide and local analyses have validated satellite-derived LST by ground-truthing, which typically compares exposure assessments in a specific site covered by both LST images and ground stations. The findings indicate that cross-section and time-series correlations between LST and average Tair are often strong.[4] This relationship nevertheless depends on surface type, insolation, and elevation. LST becomes relatively higher than maximum Tair in more sparsely vegetated and bare areas and under stronger insolation (Good et al., 2017; Lian et al., 2017). This is the case of urban regions in the tropics, where LST often surpasses Tair records. Consistent with that, LST has been used as a reliable marker of heat islands and weather anomalies.[5]

In the MODIS data product we use, the city of Rio de Janeiro comprises 1,524 pixels. We use these pixels to construct population-weighted daily temperature and heat stress measures for the city's neighborhoods. Rio has officially 164 neighborhoods, but this has changed over time with the creation of new neighborhoods. To maintain spatial compa-

[2]The main alternative to MODIS is the LANDSAT satellite, which produces images once every 16 days. Since we need daily data, MODIS is better suited for our purposes.

[3]In the economics literature on the effects of temperature, this approach seems to be rare. We are only aware of the work of Aragón et al. (2021), which studies the impact of temperature on small farmers' input decisions in Peru.

[4]In particular, LST is more closely related to maximum Tair during the day.

[5]For instance, Good et al. (2017) analyze the August 2003 European heat wave and show that LST maps closely resemble the equivalent Tair over time and space.

rability, we aggregate the 164 neighborhoods into the original 144 units as defined in the early 2000s, and further restrict our analysis to neighborhoods with population of individuals above 60 years old greater than 500, resulting in a final sample of 137 neighborhoods.[6] We use geo-referenced information from Rio's 10,504 census tracts to assign each census tract to a pixel.[7] This procedure creates a daily panel of census tracts. Next, we use geo-referenced information on neighborhoods to assign each census tract to a neighborhood. Finally, we aggregate LST at the neighborhood level by taking the unweighted average over the neighborhood's census tracts. Since census tracts are designed to have roughly equal population size (IBGE, 2010), taking the unweighted average provides a close approximation to a population-weighted average temperature for each neighborhood-day.

Using satellite data to measure heat stress has both advantages and limitations. One major advantage is that satellite-derived land surface temperature (LST) offers high spatial resolution and full geographic coverage, allowing us to systematically track daily temperature at a much finer scale than weather stations typically permit. However, a key limitation is that satellites cannot retrieve reliable temperature readings under thick cloud coverage. As a result, missing LST observations can be pervasive, posing a challenge for the construction of a complete panel of neighborhood-level exposure. In our case, missings correspond to around half the total number of day-by-neighborhood observations.[8]

To address this issue, we adopt a two-step imputation strategy commonly used in the literature (e.g., Yu et al., 2022; Shiff et al., 2021; Auffhammer and Kellogg, 2011; Schlenker and Roberts, 2009). First, we use ground-based weather station data available in Rio de Janeiro and, for each neighborhood, we regress observed satellite-derived LST on contemporaneous weather variables such as air temperature, precipitation, wind speed, and humidity. We then use the fitted values from these regressions as our benchmark measure of LST exposure. Details on this imputation procedure, alternative imputation methods, and further analyses are provided in Appendix Section C.[9]

Finally, we construct temperature measures at a monthly frequency to compute indicators

[6]Small denominators lead to excessive volatility in mortality rates. We will consider that Rio has 137 neighborhoods for the remainder of the paper.

[7]We assign the census tract's centroid to the pixel by using centroid coordinates. Census tracts are defined by the Brazilian Institute of Geography and Statistics (IBGE) and typically contain around 100 households. Residential areas in Rio are fairly dense. Therefore, census tracts are small relative to 1km x 1km pixel size. A census tract is typically entirely contained in one pixel.

[8]There are 405,397 missing cases out of the total number of observations in our daily sample (N = 731,302). The number of missing cases therefore corresponds to 55% of the total, in line with values found in previous studies (Zhang et al., 2022).

[9]As also discussed in Appendix Section C, LST products that rely on spatial interpolation are subject to smoothing geographical variation and compressing cross-neighborhood temperature differences (e.g. Zhang et al., 2022).

for heat stress at the month–neighborhood level.[10] As our first indicator, we count the number of days in a neighborhood-month in which LST exceeds a given threshold. We use $LST = 43C$ as our benchmark for defining high heat exposure. As discussed below, this value corresponds approximately to $Tair = 36C$ under typical conditions of high solar radiation and low cloud cover, and matches the 95th percentile of the daily LST distribution in our data. This benchmark isolates heat events that are unlikely to reflect normal seasonal fluctuations, and is grounded in studies on tropical and subtropical climates which indicate that air temperature inflections around 36C are associated with increased health risks, particularly for older adults (Kephart et al., 2022a; Li et al., 2014; Falchetta et al., 2024). Yet, while $LST = 43C$ provides a conservative but epidemiologically meaningful cutoff for capturing harmful heat exposure, we also construct alternative thresholds to explore non-linearities in the heat–mortality relationship and measures of heat waves. To define a heat wave lasting k days ($k = \{3, 5, 7\}$), we count the number of consecutive days with temperature greater than T in the neighborhood-month.

2.4 LST and Air Temperature Data in Rio de Janeiro

As mentioned above, the existing evidence indicates that LST strongly correlates with Tair. We follow the literature and verify this relationship in our data by ground-truthing. Figure A1 shows how LST and Tair are related in Rio de Janeiro by using air temperature measurements from one of the ground stations in the city and comparing it to the satellite data we use for the exact location and the same time.[11] The plots show that the two measures are strongly correlated. LST is typically higher than air temperature, as expected during the day in urban areas under high insolation (Naserikia et al., 2023; Berg et al., 2021). A simple time-series OLS regression of LST on Tair at the daily level yields $\widehat{LST} = 1.565 + 1.146Tair$. This indicates that LST records in our data are expected to average approximately 14.6% higher than Tair, plus 1.5C in absolute terms (e.g., 36C in Tair is roughly equivalent to 42.8C in LST).

Table 1, panel A, shows descriptive statistics of our main temperature variables. The sample size corresponds to 23,016 month-by-neighborhood observations [14 years $\times$ 12 months $\times$ 137 neighborhoods]. The average neighborhood-month LST is 32.6C (SD 4.92C), with about 10 days per month recording LST up to 31C, 14 days with LST between 31C-43C, and approximately 2 days with LST higher than 43C (see Figure A2). We document substantial variation in LST across neighborhoods, as shown on the right map of Figure

[10] We use monthly rather than daily data to smooth noise in neighborhood-level mortality and mitigate the confounding influence of short-run displacement effects (Deschênes and Greenstone, 2011).

[11] For this comparison, we choose the ground station of Alerta Rio System, located in the neighborhood of São Cristóvão, as it provides hourly temperature, over a longer span of time among the existing stations.

1. For a given point in time, the average difference between the maximum and minimum LST across neighborhoods is approximately 13C, reaching about 17C in December-February. Consistent with official measures of Tair and complementary literature (e.g., Neiva et al., 2017; Rio de Janeiro, 2016), we find that average LST in neighborhoods in the North and West Zones are usually higher, in contrast to records in areas in the South Zone.

2.5 Innovations in Health Policies

The federal government designed the Family Health Program (FHP) in 1994 to achieve universal health coverage focusing on prevention and expansion of primary health care services (PHC). FHP decentralized the health care model through professional health care teams directly intervening at the community level in pre-determined catchment areas.[12] By the mid-2000s, FHP coverage in Brazil had already achieved more than 50% of the population, nearly 8 times higher than in the city of Rio. In December 2008, the city had only 126 Health Family teams, which covered approximately 7.1% of the population. This represented the second-lowest Family Health Program coverage of all capitals in the country. More generally, up to 2009, the city of Rio had the lowest public spending in health among all Brazilian capitals, 83% of the city's health municipal budget were allocated towards large hospitals, and there was an insufficient number of doctors in family medicine (Soranz et al., 2016; Harzheim et al., 2009).

In 2009, the newly elected mayor started a restructuring of the public health system in order to improve PHC coverage (Soranz et al., 2016; Campos et al., 2016). The Municipal Health Secretary (SMS, for *Secretaria Municipal de Saúde*) expanded the FHP coverage by improving the existing health facilities' infrastructure and by increasing the number of teams settled in these clinics. The most significant action came through the construction of new health units, called the Family Health Clinics (CF, for *Clínica da Família*). These units were built as large primary health care facilities that could allocate five or more family health teams, which could rely on equipment, services and technologies enabling greater service resolution (Harzheim et al., 2009). The number of FHP teams increased to 718 in 2012, reaching coverage of 40% of the population.

In parallel, starting in 2007 the state of Rio de Janeiro launched a new emergency healthcare policy with the opening of emergency care units, the *Unidades de Pronto Atendimento* (UPAs). These are facilities that provide curative, basic to intermediary complex emergency care services. Equipped with X-rays, electrocardiography, labs, and observation beds, they provide care for acute and chronic conditions, first aid for trauma and surgi-

[12] Each health team is responsible for 3,000 to 4,500 individuals within a determined geographic area, and should provide continuous assistance to families, including prevention, health promotion, health counseling, and home visits. See Bhalotra et al. (2019) for a detailed description of the program.

cal cases, and consultations for lower-severity issues (O'Dwyer et al., 2013; Konder and O'Dwyer, 2016). Their creation aimed to expand the availability of emergency care by reducing the average distance to emergency units. By the end of the policy rollout, UPAs had become the closest emergency facility for 50% of the population in the state of Rio de Janeiro, with 30 units introduced in the city of Rio alone (Bhalotra et al., 2023). Figure A3 shows the expansion of both Family Health Clinicas and UPAs across neighborhoods since the beginning of our sample period.

Both FHP and UPAs are part of the Unified Health System (*Sistema Único de Saúde*, SUS), a tax-funded, single-payer system that provides free-of-charge services at the point of care. However, while access to FHP strictly depends on whether patients reside within their catchment areas, UPAs are 24/7 walk-in facilities open to all.

3 Data on Health Outcomes and Auxiliary Information

Our empirical analysis uses a monthly neighborhood-level panel spanning January 2003 to December 2016, incorporating data on temperature and heat stress indicators, health outcomes, and access to healthcare services.[13]

3.1 Health Outcomes

We rely on data from the Brazilian National System of Mortality Records (Datasus/SIM) to compute mortality rates. SIM gathers information on every death officially registered in Brazil. We had access to administrative mortality data from the Rio de Janeiro's Municipal Health Secretary, which provided the counting of deaths by descendant's neighborhood of residence, age, and by diagnostic identified according to the International Classification of Diseases, 10th Revision (ICD-10). We follow Vaidyanathan et al. (2020) and Lee et al. (2023) to define heat-related chronic conditions as deaths due to cardiovascular, respiratory, endocrine, nutritional and metabolic causes.[14]

Data on the population by age groups are obtained from IBGE (2010), which provides the counting of residents by neighborhood in 2010 (the mid-point year of our period of analysis).[15] We compute mortality rates for each neighborhood and month by counting the deaths of individuals by age and cause of death and dividing by population in the

[13]Although Land Surface Temperature (LST) data are available from July 2002 onward, we begin our sample in January 2003 to capture 14 complete calendar years. We conclude in December 2016, shortly after the full implementation of the two major health system reforms in Rio discussed above.

[14]Chapters IX, X, and IV define these causes in ICD10, respectively.

[15]More specifically, the denominator of mortality rates is computed for 2010, and thus, it is fixed over time. We check whether the results are sensitive to this definition by adding to our specifications a combination of neighborhood-by-year fixed effects and neighborhood-specific linear trends.

same age group (per 100,000). Monthly mortality rates at the neighborhood level are then merged with LST variables to form our panel data at the neighborhood-by-month level.

Table 1, panel B, shows descriptive statistics for mortality causes. We observe a monthly neighborhood average of 181 deaths per 100,000 individuals aged 60 or older. Among these causes, mortality is more related to circulatory conditions (mean 110.4), followed by respiratory causes (mean 47.4) and by endocrine and metabolic causes (mean 23.1).

3.2 Access to Healthcare

We identify and geocode healthcare facilities in Rio, including hospitals, primary care, and emergency care units. We also geocode the catchment areas of Family Health Clinics, and compute the average distances from each census tract to the closest emergency care (ER) facility to measure access to health services. An emergency care facility is defined as either a hospital with a 24/7 walk-in ER or an UPA. Given that UPAs were the main innovation in policy regarding emergency care, changes in distances to ER come mainly due to the opening of UPAs.

The catchment areas of each Family Health Clinics were obtained from the Secretariat of Health Surveillance of the State of Rio de Janeiro. We calculate the share of the population in each neighborhood that reside in a Family Health Clinic catchment area by intersecting the catchment areas of each facility with the neighborhood shapefile provided by DataRio. Data on location and timing of the introduction of UPAs were provided by *Sala de Apoio a Gestao Estrategica* (SAGE), from the State of Rio de Janeiro, and intersected with the neighborhood shapefile from DataRio.

3.3 Auxiliary Data and Controls

We use other auxiliary datasets in our analysis. First, we add a series of controls at the neighborhood-month level to our analysis to absorb the potential influence of concurrent policies, with a focus on those related to security and public transportation.[16] This includes mapping Pacifying Police Units, which are part of a law enforcement program aimed at restoring territorial control of some areas in the city that were previously under the power of drug gangs. Data on the location and timing of the introduction of these units were obtained from the Institute of Public Security (ISP). We intersected them with the shapefile of neighborhoods to determine the number of units in each neighborhood-year-month. We also collected indicators for access to transportation. Data on Bus Rapid Transit (BRT) and the number of subway stations in each neighborhood were obtained from

[16]These policies were substantially expanded during our period of analysis in preparation for the 2014 World CUp and 2016 Olympics, which was held in Rio de Janeiro.

Pereira Passos Institute. We followed the same procedure of intersecting the shapefile of neighborhoods to compute the number of stations in each neighborhood-year-month.

Second, we add precipitation measures at the neighborhood level. We rely on daily rainfall data from the Alerta Rio system, which operates 33 rain ground stations across the city of Rio de Janeiro. Each of these stations is associated with a predefined pluviometric zone, as delineated by the Alerta Rio system.[17] We spatially matched these zones to neighborhood boundaries. Specifically, for each neighborhood, we calculated the area of intersection with each pluviometric zone and used these areas as weights to compute a daily precipitation average. This produced a neighborhood-level daily time series of rainfall based on the weighted contribution of overlapping zones. We then aggregated these data to the monthly level by counting the number of days in each month that precipitation fell into one of four bins based on the empirical distribution of the precipitation variable: no rain (0mm of rain), light rain (1mm to 4mm), moderate rain (5mm to 24mm) , and heavy rain (above 25mm). These categories allow us to capture not only total rainfall, but also the intensity and frequency of precipitation events.

4 Empirical Model

This paper examines the relationship between heat stress and mortality rates within the city of Rio de Janeiro. Our empirical strategy follows fixed effects models commonly used to estimate the causal impacts of temperature on health outcomes in different contexts (Deschenes, 2014; Deryugina and Hsiang, 2014; Barreca et al., 2016; Cohen and Dechezleprêtre, 2022). The equation below adapts this class of models to a neighborhood-by-month setting and serves as our benchmark specification:

$$h_{iym} = \sum_{j=1}^{J} \lambda_j B^j_{iym} + \alpha_{iy} + \delta_{im} + \beta_i t + \gamma' Z_{iym} + \epsilon_{iym} \tag{1}$$

Where h_{iym} denotes the health outcome in neighborhood i, year y, and month m. The term B^j_{iym} captures the number of days in a neighborhood-year-month when the daily LST falls into bin B^j. We construct bins ranging 6°C, starting from LST equal to 24°C or below, and ending in LST equal to 43°C or above. The omitted category is the bin with LST between 31°C and 37°C. In most tables, we focus on higher temperatures and report only the coefficients related to the number of days with LST above 43°C, leaving the complete set of estimates to auxiliary figures. In additional exercises we use alternative measures of heatwaves, alternative thresholds for the bin with the highest temperature, and different

[17] The polygons defining each pluviometric zone are publicly available at https://data.rio.

ranges for the bins.

Equation 1 includes neighborhood-by-year α_{iy} and neighborhood-by-calendar month δ_{im} fixed effects. The variable t represents a linear time trend (year-month). In some specifications we add neighborhood specific linear trends $\beta_i t$ to account for long-run variation in mortality and temperature. Time-varying control variables in Z_{iym} account for potentially confounding public policy interventions at the neighborhood level, including the share of the population covered by family health clinics, the average distance to the closest ER, and the number of Pacifying Police Units, BRT stations, and subway stations. We flexibly account for precipitation by incorporating bins of varying precipitation levels.

All regressions are weighted by neighborhood population size to reduce noise from small population cells. Standard errors are clustered at the neighborhood level. We also test robustness to clustering at broader spatial levels to account for serial correlation in weather shocks within larger micro-climatic areas.

In this setting, exposure to heat stress is neighborhood-specific. Neighborhood-by-year fixed effects account for time-varying neighborhood-level confounders such as fluctuations in economic conditions and general population health. Neighborhood-by-calendar month fixed effects control for month-specific temperature averages that vary across neighborhoods. This is analogous to controlling for typical seasonality at the neighborhood level. Neighborhood-by-calendar month fixed effects also control for seasonal variation in socioeconomic or employment conditions at the neighborhood level. Conditional on these two sets of fixed effects, plus controls, residual temperature variation is plausibly idiosyncratic and unexpected for that location and time. In our benchmark setup, identification therefore relies on exogenous temperature deviations from neighborhood-month historical averages, conditional on yearly neighborhood averages, neighborhood-specific linear trends and other time-varying controls.[18]

4.1 Common Effects *versus* Local Effects

We will compare the benchmark model to a more saturated specification that introduces time fixed effects (year-month). The latter absorb any city-wide shocks, including average

[18] Jones et al. (2026) have recently raised concerns about potential biases in temperature–outcome regressions using temperature bins, particularly in long panels and cross-regional settings. Our empirical design substantially mitigates these concerns along three dimensions. First, our analysis relies on a short panel spanning 2003–2016, limiting the accumulation of mechanically induced differential trends across bins. Second, identification exploits exclusively within-city variation across neighborhoods in Rio de Janeiro, where baseline temperature gradients and long-run structural heterogeneity are markedly smaller than in cross-regional contexts. Third, our baseline specification incorporates neighborhood-by-year fixed effects, absorbing all neighborhood-specific annual shocks and trends. This design effectively implements an extreme version of the short-panel and local-trend corrections proposed by the authors, restricting identification to highly localized, short-run deviations in temperature.

effects of heat waves, and shift identification toward cross-neighborhood differences in heat exposure within the same month. In other words, when we add time fixed effects, the remaining variation used for identification stems from relative differences in heat exposure within the city, across neighborhoods in the same time period.

Divergence in estimates between the two specifications will, therefore, allow us to assess the extent to which the estimated effects are driven by city-wide (common across neighborhoods) *versus* location-specific (heterogeneous across neighborhoods) variation in heat exposure, revealing the degree of spatial localization in heat impacts. More specifically, the comparison provides an informative contrast between two sources of variation in heat exposure. The benchmark specification (without time fixed effects) identifies effects using both city-wide temperature fluctuations and cross-neighborhood differences in exposure. The saturated specification (with time fixed effects) absorbs any common temporal shocks and identifies effects solely from spatial variation within a given month. The divergence between these estimates offers an approximate decomposition of the total mortality effect into: (i) a component attributable to variation in city-wide heat stress over time; and (ii) a component attributable to spatially heterogeneous exposure across neighborhoods at a given point in time. If the estimated coefficients remain stable across the two specifications, this implies that most of the variation in heat-related mortality is driven by differences across neighborhoods (i.e., localized effects). Conversely, if the coefficient magnitude declines substantially when conditioning on time fixed effects, this would indicate that a sizable portion of the heat-related mortality response reflects city-wide effects (e.g., shared exposure to an unusually hot month). In Appendix Section B we provide a more formal discussion on this interpretation.

5 Main Results

5.1 Main Effects on Mortality Rates

Table 2 presents the effects of four alternative measures of heat stress on mortality, with a focus on deaths from chronic diseases among the elderly. The first row reports results for our benchmark measure, the monthly number of hot days defined as those in which LST exceeds 43C. Column 1 includes year fixed effects, neighborhood-by-calendar-month fixed effects, and time-varying controls. Column 2 adds neighborhood-by-year fixed effects, yielding our benchmark specification.

Point estimates in columns 1 and 2 are similar. The coefficient in column 2 indicates that one additional hot day increases the mortality rate by 1.016 deaths per 100,000 individuals aged 60 and over. To interpret the magnitude of this effect, note that the average elderly

mortality rate due to chronic diseases is 181 per neighborhood-month. Given a mean mortality rate of 181 per 100,000, this effect corresponds to an increase of approximately 0.56%. Scaling to the city of Rio de Janeiro, which is home to nearly one million elderly residents, this estimate translates into roughly 10 additional deaths for each day above 43C, relative to an additional day in the omitted temperature bin (LST between 31°C and 37°C).

In the final column of Table 2, we include time fixed effects, which absorb monthly shocks common to all neighborhoods, such as the widespread effect of heat waves, and restrict identification to within-month, cross-neighborhood variation in exposure. This allows us to assess the extent to which heat-related mortality is driven by localized *versus* city-wide temperature variation. The point estimate in column 3 represents around 60% the coefficient of column 2. This result indicates that a meaningful share of the heat–mortality relationship is attributable to variation in exposure within the city.[19]

At the same time, the results also suggest that city-wide average temperatures also play a relevant role. Figures 2a and 2b provide a likely explanation for that. Figure 2a shows that during hotter months, most neighborhoods experience temperatures above their historical medians. Figure 2b reveals that the coefficient of variation in neighborhood-level temperature declines as city-wide temperature rises. This pattern suggests that extreme heat events compress spatial variation in exposure, making the city thermally more homogeneous.

The lower panel of Table 2 further supports this mechanism. Rows 2 through 4 report effects for heat waves defined as three, five, or seven consecutive days above 43C. In column 2, the estimated effects increase monotonically with heat-wave duration, reflecting the cumulative toll of sustained heat. In contrast, column 3 shows that the difference between columns 2 and 3 grows with heat-wave length, consistent with the idea that as average temperatures rise and exposure converges across neighborhoods, localized differences account for a smaller share of excess mortality, while common city-wide shocks become more important.

Figure 3 complements this analysis by reporting the full set of coefficients from the bin-based specification, of which only the >43C bin estimate is shown in Table 2. The figure reveals a nonlinear exposure–response pattern. Estimated effects are relatively flat at moderate temperatures, but rise in the upper tail of the temperature distribution. This upward trend remains visible even after introducing time fixed effects, though with the expected attenuation. Important to note that confidence intervals widen in the extreme bins due to

[19]For completeness, Table A1 shows that the effects of heat exposure are concentrated among chronic disease deaths in both specifications. There is no statistically detectable impact on mortality from infectious diseases or external causes.

more limited variation and lack of power, but point estimates at the top end remain elevated. Taken together, these results underscore the nonlinear nature of the heat–mortality relationship and the importance of both localized exposure and city-wide shocks in driving the observed effects.

5.2 Benchmarking

Tables A2 and A3 report results using alternative temperature thresholds and bin ranges. Estimates remain statistically significant and relatively stable across specifications. To facilitate comparison with existing studies based on air temperature (Tair), we note that a land surface temperature (LST) of 43C corresponds approximately to a Tair of 36C, while an LST of 39C approximates Tair of 32C.

At the global level, Carleton et al. (2022) estimate that a day above $Tair = 35C$ raises mortality by 4.7 per 100,000 for individuals aged 64+. Deschênes and Greenstone (2011) find increases in the annual age-adjusted elderly mortality rate in the U.S. by 0.35% on days above $Tair = 32C$ relative to a temperate day. Barreca et al. (2016) estimate that, in the earlier decades of their US sample, an additional day above $Tair = 32C$ was associated with increases in all-age mortality of up to 0.5 to 0.9 deaths per 100,000, though sensitivity to heat declined over time. Cohen and Dechezleprêtre (2022) show that a day above $Tair = 32C$ increases deaths from chronic causes among individuals aged 75+ by 0.108 per 100,000 in Mexico.

At our 39C threshold (32C Tair), we find that an additional hot day increases monthly mortality among individuals aged 60+ by 0.513 per 100,000, while at 43C (36C Tair) the effect is nearly twofold. While differences in temperature measurement, age groups, and outcome definitions complicate direct comparisons, we can place our results within the mid to lower range of the existing estimates. Several conjectures may help explain this pattern. First, our setting, a tropical city where high temperatures are frequent, may involve greater baseline physiological or behavioral adaptation to heat, attenuating mortality responses relative to temperate regions. Second, although we focus on elderly individuals, who are generally less mobile and less active in the labor market, intra-urban mobility may still allow avoidance of the hottest areas, reducing overall exposure. Third, as discussed below and in Appendix Section C, our benchmark LST imputation method is relatively conservative, and residual measurement error may attenuate the estimated effects. Nonetheless, columns 3 of Tables A2 and A3 reveal that substantial spatial heterogeneity in heat-related mortality persists within the city. Localized differences in exposure account for a meaningful share of total impacts, suggesting that even in climate-adapted and relatively mobile populations, neighborhood-level disparities remain a critical dimension of

climate vulnerability.

5.3 Robustness Checks

We now assess robustness with a focus on dynamics, inference, and measurement. Figure A4 presents two complementary exercises. The top row evaluates potential temporal displacement in deaths due to heat exposure, commonly referred to as "harvesting effects". Specifically, we examine whether temperature shocks in a given month are followed by reductions in mortality in subsequent months (m+1 and m+2). The bottom row tests whether temperature shocks in prior months (m-1 and m-2) are correlated with current mortality rates. Across all panels, we find no systematic or statistically significant patterns, suggesting that our contemporaneous estimates are unlikely to be driven by short-term mortality displacement or delayed temperature effects.

In Table A4, we examine the sensitivity of our inference to spatially correlated errors. Spatial autocorrelation is a potential concern in our setting, as nearby neighborhoods may experience similar weather conditions and share unobserved health shocks. To address this, we compute Conley standard errors, which account for correlation in the error structure based on geographic distance. We report results using distance cutoffs of 1.3km and 3.3km, corresponding to the median and the 90th percentile of the distribution of minimum distances between neighborhood centroids, respectively. Across all specifications, the adjusted standard errors remain relatively stable and the coefficients retain their statistical significance. This provides further reassurance that our estimates are not driven by spatial clustering in the data.

While concerns about reverse causality or omitted variable bias are limited in our empirical setting, a potential caveat relates to measurement error in the LST data. Because LST values retrieved from MODIS satellites may be missing under cloudy-sky conditions, the resulting measurement error is likely to be non-classical, potentially biasing our estimates. To address this issue, we adopt commonly used imputation methods for MODIS-based LST data. Appendix C provides technical details on the imputation procedure and presents auxiliary exercises testing the robustness of our findings to alternative strategies. The results remain qualitatively stable across methods, with our benchmark specification yielding the most conservative estimates. Importantly, the difference in effect size between specifications with and without time fixed effects remains consistent across imputation approaches, ranging from 35% to 58%. This pattern reinforces the view that both localized exposure and city-wide heat shocks contribute meaningfully to the observed mortality effects.

We discuss and assess alternative sources of data and methods to measure local-level

temperature in Appendix C. In particular, we show that alternative LST products that rely on spatial interpolation tend to smooth geographical variation and compress cross-neighborhood temperature differences (e.g. Zhang et al., 2022), which mechanically dampens local variation, especially under more spatially uniform heat events, thereby limiting their suitability for detecting fine-grained heterogeneity in exposure and impacts.

6 Mitigation through Locally Targeted Health Policies?

The preceding results highlight that a substantial share of heat-related mortality is driven by neighborhood-specific exposure. This finding naturally raises the question of whether localized investments in public health can mitigate the adverse effects of extreme heat. In particular, if heat stress triggers health deterioration via chronic conditions, then preventive care and timely emergency access may serve as key adaptation tools, especially in large cities where exposure varies across space.

Table 3 examines this conjecture by interacting our benchmark measure of days with high LST with two key health system indicators. First, we consider the share of the local population covered by the Family Health Program (FHP). Second, we include the average distance to the nearest emergency room (ER) as a proxy for access to emergency care.

Column 1, which does not include time fixed effects, suggests that preventive care mitigates the mortality risks associated with heat under both localized and city-wide temperature fluctuations. The interaction term between hot days and FHP coverage is negative and statistically significant. A one standard deviation increase in FHP coverage (0.21) reduces the estimated mortality impact of a hot day by approximately 45%. In contrast, the interaction with ER distance is positive and significant, suggesting that distance to emergency services is associated with larger heat-related mortality effects. Column 2 adds time fixed effects, absorbing time-specific shocks shared across all neighborhoods (e.g., city-wide heat waves). In this more saturated specification, the mitigating effect of preventive care becomes statistically insignificant. In contrast, access to emergency care continues to matter even under city-wide shocks, with point estimates remaining stable.

Taken together, these findings suggest that localized preventive care is effective in reducing mortality when exposure to city-wide fluctuations dominates. However, the effectiveness of local access to preventive care narrows while proximity to emergency services gains in importance when neighborhood-specific differences in exposure dominate. This underscores the importance of combining both localized preventive infrastructure and broad emergency readiness as complementary strategies for adapting to increasing frequency of heat waves.

7 Conclusion

This paper examined the health effects of extreme heat within a large tropical city, leveraging novel, high-resolution satellite imagery to measure temperature variation across neighborhoods in Rio de Janeiro. We showed that heat stress meaningfully increases mortality among older adults from chronic diseases, and that a relevant part of this effect is driven by localized exposure differences within the city. In our benchmark specification, one additional hot day raises elderly mortality by approximately 0.56% of the monthly mean. The attenuation of effect sizes under saturated fixed effects models suggests that nearly three-fifths of the impact of extreme heat can be traced to micro-spatial variation in exposure.

By documenting substantial within-city heterogeneity in heat–mortality relationships, our findings contribute to a growing literature on climate, health, and spatial inequality. While prior work has typically relied on variation across cities, states, or countries, we bring new evidence from intra-urban settings, where most of the world's population will reside in coming decades. From a policy perspective, our results underscore the importance of designing localized responses to environmental shocks. Preventive care appears most effective under city-wide heat waves, while access to emergency services becomes particularly important under neighborhood-specific variation in high-temperature exposure. These findings suggest that localized and systemic responses should be seen as complementary, not substitutive, in preparing urban health systems for climate risks. More broadly, the results highlight the need for urban adaptation strategies that combine targeted infrastructure investments (e.g., cooling centers, green spaces) with broader preparedness efforts (e.g., public alerts, heat-risk communication).

Finally, the paper reveals the value of combining satellite-based remote sensing with disaggregated health outcomes to better understand spatial inequalities in environmental risk. Future research can extend this approach to evaluate cost-effectiveness of alternative mitigation strategies, such as subsidies for cooling technologies, zoning reforms, or climate-sensitive urban planning, and to study other settings where sub-city variation in exposure may matter. As the burden of extreme heat accelerates, especially in tropical and low- and middle-income regions, improving our ability to measure, target, and mitigate localized climate risk will be essential.

References

Achebak, H., Devolder, D., and Ballester, J. (2018). Heat-related mortality trends under recent climate warming in Spain: A 36-year observational study. *PLoS medicine*, 15(7):e1002617.

Aragón, F. M., Oteiza, F., and Rud, J. P. (2021). Climate change and agriculture: Subsistence farmers' response to extreme heat. *American Economic Journal: Economic Policy*, 13(1):1–35.

Auffhammer, M., Hsiang, S. M., Schlenker, W., and Sobel, A. (2013). Using Weather Data and Climate Model Output in Economic Analyses of Climate Change. *Review of Environmental Economics and Policy*, 7(2):181–198.

Auffhammer, M. and Kellogg, R. (2011). Clearing the air? the effects of gasoline content regulation on air quality. *American Economic Review*, 101(6):2687–2722.

Azevedo, J. A., Chapman, L., and Muller, C. L. (2016). Quantifying the Daytime and Night-Time Urban Heat Island in Birmingham, UK: A Comparison of Satellite Derived Land Surface Temperature and High Resolution Air Temperature Observations. *Remote Sensing*, 8(2):153.

Ballester, J., Quijal-Zamorano, M., Méndez Turrubiates, R. F., Pegenaute, F., Herrmann, F. R., Robine, J. M., Basagaña, X., Tonne, C., Antó, J. M., and Achebak, H. (2023). Heat-related mortality in Europe during the summer of 2022. *Nature medicine*, 29(7):1857–1866.

Banerjee, R. and Maharaj, R. (2020). Heat, infant mortality, and adaptation: Evidence from India. *Journal of Development Economics*, 143:102378.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2016). Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy*, 124(1):105–159.

Berg, A. et al. (2021). Linking remotely sensed land surface temperature with ground-based air temperature across urban areas. *Remote Sensing*, 14(1):165.

Bhalotra, S., Rocha, R., Soares, R. R., et al. (2019). Can universalization of health work? evidence from health systems restructuring and expansion in Brazil. *IZA DP*, 12111.

Bhalotra, S. R., Nunes, L., and Rocha, R. (2023). Emergency care centers, hospital performance and population health. *Working paper*.

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2017). Weather, climate change and death in India. *University of Chicago working paper*.

Burke, M., Wilson, A. J., Avirmed, T., Wallstein, J., Martins, M. C., Behrer, P., Callahan, C. W., Childs, M., Choi, J., French, K., et al. (2025). Understanding and addressing temperature impacts on mortality. Technical report, National Bureau of Economic Research.

Campos, C. E. A., Cohn, A., and Brandão, A. L. (2016). Trajetória histórica da organização sanitária da cidade do Rio de Janeiro: 1916-2015. cem anos de inovações e conquistas. *Ciência & Saúde Coletiva*, 21:1351–1364.

Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Kopp, R. E., McCusker, K. E., Nath, I., et al. (2022). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *The Quarterly Journal of Economics*, 137(4):2037–2105.

Chen, F., Liu, Y., Liu, Q., and Qin, F. (2015). A statistical method based on remote sensing for the estimation of air temperature in China. *International Journal of Climatology*, 35(8):2131–2143.

Cohen, F. and Dechezleprêtre, A. (2022). Mortality, temperature, and public health provision: evidence from Mexico. *American Economic Journal: Economic Policy*, 14(2):161–92.

De Ridder, K., Bertrand, C., Casanova, G., and Lefebvre, W. (2012). Exploring a New Method for the Retrieval of Urban Thermophysical Properties Using Thermal Infrared Remote Sensing and Deterministic Modeling. *Journal of Geophysical Research: Atmospheres*, 117(D17).

Deryugina, T. and Hsiang, S. (2014). Does the Environment Still Matter? Daily Temperature and Income in the United States. Technical Report w20750, National Bureau of Economic Research, Cambridge, MA.

Deschenes, O. (2014). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, 46:606–619.

Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–85.

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

Dousset, B., Gourmelon, F., Laaidi, K., Zeghnoun, A., Giraudet, E., Bretin, P., Mauri, E., and Vandentorren, S. (2011). Satellite Monitoring of Summer Heat Waves in the Paris Metropolitan Area. *International Journal of Climatology*, 31(2):313–323.

Falchetta, G. et al. (2024). Global projections of heat exposure risk for older adults. *Nature Communications*, 15:1–10.

Good, E. (2015). Daily Minimum and Maximum Surface Air Temperatures from Geostationary Satellite Data. *Journal of Geophysical Research: Atmospheres*, 120(6):2306–2324.

Good, E. J., Ghent, D. J., Bulgin, C. E., and Remedios, J. J. (2017). A Spatiotemporal Analysis of the Relationship Between Near-Surface Air Temperature and Satellite Land Surface Temperatures Using 17 Years of Data from the ATSR Series. *Journal of Geophysical Research: Atmospheres*, 122(17):9185–9210.

Hajat, S., Vardoulakis, S., Heaviside, C., and Eggen, B. (2014). Climate change effects on human health: projections of temperature-related mortality for the uk during the 2020s, 2050s and 2080s. *J Epidemiol Community Health*, 68(7):641–648.

Han, J., Fang, S., Mi, Q., Wang, X., Yu, Y., Zhuo, W., and Peng, X. (2024). A time-continuous land surface temperature (lst) data fusion approach based on deep learning with microwave remote sensing and high-density ground truth observations. *Science of The Total Environment*, 914:169992.

Harzheim, E., Lima, K. M., and Hauser, L. (2009). Reforma da atenção primária à saúde na cidade do rio de janeiro: avaliação dos primeiros três anos de clínicas da família: pesquisa avaliativa sobre aspectos de implantação, estrutura, processo e resultados das clínicas da família na cidade do rio de janeiro.

Hooker, J., Duveiller, G., and Cescatti, A. (2018). A global dataset of air temperature derived from satellite remote sensing and weather stations. *Scientific data*, 5(1):1–11.

Hu, L. and Brunsell, N. A. (2013). The Impact of Temporal Aggregation of Land Surface Temperature Data for Surface Urban Heat Island (SUHI) Monitoring. *Remote Sensing of Environment*, 134:162–174.

IBGE, I. (2010). Censo demográfico 2010. *IBGE: Insituto Brasileiro de Geografia e*.

IETS (2015). *Painel Regional: Rio de Janeiro de Bairros*. Observatório Sebrae/RJ & IETS.

Johnson, D. P. and Wilson, J. S. (2009). The Socio-Spatial Dynamics of Extreme Urban Heat Events: The Case of Heat-Related Deaths in Philadelphia. *Applied geography*, 29(3):419–434.

Jones, B., Moscona, J., Olken, B. A., and von Dessauer, C. (2026). With or without u? binning bias and the causal effects of temperature extremes. Technical report, National Bureau of Economic Research.

Kephart, J. L. et al. (2022a). Global associations between extreme heat and mortality: A multi-country study. *Nature Medicine*, 28:1–8.

Kephart, J. L., Sánchez, B. N., Moore, J., Schinasi, L. H., Bakhtsiyarava, M., Ju, Y., Gouveia, N., Caiaffa, W. T., Dronova, I., Arunachalam, S., et al. (2022b). City-level impact of extreme temperatures and mortality in Latin America. *Nature medicine*, 28(8):1700–1705.

Kilibarda, M., Hengl, T., Heuvelink, G. B., Gräler, B., Pebesma, E., Perčec Tadić, M., and Bajat, B. (2014). Spatio-temporal Interpolation of Daily Temperatures for Global Land Areas at 1 km Resolution. *Journal of Geophysical Research: Atmospheres*, 119(5):2294–2313.

Konder, M. and O'Dwyer, G. (2016). The integration of the Emergency Care Units (UPA) with healthcare services in the city of Rio de Janeiro, Brazil. *Interface*, 20(59):879–892.

Lee, M. J., McLean, K. E., Kuo, M., Richardson, G. R., and Henderson, S. B. (2023). Chronic diseases associated with mortality in british columbia, canada during the 2021 western north america extreme heat event. *GeoHealth*, 7(3):e2022GH000729.

Li, B., Liang, S., Ma, H., Liu, X., He, T., and Zhang, Y. (2024). Generation of global 1 km all-weather instantaneous and daily mean land surface temperature from modis data. *Earth System Science Data Discussions*, 2024:1–45.

Li, T. et al. (2014). Temperature and mortality in 31 chinese cities: A threshold multicity time-series analysis. *Environmental Health*, 13:65.

Lian, X., Zeng, Z., Yao, Y., Peng, S., Wang, K., and Piao, S. (2017). Spatiotemporal Variations in the Difference Between Satellite-Observed Daily Maximum Land Surface Temperature and Station-Based Daily Maximum Near-Surface Air Temperature. *Journal of Geophysical Research: Atmospheres*, 122(4):2254–2268.

Matsuura, K. and Willmott, C. (2018). Terrestrial air temperature: 1900-2017 gridded monthly time series. Technical report, University of Delaware. University of Delaware. http://climate.geog.udel.edu/ climate/.

Mildrexler, D. J., Zhao, M., and Running, S. W. (2011). A Global Comparison Between Station Air Temperatures and MODIS Land Surface Temperatures Reveals the Cooling Role of Forests. *Journal of Geophysical Research: Biogeosciences*, 116(G3).

Naserikia, M. et al. (2023). Urban form and the dynamic relationship between land surface and air temperature. *Science of The Total Environment*, 894:165–183.

Neiva, H. D. S., Da Silva, M. S., and Cardoso, C. (2017). Analysis of Climate Behavior and Land Use in the City of Rio de Janeiro, RJ, Brazil. *Climate*, 5(3):52.

O'Dwyer, G., Konder, M., Machado, C., Alves, C., and Alves, R. (2013). The current scenario of emergency care policies in Brazil. *BMC health services research*.

Rajasekar, U. and Weng, Q. (2009). Spatio-temporal modelling and analysis of urban heat islands by using landsat tm and etm+ imagery. *International Journal of Remote Sensing*, 30(13):3531–3548.

Rio de Janeiro (2016). Climate change adaptation strategy for the city of Rio de Janeiro. *Prefeitura do Rio de Janeiro*.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598.

Serra, A. B. and Ratisbonna, L. (1941). *O Clima do Rio de Janeiro*. Ministerio da Agricultura, Serviço de Meteorologia.

Shiff, S., Helman, D., and Lensky, I. M. (2021). Worldwide continuous gap-filled modis land surface temperature dataset. *Scientific Data*, 8(1):1–10.

Soranz, D., Pinto, L. F., and Penna, G. O. (2016). Eixos e a reforma dos cuidados em atenção primária em saúde (rcaps) na cidade do Rio de Janeiro, Brasil. *Ciência & Saúde Coletiva*, 21:1327–1338.

UN (2018). 2018 revision of world urbanization prospects.

Vaidyanathan, A., Malilay, J., Schramm, P., and Saha, S. (2020). Heat-related deaths—united states, 2004–2018. *Morbidity and Mortality Weekly Report*, 69(24):729.

Wan, Z. (1999). Modis land-surface temperature algorithm theoretical basis document (lst atbd). *National Aeronautics and Space US Department of Commerce*.

Wan, Z., Hook, S., and Hulley, G. (2015). MYD11A1 MODIS/Aqua Land Surface Temperature/Emissivity Daily L3 Global 1km SIN Grid V006. [Data set] NASA EOSDIS Land Processes DAAC.

White-Newsome, J. L., Brines, S. J., Brown, D. G., Dvonch, J. T., Gronlund, C. J., Zhang, K., Oswald, E. M., and O'Neill, M. S. (2013). Validating Satellite-Derived Land Surface Temperature with in Situ Measurements: A Public Health Perspective. *Environmental Health Perspectives*, 121(8):925–931.

Yin, Y., He, L., Wennberg, P. O., and Frankenberg, C. (2023). Unequal exposure to heatwaves in Los Angeles: Impact of uneven green spaces. *Science*, 9(17):eade8501.

Yu, P., Zhao, T., Shi, J., Ran, Y., Jia, L., Ji, D., and Xue, H. (2022). Global spatiotemporally continuous modis land surface temperature dataset. *Scientific Data*, 9(1):143.

Zhang, T., Zhou, Y., Zhao, K., Zhu, Z., Chen, G., Hu, J., and Wang, L. (2022). A global dataset of daily near-surface air temperature at 1-km resolution (2003–2020). *Earth System Science Data Discussions*, 2022:1–18.

Figures and Tables

Figure 1: Rio de Janeiro: Land Use and Heat Map

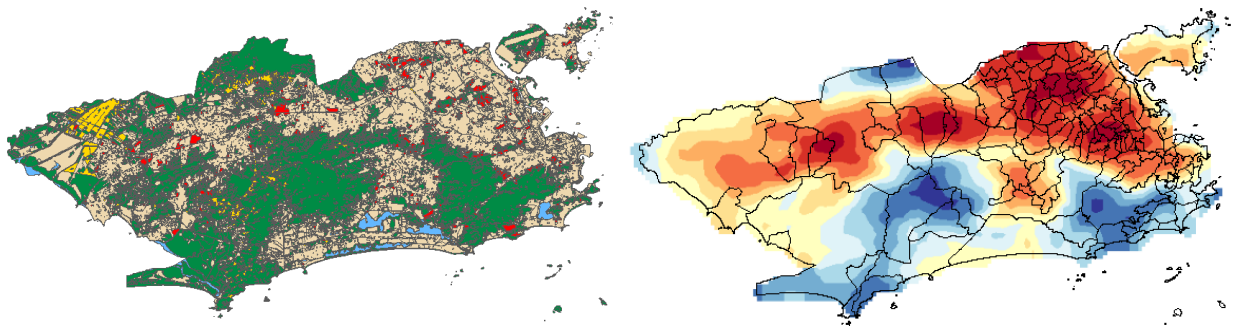

(a) Land Use (b) Average LST

Data on land use are publicly available from Institute Pereira Passos data lake (`data.rio`). We construct average LST at the pixel level over the entire period of analysis. In subfigure (a), areas under forested massifs are represented by green, urban areas are represented by light yellow, light blue shows lagoons, and red represents slums. In subfigure (b), darker blue colors mean colder pixels, while darker red colors mean hotter pixels.

Figure 2: Variation of LST temperature

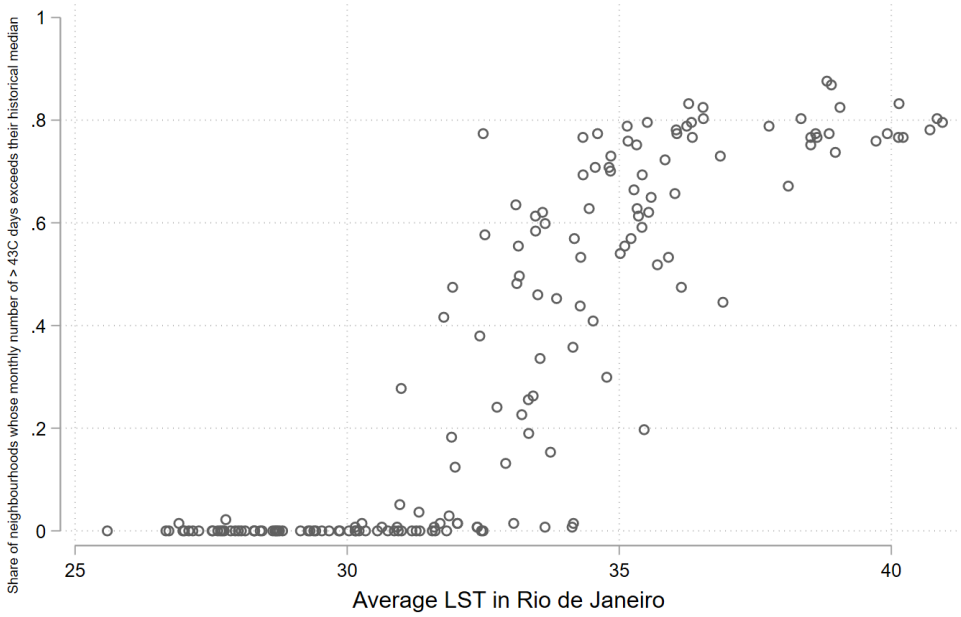

(a) Number of neighborhoods with high temperatures and average LST temperature in Rio

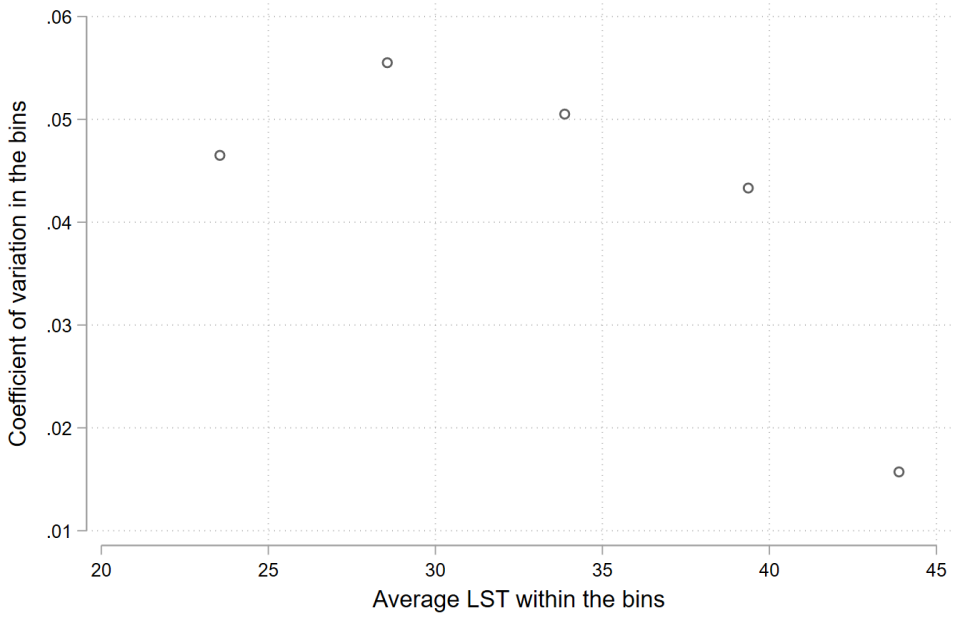

(b) Coefficient of variation and average LST in the bins

Figure 2a shows the relationship between the number of neighborhoods with temperature above 43C that are above the median value of the variable number of days above 43C and the average LST temperature in Rio de Janeiro. For Figure 2b, we created 5 bins for neighborhoods' temperature: below 25C, between 25C and 31C, between 31C and 37C, between 37C and 43C, and above 43C. For each bin, we calculated the coefficient of variation and the average LST temperature. Figure 2b shows the relationship between these two values.

Table 1: Summary Statistics of Neighborhood-Month Panel

	Mean	sd
Panel A - Temperature Measures		
Number of Days with High Temperature	1.77	3.67
Number of 3-day Heat Waves	0.15	0.51
Number of 5-day Heat Waves	0.04	0.23
Number of 7-day Heat Waves	0.02	0.14
Panel B - Health Outcomes		
Chronic death rate per 100,000 individuals aged 60+	180.95	100.26
Circulatory death rate per 100,000 individuals aged 60+	110.42	71.83
Respiratory death rate per 100,000 individuals aged 60+	47.41	43.25
Endometab. death rate per 100,000 individuals aged 60+	23.12	30.06
Infectious death rate per 100,000 individuals aged 60+	15.20	23.56
External death rate per 100,000 individuals aged 60+	12.84	21.40
Panel C - Controls		
Share population covered FHC	0.09	0.21
Avg. Distance to ER (km)	2.42	2.14
Number Bus Rapid Transit stations in the neighborhood	0.18	0.98
Number subway stations in the neighborhood	0.02	0.18
Number of Pacifying Police Units in the neighborhood	0.24	0.70
Number of days with no rain	18.46	4.90
Number of days with light rain	6.72	3.17
Number of days with moderate rain	4.16	2.57
Number of days with heavy rain	1.09	1.30

Notes: Data are organized as a neighborhood-by-month panel. Temperature measures are constructed from daily land surface temperature observations aggregated to the monthly level. Heat waves are defined as consecutive days with maximum daily land surface temperature exceeding 43C for at least 3, 5, or 7 days. Health outcomes are monthly mortality rates per 100,000 individuals aged 60 or older, by cause of death, using death registry data. Control variables include measures of healthcare access (coverage by community health facilities, average distance to the nearest emergency room, and counts of bus, subway, and UPPs stations in the neighborhood) and weather controls (number of days with no rain, light rain, moderate rain, and heavy rain).

Table 2: Effects of Hot Days on Mortality due to Chronic Causes

Shock	Chronic Deaths per 100,000 individuals aged 60+		
	(1)	(2)	(3)
1. Number of Days with High Temperature	1.051	1.016	0.590
	(0.200)***	(0.198)***	(0.277)**
2. Number of 3-day Heat Waves	6.490	6.470	4.057
	(1.140)***	(1.119)***	(1.268)***
3. Number of 5-day Heat Waves	8.045	8.102	2.618
	(2.084)***	(2.074)***	(2.371)
4. Number of 7-day Heat Waves	12.418	12.079	3.823
	(3.620)***	(3.627)***	(4.013)
Year	✓		
Neighborhood x calendar Month	✓	✓	✓
Neighborhood x Year		✓	✓
Year x Month (time)			✓
Controls	✓	✓	✓
Observations	23,016	23,016	23,016
Mean dep. var.	181.0	181.0	181.0

Notes: Data is a monthly panel of neighborhoods. Each row shows the effect of an alternative measure of a temperature shock. In row 1, the shock variable is the number of days with LST greater than 43C for a neighborhood i in year t and month m. In this specification, we also control for other bins of temperature, and the omitted category is the bin with temperatures between 31C and 37C. In rows 2 through 4, the shock is the number of events with 3, 5, or 7 consecutive days of temperatures exceeding 43 C.
Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
Standard errors clustered at neighborhood level in parenthesis. All regressions are weighted by the population aged 60+ in each neighborhood.
We include specific neighborhood linear trend in specification in column (1) to account for long-run changes between temperature and mortality. These trends are collinear in specifications (2) and (3). Controls include the share of population covered by family health clinics (CSF), the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month.
* significant at 10%; ** significant at 5%; *** significant at 1%.

Figure 3: Effects of Daily Temperatures on Mortality due to Chronic Causes in Population 60+ by bins of temperature

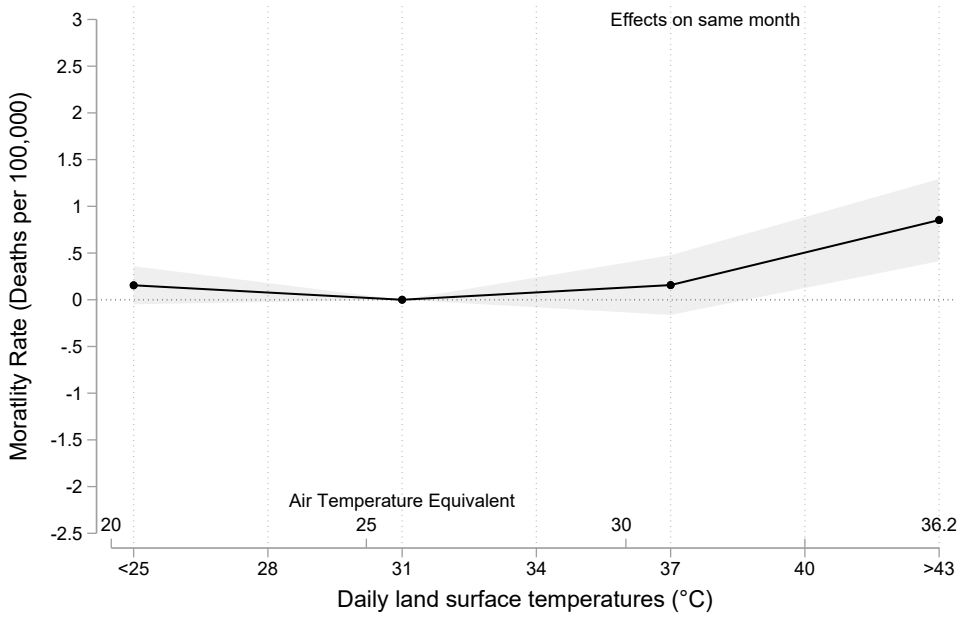

(a) Without time FE

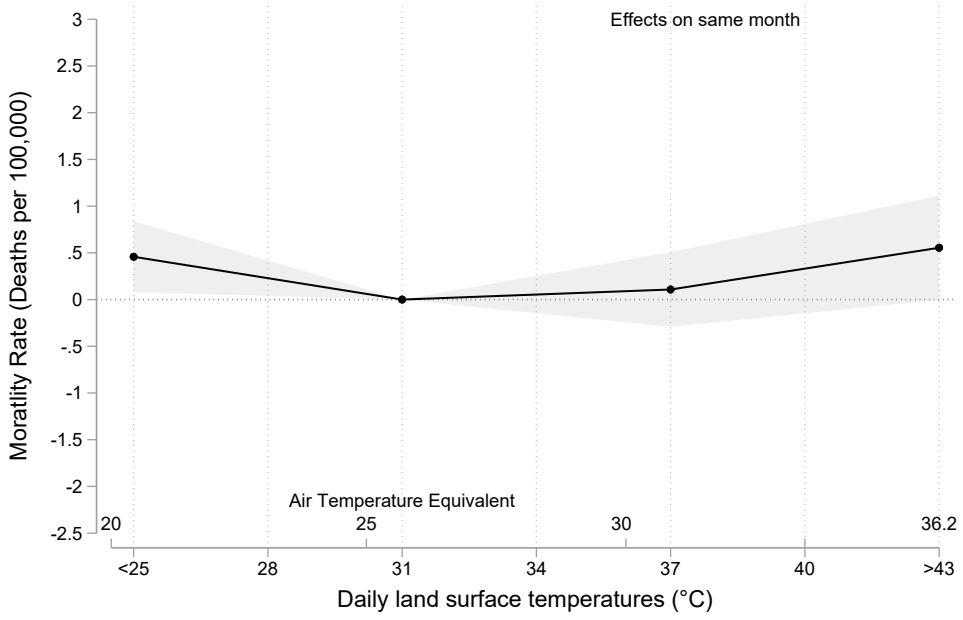

(b) With time FE

Estimated contemporaneous effects of daily temperatures, binned in six-degree days on mortality rates due to chronic conditions. The response function is normalized, with the 31C–37C temperature category set to zero. This ensures that each estimate represents the estimated impact of an additional day falling within a specific temperature bin on the monthly mortality rate, relative to days with temperatures between 31C to 37C.

Table 3: Mitigation Policies to Temperature shocks on Mortality due to Chronic Causes

	Chronic Deaths (per 100,000 individuals aged 60+)	
	(1)	(2)
Number of Days with High Temperature	0.818	0.217
	(0.332)**	(0.382)
Number of Days with High Temperature × Preventive Care	-1.756	-0.383
	(0.658)***	(0.700)
Number of Days with High Temperature × Distance to ER	0.243	0.248
	(0.134)*	(0.147)*
Neighborhood x Year	✓	✓
Neighborhood x calendar Month	✓	✓
Year x calendar Month (time)		✓
Controls	✓	✓
Observations	23,016	23,016
Mean dep. var.	181.0	181.0

Notes: Data is a monthly panel of neighborhoods.
Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
Standard errors clustered at neighborhood level in parenthesis. All regressions are weighted by the population aged 60+ in each neighborhood.
Linear trend refers to a specific neighborhood linear trend. Controls include the share of the population covered by family health clinics, the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month.
* significant at 10%; ** significant at 5%; *** significant at 1%.

A Additional Figures and Tables

Figure A1: LST and Air temperature in Rio de Janeiro

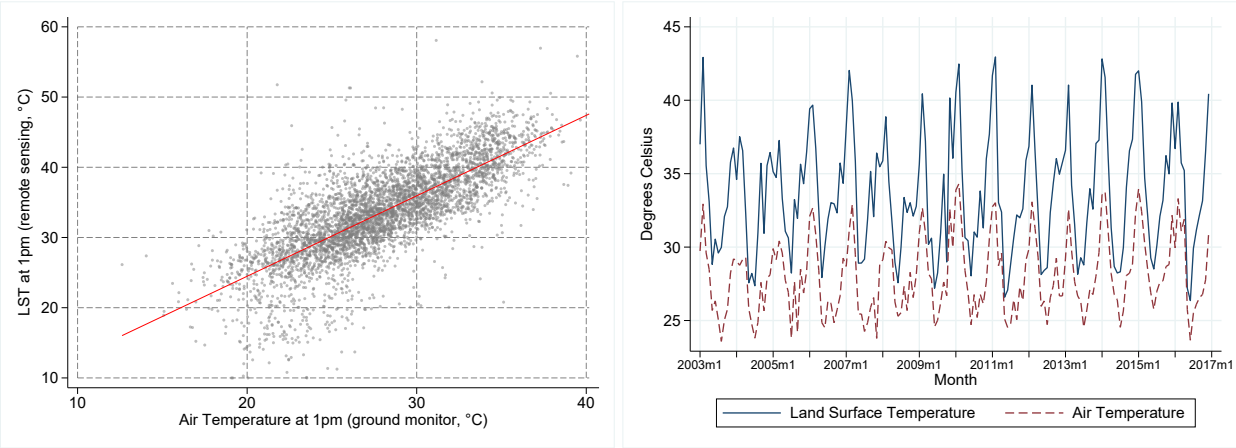

(a) Daily Measures (b) Monthly Averages

These figures use temperature data from two sources. Land surface temperature is measured by the MODIS sensor aboard of NASA's Aqua satellite, which overpasses Rio de Janeiro approximately around 1pm local time daily. We construct our benchmark imputation land surface temperature measures by neighborhood according to the procedure described in Appendix C. We select the São Cristóvão, for comparison with the air temperature data. Air temperature is measured by São Cristóvão's ground station, which provides hourly temperature readings. We select the 1pm reading for comparison with the land surface temperature data. Figure A1a shows the daily measures of LST and air temperature at 1pm. The red line represents the linear fit $\widehat{LST_t} = 1.565 + 1.146 AirTemperature_t$ (R^2 = 0.58). Figure A1b aggregates these data by taking monthly averages. In the Appendix, we provide further analysis of correlations between different temperature measures.

Figure A2: Neighborhood-Day Monthly Temperature Distribution

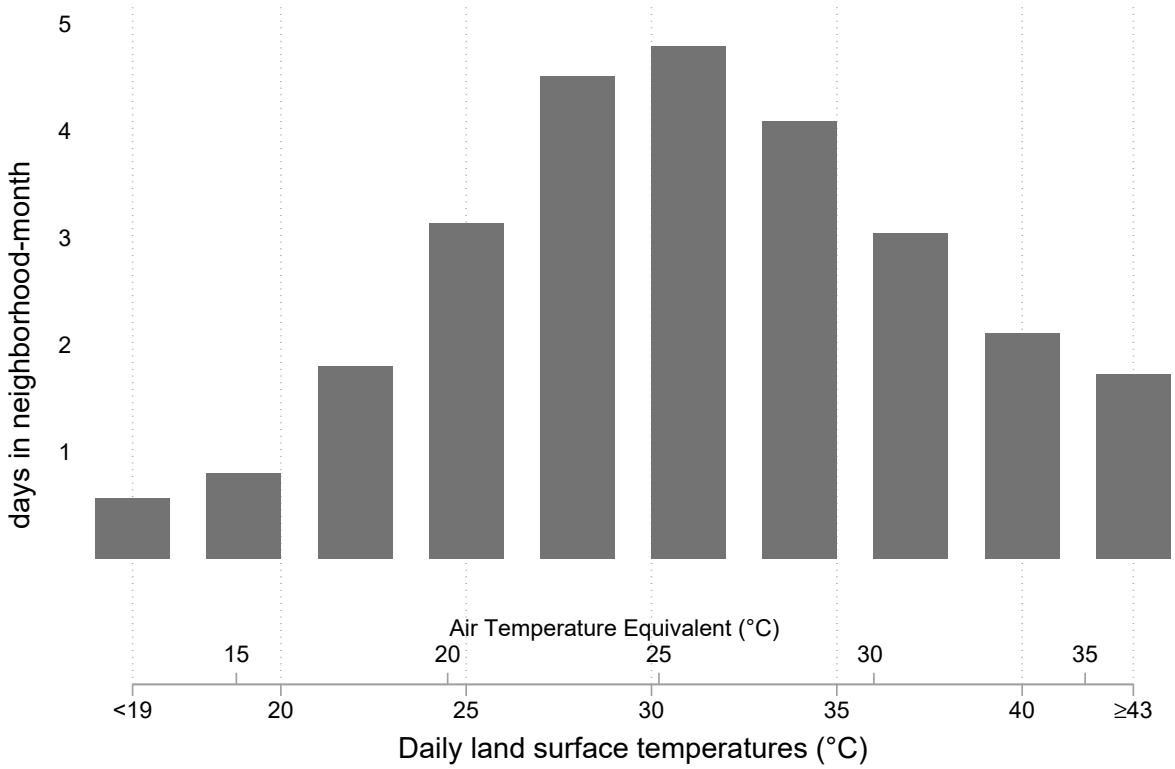

Histogram of neighborhood-day observations by daily land surface temperature bin. The outcome is expressed as the average number of days per neighborhood-month falling within each 3-degree range LST interval. The lower axis reports land surface temperature bins (°C), with the secondary axis providing the corresponding air temperature equivalents.

Figure A3: Expansion of health infrastructure over time and across neighborhoods in Rio

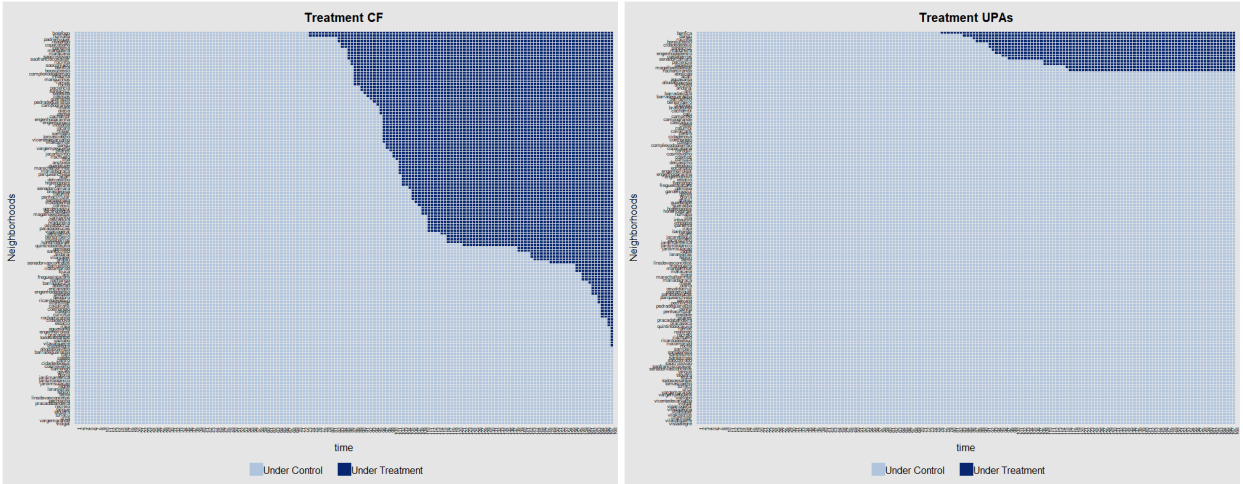

(a) (b)

This figure shows the variation over time and across space that we explore to construct the measures of distance to ER and percentage of individuals covered by primary health care. *CF* refers to the Primary Health Care units (*Clinicas da Familia*) and *UPAs* to ER rooms that were open during the sample period. The y-axis refers to the neighborhoods and the x-axis to the time since the start of the panel.

Figure A4: The Dynamic Effects of Daily Temperatures on Mortality due to Chronic Causes in Population 60+

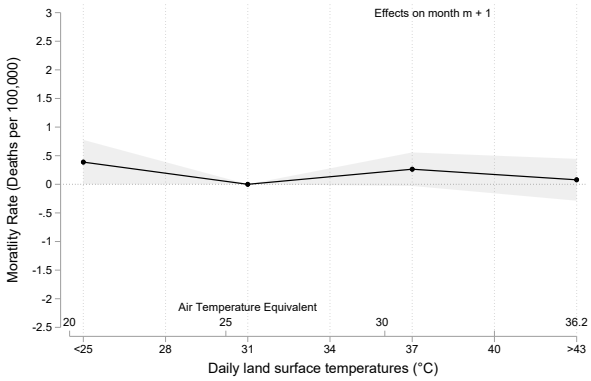

(a) Effects on month m + 1

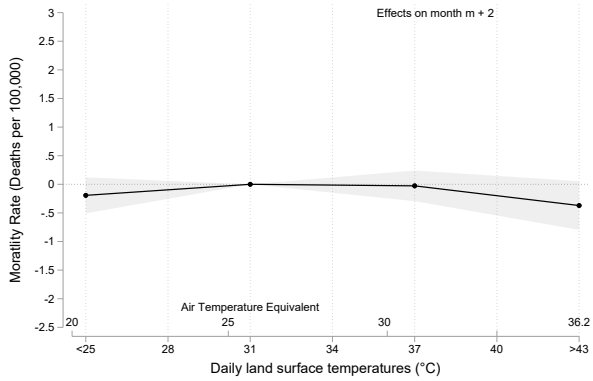

(b) Effects on month m + 2

This subfigure plots the coefficients from the specification (1) for the effects of contemporaneous daily temperatures, binned into six-degree bins, on mortality due to chronic causes in the next two months. The omitted category is the bin with temperatures between 31C and 37C. This specification allows for displacement (or "harvesting") and delayed effects of temperature exposure.

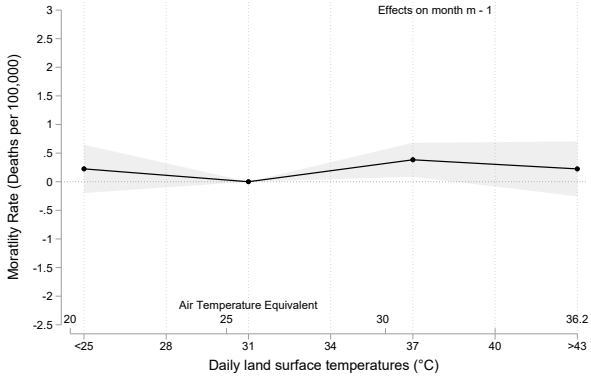

(c) Effects on month m - 1

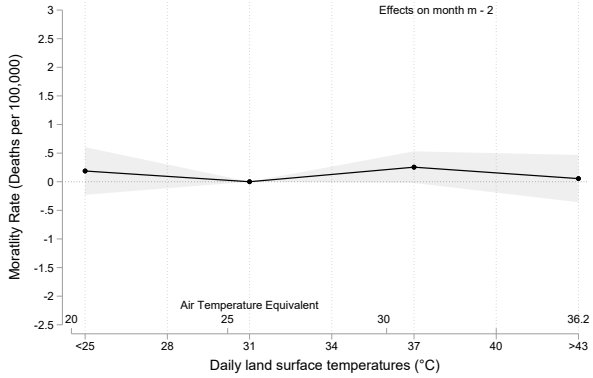

(d) Effects on month m - 2

This subfigure plots the coefficients from the specification (1) for the effects of contemporaneous daily temperatures, binned into six-degree bins, on mortality due to chronic causes in the previous two months. The omitted category is the bin with temperatures between 31C and 37C. Time fixed effects are added in the regression. The coefficients associated with temperature in this specification should not affect the current mortality rate. This works as a placebo test.

Figure A5: The Dynamic Effects of Daily Temperatures on Mortality due to Chronic Causes in Population 60+ with Time FE

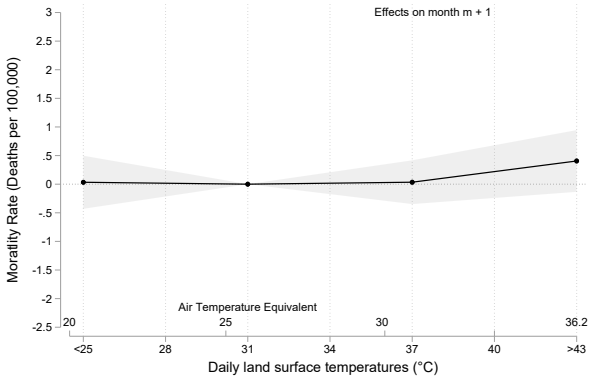

(a) Effects on month m + 1

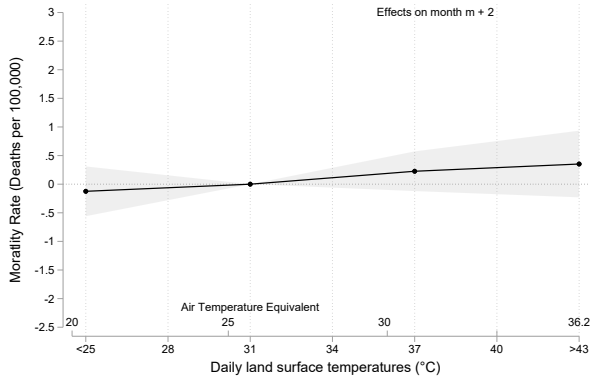

(b) Effects on month m + 2

This subfigure plots the coefficients from the specification (1) for the effects of contemporaneous daily temperatures, binned into three-degree bins, on mortality due to chronic causes in the next two months. This specification allows for displacement (or "harvesting") and delayed effects of temperature exposure.

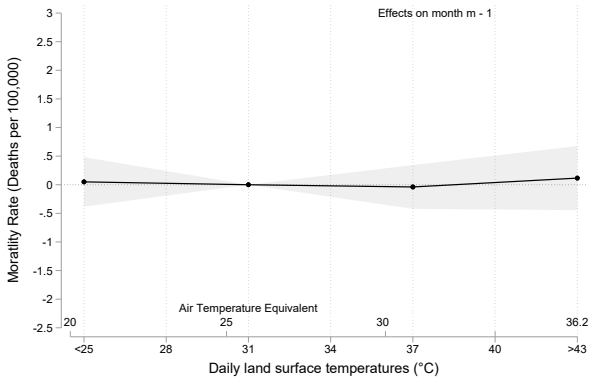

(c) Effects on month m - 1

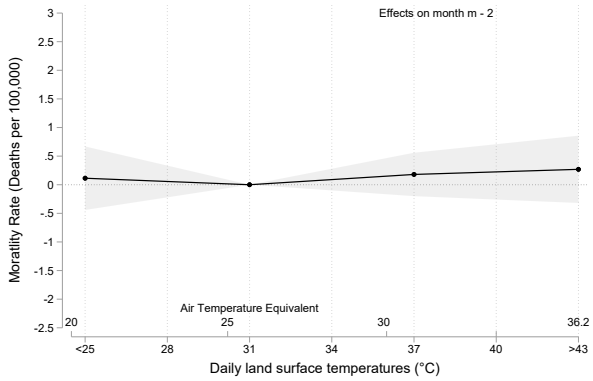

(d) Effects on month m - 2

This subfigure plots the coefficients from the specification (1) for the effects of contemporaneous daily temperatures, binned into three-degree bins, on mortality due to chronic causes in the previous two months. Time fixed effects are added in the regression. The coefficients associated with temperature in this specification should not affect the current mortality rate. This works as a placebo test.

Table A1: Effects of Hot Days on Mortality by Cause of Death on Population 60+

	Chronic	Infectious	External
Panel A - Without time fixed effects			
Number of Days w/ High Temp.	1.016	-0.050	0.035
	(0.198)***	(0.053)	(0.048)
Panel B - With time fixed effects			
Number of Days w/ High Temp.	0.590	-0.037	-0.020
	(0.277)**	(0.071)	(0.069)
Mean	181.0	15.20	12.84
Observations	23,016	23,016	23,016
Neighborhood x calendar Month	✓	✓	✓
Neighborhood x Year	✓	✓	✓
Controls	✓	✓	✓

Notes: Data is a monthly panel of neighborhoods. Each column is a separate regression for the dependent variable listed. All regressions are weighted by the corresponding subgroup population in the neighborhood. Standard errors clustered at the neighborhood level in parentheses.Chronic deaths include circulatory (ICD-10 Chapter IX), respiratory (Chapter X), and endocrine/metabolic diseases (Chapter IV). Infectious = Chapters I; External = Chapters XX. Controls include Family Health Clinic coverage in the neighborhood, the average distance of the neighborhood to the closest ER room, and public transport infrastructure. Regressions include neighborhood-specific linear trends.
* $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table A2: Robustness for different thresholds for the highest temperature bin: effects of Temperature on Mortality due to Chronic Causes

Shock	Mortality Rate (per 100,000 individuals aged 60+)		
	(1)	(2)	(3)
Panel A: 43C+ (benchmark)			
Number of Days with High Temperature	1.051	1.016	0.590
	(0.200)***	(0.198)***	(0.277)**
Panel B: 41C+			
Number of Days with High Temperature	0.719	0.628	0.420
	(0.156)***	(0.156)***	(0.215)*
Panel C: 39C+			
Number of Days with High Temperature	0.592	0.513	0.361
	(0.137)***	(0.141)***	(0.197)*
Year	✓		
Neighborhood x calendar Month	✓	✓	✓
Neighborhood x Year		✓	✓
Year x Month (time)			✓
Controls	✓	✓	✓
Observations	23,016	23,016	23,016
Mean dep. var.	181.0	181.0	181.0

Notes: Notes: Data is a monthly panel of neighborhoods. Each panel shows the effect of an alternative threshold for the bin with the highest temperature. In Panel A, we show the benchmark, using the number of days with LST above 43C, and the omitted category is the bin with temperatures between 31C and 37C. In panel B, the threshold is LST 41C, which represents approximately the 90-percentile of the distribution of average LST across the neighborhood and over time. The omitted category is the bin with temperatures between 35C and 41C. Panel C uses LST 39C as the threshold, which represents roughly the 85-percentile of the distribution of LST. The omitted category is the bin with temperatures between 33C and 39C. Because we keep the bin width constant across panels, the omitted temperature range necessarily shifts as the upper bound of the top bin changes. The omitted category thus always represents roughly the same intermediate temperature range, ensuring that the contrasts remain comparable even though the numerical limits of the reference bin differ across panels.
Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
Standard errors clustered at neighborhood level in parenthesis. All regressions are weighted by the population aged 60+ in each neighborhood.
We include specific neighborhood linear trend in specification in column (1) to account for long-run changes between temperature and mortality. These trends are collinear in specifications (2) and (3). Controls include the share of population covered by family health clinics (CSF), the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month.
* significant at 10%; ** significant at 5%; *** significant at 1%.

Table A3: Robustness for different ranges for the bins: effects of Temperature on Mortality due to Chronic Causes

Shock	Mortality Rate (per 100,000 individuals aged 60+)		
	(1)	(2)	(3)
Panel A: range of 2C			
Number of Days with High Temperature	0.957	0.933	0.553
	(0.215)***	(0.213)***	(0.331)*
Panel B: range of 3C			
Number of Days with High Temperature	0.993	0.974	0.668
	(0.205)***	(0.204)***	(0.310)**
Panel C: range of 4C			
Number of Days with High Temperature	1.000	0.975	0.583
	(0.202)***	(0.198)***	(0.278)**
Year	✓		
Neighborhood x calendar Month	✓	✓	✓
Neighborhood x Year		✓	✓
Year x Month (time)			✓
Controls	✓	✓	✓
Observations	23,016	23,016	23,016
Mean dep. var.	181.0	181.0	181.0

Notes: Notes: Data is a monthly panel of neighborhoods. Each panel shows the effect of an alternative range for the bins. In Panel A, the range is 2C and the omitted category is the bin with temperatures between 31C and 33C. In panel B, the range is 3C and the omitted category is the bin with temperatures between 31C and 34C. Panel C shows the effects for a 4C-range and the bin omitted has temperatures between 31C and 35C. The high-temperature threhsolds remains fixed at above 43C, for all panels.
Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
Standard errors clustered at neighborhood level in parenthesis. All regressions are weighted by the population aged 60+ in each neighborhood.
We include specific neighborhood linear trend in specification in column (1) to account for long-run changes between temperature and mortality. These trends are collinear in specifications (2) and (3). Controls include the share of population covered by family health clinics (CSF), the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month.
* significant at 10%; ** significant at 5%; *** significant at 1%.

Table A4: Effects of Hot Days on Mortality due to Chronic Causes – Robustness for spatial correlation using Conley standard errors

Shock	Chronic Deaths per 100,000 individuals aged 60+		
	(1)	(2)	(3)
1. Number of Days with High Temperature	1.051	1.016	0.590
	(0.200)***	(0.198)***	(0.277)**
	[0.190]***	[0.187]***	[0.266]**
	{0.178}***	{0.198}***	{0.300}**
2. Number of 3-day Heat Waves	6.490	6.470	4.057
	(1.140)***	(1.119)***	(1.268)***
	[1.200]***	[1.179]***	[1.296]***
	{0.989}***	{0.995}***	{1.141}***
3. Number of 5-day Heat Waves	8.045	8.102	2.618
	(2.084)***	(2.074)***	(2.371)
	[1.923]***	[2.125]***	[2.389]
	{2.387}***	{2.641}***	{2.865}
4. Number of 7-day Heat Waves	12.418	12.079	3.823
	(3.620)***	(3.627)***	(4.013)
	[3.892]***	[4.082]***	[4.408]
	{4.142}***	{4.171}***	{3.654}
Year	✓		
Neighborhood x calendar Month	✓	✓	✓
Neighborhood x Year		✓	✓
Time			✓
Controls	✓	✓	✓
Observations	23,016	23,016	23,016
Mean dep. var.	181.0	181.0	181.0

Notes: Data is a monthly panel of neighborhoods. Each row shows the effect of an alternative measure of a temperature shock. In row 1, the shock variable is the number of days with LST greater than 43C for a neighborhood i in year t and month m. In rows 2 through 4, the shock is the number of events with more than 3, 5, or 7 consecutive days of temperatures exceeding 43 C.
Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
Standard errors clustered at neighborhood level in parentheses; Conley standard errors are presented in square and curly brackets. The cutoffs are 1.3km and 3.3km, which represent the median and the 90th percentile of the distribution of the minimum distance among the centroids of the neighborhoods. All regressions are weighted by the population aged 60+ in each neighborhood.
We include specific-neighborhood linear trend in column 1. Controls, included in all columns, include the share of population covered by family health clinics (CSF), the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month.
* significant at 10%; ** significant at 5%; *** significant at 1%.

B Specification Comparison and Interpretation

We estimate two versions of our main model. The benchmark model is:

$$h_{iym} = \alpha_{iy} + \delta_{im} + \lambda B_{iym} + \beta_i t + \gamma' Z_{iym} + \epsilon_{iym} \tag{B1}$$

where h_{iym} denotes the health outcome in neighborhood i, year y, and month m. B_{iym} captures the number of days in a neighborhood-year-month that the daily temperature falls into bin B (for simplicity, we only show one bin here). The model includes neighborhood-by-year α_{iy} and neighborhood-by-calendar month δ_{im} fixed effects. The variable t represents a linear time trend (year-month). We include specific neighborhood linear trends $\beta_i t$ in some specifications to account for long-run changes in temperature and mortality. The control variables Z_{iym}, included in all specifications, account for potentially confounding public policy interventions at the neighborhood level. This term considers the share of the population covered by Family Health Clinics, the neighborhood's average distance to the closest ER, the number of Pacifying Police Units, BRT stations, subway stations, and precipitation bins. A more saturated version includes time fixed effects:

$$h_{iym} = \alpha^*_{iy} + \delta^*_{im} + \tau^*_{ym} + \lambda^* B_{iym} + \gamma^{*\prime} Z_{iym} + \epsilon_{iym} \tag{B2}$$

To interpret the difference between the estimated coefficients λ and λ^*, note that we can approximately decompose the regressor B_{iym} as:

$$B_{iym} = \bar{B}_{ym} + \tilde{B}_{iym}$$

Where $\bar{B}_{ym}$ is the city-wide average heat exposure in year y and month m, and $\tilde{B}_{iym}$ is the deviation from this average in neighborhood i. Substituting into the benchmark model:

$$h_{iym} = \alpha_{iy} + \delta_{im} + \lambda \bar{B}_{ym} + \lambda \tilde{B}_{iym} + \beta_i t + \gamma' Z_{iym} + \epsilon_{iym}$$

Adding time fixed effects τ_{ym} absorbs $\bar{B}_{ym}$ and isolates identification from $\tilde{B}_{iym}$ alone. Therefore, the relationship between the two coefficients can be roughly expressed as $\lambda = \lambda^* + \lambda_{\text{city}}$, where λ_{city} captures the average effect on mortality of city-wide increases in heat exposure which is common to all neighborhoods in a given month, similar to the effects captured by τ^*_{ym} in specification (B2).[20]

[20] Formally, the benchmark coefficient λ is a variance-weighted combination of the structural effect of city-wide mean temperature and the effect of within-month spatial deviations. Thus, the difference $\lambda - \lambda^*$

Note that identification in both equations (B1) and (B2) relies on a rich set of fixed effects. The benchmark model includes neighborhood-by-year (α_{iy}) and neighborhood-by-calendar-month (δ_{im}) fixed effects. These absorb long-term, location-specific variation and seasonal neighborhood patterns, respectively. However, they do not capture variation in temperature that is common to all neighborhoods in a given year-month (e.g., a city-wide heat wave in January 2015). This component is only absorbed when we add time fixed effects (τ_{ym}) in equation (B2). Thus, the time fixed effects τ_{ym} control for city-wide shocks that are orthogonal to both α_{iy} and δ_{im}. This allows us to isolate the identifying variation in B_{iym} to within-month, across-neighborhood differences, yielding the decomposition in λ. In other words, the comparison between λ and λ^* therefore informs the degree to which heat-related mortality effects are spatially concentrated *versus* uniform across the city.

provides a heuristic measure of the relative importance of common versus localized heat exposure, rather than an algebraically exact decomposition.

C Dealing with Missing Values

Missing values can be pervasive in high-frequency LST data, primarily due to cloudy-sky conditions that obstruct satellite readings (Mildrexler et al., 2011; Shiff et al., 2021; Zhang et al., 2022; Yu et al., 2022; Li et al., 2024). Relying only on non-missing cases would therefore bias the analysis toward clear-sky days. In our case, there are 405,397 missing cases out of the total number of observations in our daily sample ($N = 731{,}302$). The number of missing cases therefore corresponds to about 55% of the total, in line with values found in the literature (Zhang et al., 2022). However, because the probability of missingness is strongly correlated with observable covariates that capture cloud coverage, we can address this issue by using imputation methods.

C.1 Empirical Approach

We follow commonly used methods to deal with missing observations in LST retrieved from MODIS products (Shiff et al., 2021; Zhang et al., 2022; Yu et al., 2022; Li et al., 2024; Han et al., 2024). Based on the idea of ground-truthing, where observations from ground-based weather stations are used to deal with the missing data problem, we rely on a model-based imputation that uses a regression framework incorporating spatial and temporal features to deal with the missing values in our dataset. The idea is to explore the time-series variation of several weather variables from the two ground-based weather stations with complete observations available in Rio (one from Alerta Rio System and the other from the National Institute of Meteorology), and to fit an econometric model to predict the LST values for each combination of neighborhood $\times$ day. The weather variables are: (i) from Alerta Rio System, the hourly observations of air temperature, humidity, a dummy for precipitation and wind speed from 6am to 5pm; and (ii) from the National Institute of Meteorology (INMET), daily observations of maximum air temperature, minimum air temperature, average air temperature, a precipitation dummy, relative humidity and Piche evaporation. In short, these data correspond to measures of temperature, wind speed, wind direction, precipitation, relative humidity, atmospheric pressure, and evaporation.

We perform the empirical exercise for each neighborhood in Rio, allowing for neighborhood-specific relationships among weather variables and LST. In practical terms, for each neighborhood, we regress the existing LST observations in that neighborhood on the rich set of weather covariates, and use the predicted values as our primary exposure variable. More specifically, the empirical specification we adopt for each neighborhood i is:

$$LST^i_{ymd} = \alpha^i + \gamma^i_m + \beta^i_m W_{ymd} + \epsilon_{ymd}$$

where LST^i is the time-series of LST for neighborhood i, γ_m corresponds to calendar-month fixed effects and accounts for general weather seasonality, and W are the weather variables from Alerta Rio System and INMET. The final chosen set of variables uses the set of weather covariates that minimizes the root mean squared error (RSME) of the prediction. In our benchmark method, we therefore predict the LST values for neighborhood i and use the fitted values as our LST measure at the neighborhood–day level.

We also consider alternative imputation strategies in additional analyses. First, we consider an elastic net model, which balances LASSO-style variable selection with ridge-style shrinkage. To ensure computational tractability, this approach is implemented at the daily level using summary statistics of the weather covariates (daily mean, maximum, and minimum). This allows the model to flexibly capture nonlinearities and interactions while remaining computationally efficient. Second, we employ a gap-filled model, which uses the benchmark specification to impute missing LST values only where gaps exist, while retaining observed values otherwise. This approach ensures fidelity to the raw satellite data while exploiting auxiliary information for interpolation.

C.2 Descriptives and Comparative Exercises

Figure C1 compares daily LST values obtained using the benchmark imputation method with those from observed cases. It shows that the benchmark imputation method reproduces the observed data quite closely, with no evidence of systematic upward or downward bias. Observations are clustered around the 45-degree line, indicating that imputed values are on average consistent with observed values. The main difference is a modest increase in dispersion, particularly at higher temperatures, which suggests classical prediction error rather than systematic distortion of the underlying distribution.

In Table C1 we perform comparative exercises with different imputation methods for estimating mortality effects. We observe that the benchmark specification yields the most conservative estimates. This is particularly clear in Panel A, where point estimates from alternative imputation methods are 40% to 71% greater than the benchmark coefficient. In Panel B, where time fixed effects are included, differences in estimates across methods are smaller. Importantly, the difference in effect size between specifications with and without time fixed effects remains relevant across imputation approaches. Point estimates from specifications with time fixed effects range from 36% to 48% relative to those from models without time fixed effects.

Figure C1: LST for the benchmark imputation method and for complete cases observations

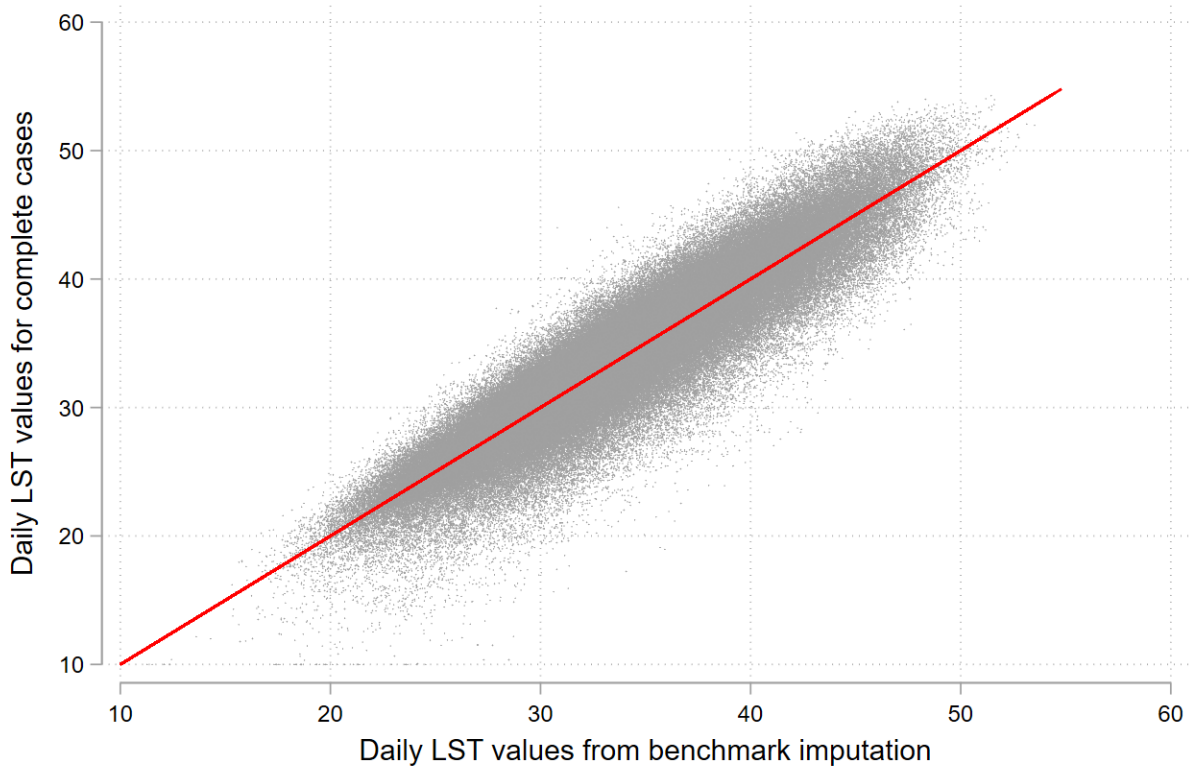

This figure shows the scatter plot for the benchmark imputation method and for complete cases observations. Each dot represents a neighborhood monthly LST average. The red line is a 45-degree line.

Table C1: Effects of LST temperature on Chronic outcomes (per 100,000 Individuals aged 60+) - Alternative Imputations

	(1) Benchmark	(2) Elastic Net	(3) Gap-Filled
Panel A: Without time fixed effects			
Number of Days w/ High Temp.	1.016	1.712	1.404
	(0.198)***	(0.259)***	(0.247)***
Panel B: With time fixed effects			
Number of Days w/ High Temp.	0.590	0.604	0.672
	(0.277)**	(0.336)*	(0.320)**
Observations	23,016	23,016	23,016
Neighborhood x Year	✓	✓	✓
Neighborhood x calendar Month	✓	✓	✓
Controls	Full	Full	Full
Mean Dep. Var.	181.0	181.0	181.0

Notes: The estimates are for bin 43+, which calculates the number of days with LST temperature greater than 43C for a neighborhood i in year y and month m. The "elastic net" specification balances LASSO-style variable selection with ridge-style shrinkage. The "gap-filled" specification uses the benchmark imputation to fill missing values. Standard errors are clustered at the neighborhood level and the regression equation is weighted by the population aged 60+ within each neighborhood. We include specific neighborhood linear trend in specification in column (1) to account for long-run changes between temperature and mortality. These trends are collinear in specifications (2) and (3). Controls include health services (Family Health Clinics coverage and average distance to the closest emergency units), and city infrastructure (e.g., Pacifying Police Units, BRT, subway stations). * significant at 10%; ** significant at 5%; *** significant at 1%.

C.3 Spatial Interpolation

An alternative to local imputation is to use spatially interpolated gridded temperature products. These datasets, such as the 1-km global air temperature product from Zhang et al. (2022), have the advantage of broad coverage but rely heavily on spatial smoothing techniques. By construction, they interpolate temperature readings across space, often compressing local highs and dampening cross-neighborhood variation. In principle, this feature may limit its suitability for studies seeking to uncover fine-grained heterogeneity in exposure and outcomes. To evaluate these trade-offs, we conduct a series of empirical comparisons between our benchmark imputation method and the interpolated temperature product from Zhang et al. (2022).

Figure C2 shows the kernel density of daily LST values under different data sources. While the imputation methods yield broadly similar distributions, the interpolated values from Zhang et al. (2022) display higher central concentration and thinner tails, consistent with spatial smoothing. This loss of dispersion becomes more evident in Figure C3, where we examine the coefficient of variation in temperature across neighborhoods.[21] In the benchmark data, cross-neighborhood variation declines but remains relevant even in hotter bins. In contrast, variation in the spatially interpolated data shrinks sharply at higher temperature levels, precisely where city-wide heat stress becomes most relevant. While the coefficient of variation in the highest bin in the benchmark data is 0.015, it reduces to 0.002 in the spatially interpolated data.[22]

This mechanical compression of spatial variation can have direct implications for the empirical estimation of mortality effects. Table C2 compares point estimates using the two LST measures. When we estimate the effect of hot days on mortality without time fixed effects (Panel A), both methods yield similar coefficients, providing a reassuring check. However, once we add time fixed effects to absorb common city-wide shocks (Panel B), estimates based on the spatially interpolated data shrink and become statistically insignificant. This divergence is consistent with the interpretation that, once common temporal variation is absorbed, residual identification hinges on spatial heterogeneity, which the spatially interpolated dataset largely eliminates.

Taken together, these results highlight a limitation of spatial interpolation. While useful

[21] We use the same rationale for the construction of the thresholds for the bins for data from Zhang et al. (2022) as in our benchmark specification. We define high-temperature days if the temperature is above the 95th percentile of the empirical distribution, and then, we define 6C bins of temperature starting from the highest bin. In the case of Zhang et al. (2022)'s data, the thresholds are above 41C, between 41C and 35C, between 35C and 29C, 29C and 23C, and below 23C.

[22] We also calculate the average cross-sectional standard deviation in a period of time. The average standard deviation in the cross-section for the benchmark specification is 4.05, while for data from Zhang et al. (2022) we find almost ten time less variation, with a standard deviation of 0.49.

for broader climatological applications, such methods may understate localized climate impacts in urban settings. By contrast, our benchmark imputation, which relies on local ground station data and preserves neighborhood-level temperature differences, is better suited to studying heterogeneous exposure and urban vulnerability under heat stress.

Figure C2: Kernel Density distribution of daily LST values for different scenarios

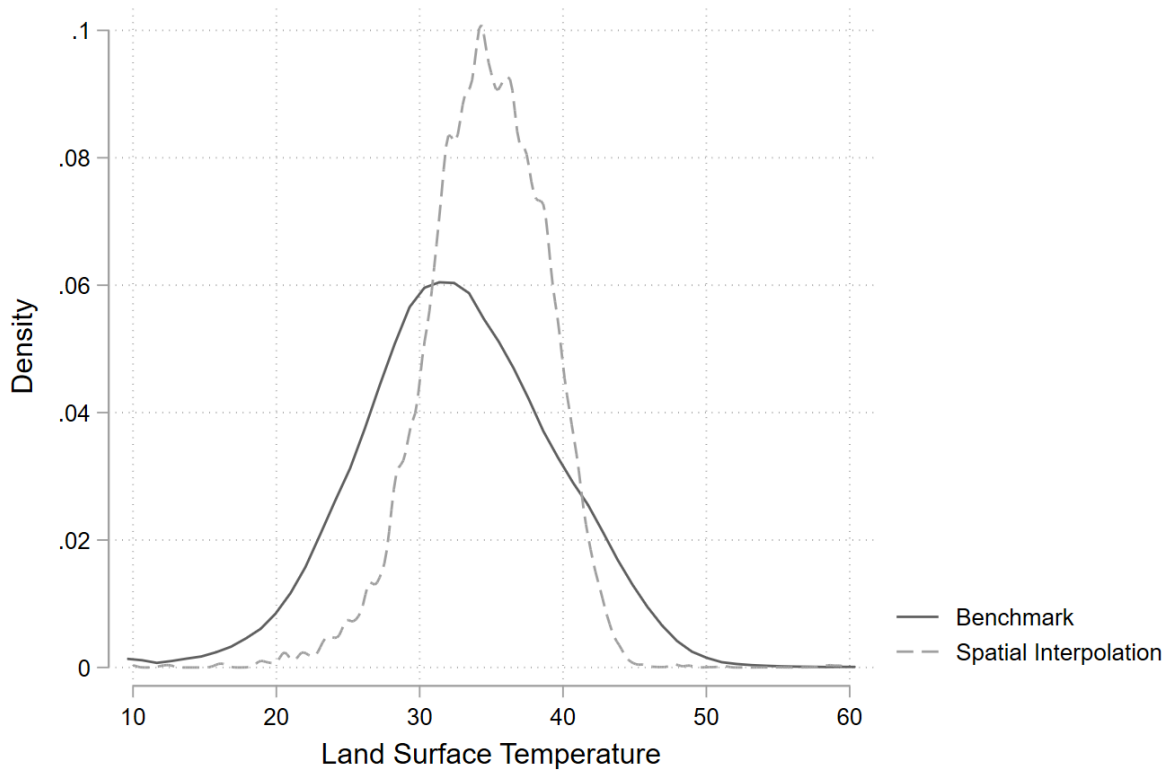

Figure C2 shows the kernel density distribution of daily land surface temperature (LST) values under different data imputation scenarios. The solid line corresponds to the study's reference scenario (Benchmark), the dashed line follows the procedure proposed by Zhang et al. (2022), and the dotted line relies exclusively on complete case observations.

Figure C3: Variation of LST temperature in Zhang et al. (2022) data

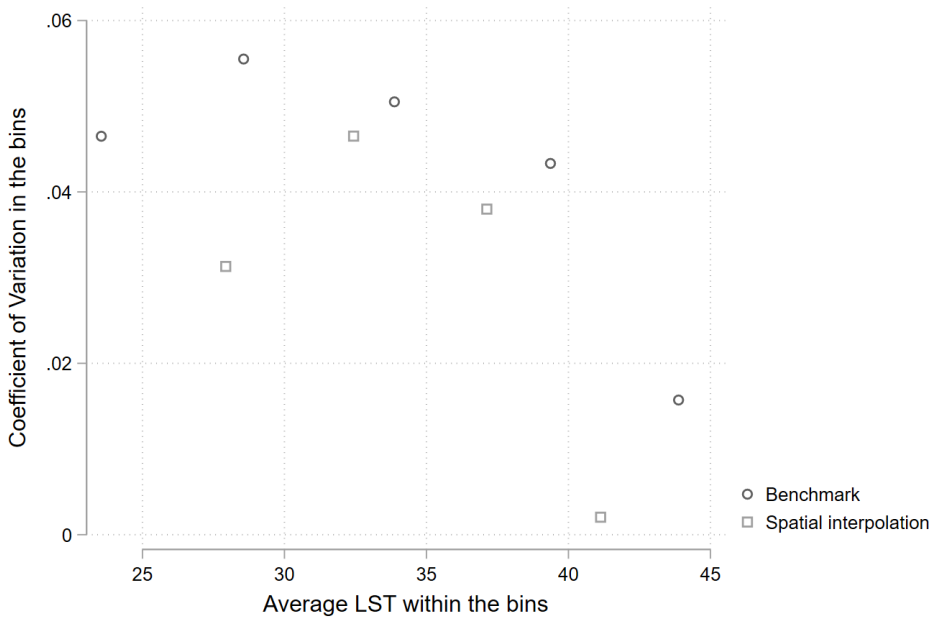

For this figure, we classify neighborhoods into four temperature bins based on the daily product of Zhang et al. (2022): below 29C, between 29C and 35C, between 35C and 41C, and above 41C. While the original thresholds also include a category below 23C, when averaging daily values to compute the coefficient of variation, no bin means fall into this lower range. The figure shows the relationship between the coefficient of variation and the average LST across bins.

Table C2: Effects of LST temperature on Chronic outcomes (per 100,000 Individuals aged 60+) - Alternative Imputations

	(1) Benchmark	(2) Spatial Interpolation
Panel A: Without time fixed effects		
Number of Days w/ High Temp.	1.016	1.193
	(0.198)***	(0.183)***
Panel B: With time fixed effects		
Number of Days w/ High Temp.	0.590	-0.161
	(0.277)**	(0.501)
Observations	23,016	23,016
Neighborhood x Year	✓	✓
Neighborhood x calendar Month	✓	✓
Controls	Full	Full
Mean Dep. Var.	181.0	181.0

Notes: Chronic deaths are defined as diseases of the circulatory system (Chapter IX), respiratory system (Chapter X) and endocrine, nutritional and metabolic diseases (Chapter IV).
The estimates are for the benchmark are bin 43C+, which calculates the number of days with LST temperature greater than 43C for a neighborhood i in year y and month m. The 'benchmark' imputation method uses the full set of weather variables in the prediction part described in Appendix C; in column (2) uses the LST derived from the air temperature values in the Zhang et al. (2022). The estimates are for bin 41C+, which represents the 95th percentile of the daily distribution LST values for Zhang et al. (2022) paper. Standard errors are clustered at the neighborhood level. All regressions are weighted by the population aged 60+ in each neighborhood.
Linear trend refers to a specific neighborhood linear trend. Controls include the share of population covered by family health clinics (CF), the average distance of the neighborhood to the closest ER room, and the number of Pacifying Police Units, bus rapid transportation stations, and subway stations in each neighborhood-month. * significant at 10%; ** significant at 5%; *** significant at 1%.

D Neighborhoods' Aggregation

In order to make our datasets consistent, we aggregated the 162 neighborhoods into 144 *bairros*. To do so, we consider recently emancipated neighborhoods as still belonging to their neighborhood of origin. We do the following corrections:

- Gericino is considered as Bangu.
- Vila Kennedy as Bangu.
- Vasco da Gama as Sao Cristovao.
- Parque Columbia as Pavuna.
- Lapa as Centro.
- Freguesia, Ribeira, Zumbia, Cacuia, Pitangueiras, Cocota, Bancarios, Jardim Guanabara, Jardim Carioca, Taua, Monero, Portuguesa and Galeao as Ilha.

The main regressions for individuals aged 60 years old or more consider only neighborhoods with population for this age group greater or equal than 500 people. Seven *bairros* are excluded: Camorim, Campo dos Afonsos, Cidade Universitaria, Grumari, Joa, Paqueta and Saude. We chose this criteria because since these neighborhoods have very low population, their mortality rates were inflated.