The Price of Cool: How Urban Heat Shapes Neighborhood Value

Hannah Kamen

January 2026

Abstract: Urban heat is a growing concern in U.S. cities, however it remains unclear how households value the benefits of neighborhood cooling (e.g., shade trees, cool pavement, etc.) at increasing distances from their home. Measuring these benefits is challenging because temperature is correlated with many other neighborhood differences, including resident income and development patterns. I use local wind patterns to identify temperature variation that is unrelated to other neighborhood attributes, allowing me to estimate how much households value cooling at different distances. Using home sales data from Maricopa County, AZ, I find that households are willing to pay $676 annually for a 1°C decrease in summer air temperatures in their immediate surroundings, and $245 annually for equivalent cooling in their broader neighborhoods. I apply these estimates to Phoenix's Cool Corridors Program and find that analyses ignoring these broader effects understate the program's total value by 48%. These findings provide novel evidence that neighborhood cooling is a local public good: private adoption underprovides cooling, and conventional evaluations can understate the gains from targeted public investments.

JEL Codes: C36, Q51, Q58

Keywords: urban heat island, hedonic pricing, instrumental variables, spatial spillovers, local public goods, climate adaptation, cost–benefit analysis, urban heat policy

Acknowledgments: This research was supported by graduate scholarship funding from the Colorado School of Mines, including the Carol Dahl Fellowship and the Pauli Scholarship. I am grateful to Jared Carbone, Justin Winikoff, and Ben Gilbert for invaluable guidance and feedback throughout the development of this paper. I also thank seminar participants at the Colorado School of Mines and at CU Boulder for helpful comments. The author has had no relevant financial relationships that could reasonably be viewed as related to the subject of this paper.

1 Introduction

Extreme heat exposure has become a growing concern in U.S. cities. Higher summer temperatures raise household costs through increased electricity spending and adverse health outcomes (Sailor and Pavlova, 2003; Zivin and Shrader, 2016), prompting cities to expand investments in cooling infrastructure such as parks, trees, cool pavement, and shade structures (EPA, 2025). Yet policymakers and researchers still have limited evidence on how much households value cooler neighborhoods, despite survey evidence that they are aware of property-level variation in urban heat (Jenerette et al. 2016). In particular, it remains unclear how far beyond the home households value the benefits of local cooling improvements. Because heat shapes daily life both inside the home and in nearby public and private spaces, cooling projects may generate benefits that extend to surrounding residents through local spillovers (Basu et al., 2024; Han et al., 2024; Zivin and Neidell, 2014). This spatial reach matters for policy analysis: when the benefit area extends beyond the homes immediately adjacent to an intervention, valuations that assume too narrow a footprint will understate the true benefits of public cooling investments.

A central challenge in estimating household willingness to pay for urban cooling is that temperature is endogenous to other neighborhood characteristics. Urban heat changes alongside many features that are difficult to observe—such as vegetation, parks, development intensity, and traffic—that evolve over fine spatial scales and over time (EPA, 2024). These factors can affect housing values both through temperature and through independent amenity channels; for example, parks provide both cooling and recreational value. At the same time, property values and microclimates may be jointly determined: higher-valued neighborhoods may attract investment in trees, irrigation, and cool pavements, while high land values can discourage land-intensive cooling amenities such as public green space. As a result, naïve regressions of home prices on neighborhood temperature are likely to conflate the value of cooling with the effects of other co-evolving neighborhood changes. These endogeneity challenges are especially salient when the goal is to value *local* heat exposure *within* a city, rather than regional climate differences *across* cities. Most existing climate amenity estimates rely on cross-metropolitan comparisons or time-invariant fixed effects that are well suited to absorbing persistent regional heterogeneity but cannot address time-varying confounders or reverse causality at the neighborhood scale (Diamond et al., 2016; Albouy et al., 2016; Sinha et al., 2018; Wrenn et al., 2024).

In this paper, I investigate how much households value urban cooling at different distances from the home and use these estimates to evaluate the benefits of an urban cooling program. I proceed in three steps. First, I estimate a hedonic property value model using home sales data from Maricopa County, Arizona paired with 1 km resolution urban-resolving air temperature data from NCAR's High-resolution Urban Meteorology for Impacts Dataset (HUMID)– one of the first datasets capable of capturing within-city temperature variation at a scale that reflects how urban heat varies across neighborhoods. To address the endogeneity of neighborhood temperature, I develop a novel instrumental variable based on fine-scale urban air advection– the horizontal transport of air by surface winds– that separates local temperature changes caused by upwind hot and cool zones from those arising within the neighborhood itself. Unlike wind-direction instruments used in the air pollution literature, which typically exploit city- or county-level variation in wind patterns (Deryugina et al., 2019; Freeman et al., 2019; Schlenker & Walker, 2016), this approach leverages parcel-level differences in orientation relative to surrounding temperature gradients and uses modeled hourly air parcel trajectories from NOAA's HYSPLIT atmospheric transport model to trace the precise paths of surface airflow over each property. The resulting instrument varies at sub-block spatial scales and is constructed entirely from *relative* temperature anomalies and angular alignment, making it orthogonal to the local temperature level by design. Because the instrument does not rely on granular spatial fixed effects for identification, I am able to estimate a spatial lag extension of the hedonic model that separately identifies household willingness to pay for cooling at multiple distance bands around the home. Second, I use a stacked difference-in-differences design to estimate how much Phoenix's Cool Corridors tree-planting program reduces local temperatures, comparing treated corridors to program-selected corridors that have not yet been implemented. Third, I combine the hedonic MWTP estimates with the DiD cooling estimates to calculate the total value of cooling benefits generated by the program– and show how that value changes depending on the spatial scale at which benefits are measured.

The results from this framework yield three main findings. Households are willing to pay $676 annually for a 1°C reduction in summer air temperatures within a few blocks of their home (500 meter radius), and $245 annually for equivalent cooling at the broader neighborhood scale (1,000 meter radius). Applying these estimates to Phoenix's Cool Corridors Program, I show that conventional single-scale,

tract-level approaches understate benefits by 48%. Underestimation occurs because such approaches do not account for the total spatial scale over which households value cooling and how that value changes with distance: cooling delivers its largest gains very near the home while also generating meaningful benefits at the neighborhood scale. Incorporating both effects is particularly important in dense cities, where tract boundaries divide otherwise continuous neighborhoods. This strategy may also be useful in other settings that require fine-scale identification of environmental exposure, such as hyper-local air quality valuation.

These findings are especially timely given recent surges in funding for public cooling investments (City of Phoenix, 2024; NOAA, 2024; New Jersey Department of Environmental Protection, 2025; EPA, 2025). They provide the first empirical evidence that policymakers should not take the geographic reach of heat-related benefits as given: the spatial scale at which cooling is measured and valued can materially change the implied benefits of marginal temperature reductions. The results also indicate that cooling infrastructure generates spillovers, which may limit the ability of market forces alone to deliver efficient outcomes, for example– if households value heat conditions beyond their property, private choices—such as landscaping, cool roofs, or commercial development do not fully account for impacts on nearby residents.

These findings are especially timely given recent surges in funding for public cooling investments (City of Phoenix, 2024; NOAA, 2024; New Jersey Department of Environmental Protection, 2025; EPA, 2025). They provide the first empirical evidence that the spatial scale at which cooling is measured and valued can materially change the implied benefits of marginal temperature reductions. Previous hedonic estimates of household willingness to pay for urban heat reduction have been produced exclusively at the census tract level (**klaiber2017**; **borsky2024**), which means that benefit–cost analyses drawing on this literature have not had access to multi-scale estimates needed to account for how cooling value varies with distance from the home. The results also speak to the role of public provision in this setting: because households value heat conditions beyond their own property, private choices—such as landscaping, cool roofs, or commercial development—do not fully internalize cooling benefits that accrue to nearby residents. These spillovers provide an economic rationale for public intervention in urban heat reduction, and the multi-scale MWTP estimates developed in this paper can help policymakers more

precisely quantify the spatial reach of such benefits when evaluating cooling programs. More broadly, the findings suggest that models which average environmental attributes over areas larger than the scale at which they vary can inadvertently treat benefits as uniform across space, raising equity concerns. This issue is particularly relevant for city governments, which are often responsible for allocating cooling investments across relatively small geographic areas.

The remainder of the paper proceeds as follows. Section 2 situates the paper in the related literature. Section 3 describes urban heat in the study area and Phoenix's Cool Corridors Program. Section 4 summarizes the data. Section 5 develops the intuition for the instrumental variable I use and explains how it is constructed. Section 6 presents the econometric framework for identifying household MWTP and the differences-in-differences design used to estimate Cool Corridors' cooling impacts. Section 7 reports the hedonic and program-evaluation results, including the baseline IV estimates. Section 8 concludes.

2 Literature

Few papers estimate households' MWTP for hyper-local reductions in urban heat, despite growing city investments in neighborhood cooling infrastructure (City of Los Angeles, 2023; EPA, 2025; City of Phoenix, 2026). Endogeneity remains a central challenge in this setting. Among the few exceptions, Klaiber, Abbott, and Smith (2017) and Borsky et al. (2024) estimate MWTP for local heat reduction, while Han et al. (2024) focus on electricity cost savings from local temperature changes rather than broader amenity value. Klaiber, Abbott, and Smith (2017) study Phoenix, Arizona and estimate MWTP at the tract level using a Hausman-Taylor model, instrumenting for aggregate tract temperatures using between-tract variation in elevation. They report a MWTP of $72 per month ($864 annually) per °F reduction in tract temperature in 2010 dollars, which corresponds to $2,292 annually per °C reduction in 2025 dollars. Because identification in this design relies on comparisons across low and high elevation tracts, the resulting estimates may also reflect other unobserved amenities that are correlated with elevation, such as air quality, views, or flood risk. Focusing on climate gentrification effects from cooling programs, Borsky et al. (2024) estimate a spatial sorting model with tract fixed effects and find that households in New York City are willing to pay $99 annually to avoid one additional

day above 95 °F. While tract fixed effects address time invariant differences across neighborhoods, this approach does not resolve endogeneity from time varying, unobserved changes in urban form that can affect both microclimates and neighborhood amenity values.

Previous work also does not fully address the possibility that households value summer heat reduction differently across space. There are clear reasons to expect this. Heat in the immediate vicinity of the home affects indoor cooling costs (Han et al. 2024), the use of outdoor residential space (Zivin and Neidell, 2014), and exposure during short trips nearby (Basu et al., 2024). Heat in the broader surrounding area may matter for commuting, errands, recreation, school, and social activities (Fan et al., 2023; Batur et al., 2024). Common econometric choices can make it difficult to study these differences. Approaches that aggregate temperature to the tract level cannot distinguish microclimate effects within tracts (Klaiber, Abbott, and Smith 2017), even though recent work shows that urban temperature varies meaningfully at sub-tract spatial scales (Chakraborty et al., 2019; Hsu et al., 2021; Wang et al., 2018; Middel et al., 2014, 2019; Myint et al., 2015). At the same time, highly granular spatial fixed effects—while useful for addressing time-invariant unobservables (Kuminoff, Parmeter, and Pope, 2010)—remove the cross-area variation needed to detect impacts that extend beyond the fixed-effect unit, so researchers face a tension between using highly granular location controls to limit omitted-variable bias and estimating amenity effects that operate over larger geographic areas.

To address endogeneity concerns, and the limits of relying on either tract averages or very fine spatial fixed effects, I develop an instrument for neighborhood temperature based on fine-scale differences in household orientation relative to nearby "hot" and "cool" zones that generate plausibly exogenous temperature variation through urban air advection. This strategy offers two advantages in the urban heat setting. First, relative to approaches that rely on granular spatial fixed effects to address time-invariant unobservables, the instrumental variable (IV) approach addresses both time-invariant and time-varying confounders, which is important because urban heat is closely linked to changes in cities over time (Middel et al., 2014; Myint et al., 2015). Second, because this instrument does not require aggregation to the tract level or identification solely from within-unit variation under highly granular fixed effects, it allows me to estimate how households value microclimate improvements at multiple spatial scales around the home.

3 Empirical Overview

I estimate how much households value cooler neighborhood temperatures using Maricopa County, Arizona home sales data from 2009–2018, which covers the core of the Phoenix–Mesa–Chandler metro area. The study area represents 1.7 million households, with a total property value exceeding $1 trillion (Maricopa County Assessor's Office, 2026). Heat exposure varies considerably across the county (Fig 1, panel (a)) and is a primary environmental concern for residents, with July average temperatures rising past past 100 °F (38 °C) in recent years (US Dept of Commerce, 2025). In 2021 the City of Phoenix collected climate survey responses from residents, and found that requests related to improved cooling infrastructure comprised the top category, with 75% of residents reporting "extreme concern" related to heat (City of Phoenix, 2021). Furthermore, research shows that residents are aware of property-level variation in local climate that result from changes in development and vegetative land cover (i.e., the urban heat island effect). Jenerette et al. (2016) combine high resolution land cover and temperature data with survey responses from Phoenix residents and show that residents noticed even micro-scale variation in urban heat. Evidence from energy expenditure data and community responses suggests that heat saliency stems from both electricity expenses required for cooling,[1] as well as the impact that heat has on activities beyond the home, including comfort and safety during work and school commutes, as well as recreation quality in parks (Heat Action Planning Guide for the Neighborhoods of Greater Phoenix, 2017). In response to these risks, the City Council of Phoenix formed its Office of Heat Response and Mitigation in 2021, and launched Phoenix's Cool Corridors Program in 2022. Since then the city approved a $60M Tree and Shade Master Plan in 2024, targeted at further cooling infrastructure investment via shade structures, cool pavement and tree planting.

[1] Appendix Figure A1 provides descriptive evidence that hotter microclimates are associated with higher household cooling costs. I merge tract-level annual household electricity expenditures from the Department of Energy's Low-Income Energy Affordability Dataset (LEAD, 2022) with minimum July nighttime temperatures from the NCAR high-resolution urban climate dataset, and compute expenditures per square foot using the median home size in the Maricopa County parcel data (1,623 sq ft). In 2020, households in the hottest nighttime-temperature decile spent more than 12% more on electricity per square foot annually than those in the coolest decile (about $0.15/sq ft, or $227 per year). Because these data are observed at the tract level, and aggregate electricity costs across all uses and months, they likely understate the magnitude of household-level savings from cooler microclimates, but they motivate a channel through which localized heat exposure can generate meaningful welfare losses and thus household willingness to pay for neighborhood cooling; the hedonic estimates in this paper reflect both energy-cost savings and broader amenity benefits such as comfort and outdoor usability.

In this paper, I use Phoenix's Cool Corridors Program to operationalize my estimates of household MWTP for urban cooling. The program is a heat-mitigation initiative led by the Street Transportation Department under the city's Climate Action Plan Framework and Tree and Shade Master Plan (City of Phoenix, 2025). Since its launch in 2022, the city has allocated $1.4 million in annual funding from the city's Capital Improvement Program (CIP) through 2026. The funding is primarily allocated to tree planting along quarter to half-mile segments of prioritized pedestrian and cyclist routes (i.e.,"Cool Corridors"), with a target of roughly 200 drought-tolerant, low-water-use trees per mile in heavily traveled areas (City of Phoenix, 2025). Corridor locations are selected using a multi-criteria ranking system that emphasizes neighborhood heat vulnerability, transit dependence, pedestrian activity, and existing heat conditions (City of Phoenix, 2025). Funding has been largely productive: as of September 2024, the City reported 2,945 trees planted since the program began.[2] Figure 1 maps the 94 corridor locations (63 completed and 31 anticipated) within Phoenix city boundaries as of 2025.

A key challenge to constructing policy benefit measures in this context is the plausible correlation between urban heat and other, often unobservable, neighborhood characteristics that households also value. Figure 2 uses available measures of tract income and land cover, as well the main metric for air temperature used in this paper to motivate the use of an econometric strategy that addresses likely sources of omitted variable bias, and is not limited to identifying effects at the scale of granular spatial fixed effects. In particular, development changes over time and cross-sectional patterns of median tract income appear correlated with the temporal and spatial changes in urban air temperatures shown in Figure 2 (a). Although these particular measures are observable, Figure 2 suggests that urban heat co-evolves with associated unobservables that may be realized over scales smaller than the tract level, and which may change over time. Furthermore, the observed fine-scale spatial heterogeneity of urban heat exposure shown in Figure 2 (a) motivates the possibility that households may respond to urban heat at multiple spatial scales- extending beyond private cooling benefits (i.e. cooling, comfort on private property) on their own parcel to include adjacent public spaces and neighborhood amenities where they

[2]Because my data end in 2025, my difference-in-differences estimates capture cooling over the first three summers after planting, before trees reach maturity. The urban climate literature finds that even individual trees can measurably reduce near-surface air temperatures through shade and evapotranspiration, and that cooling effectiveness increases with canopy density and leaf area as trees grow. Thus, I interpret my cooling and benefit estimates as conservative relative to longer-run impacts once canopy is fully developed (Meili et al., 2021; Rahman et al., 2020; McPherson, 2014; Cai et al., 2022).

may work or recreate.

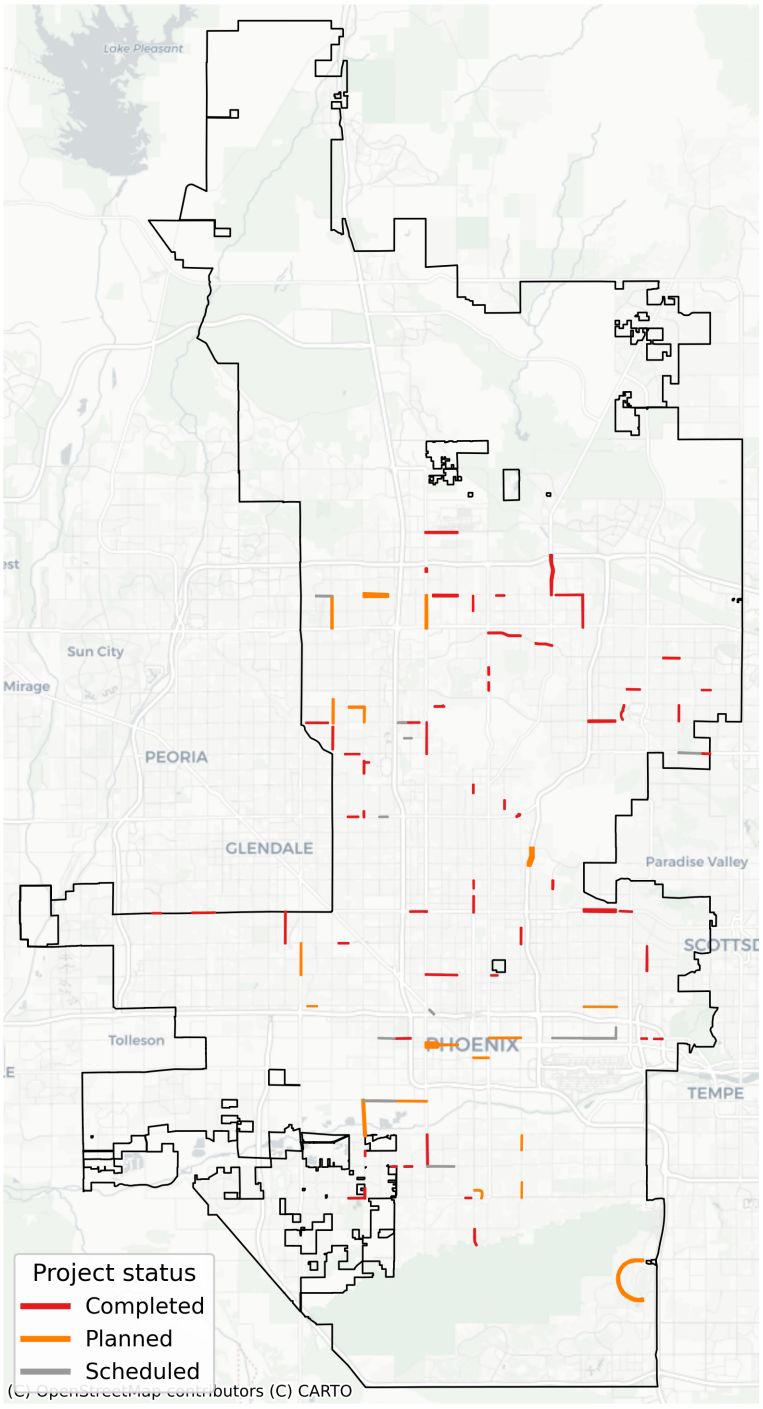
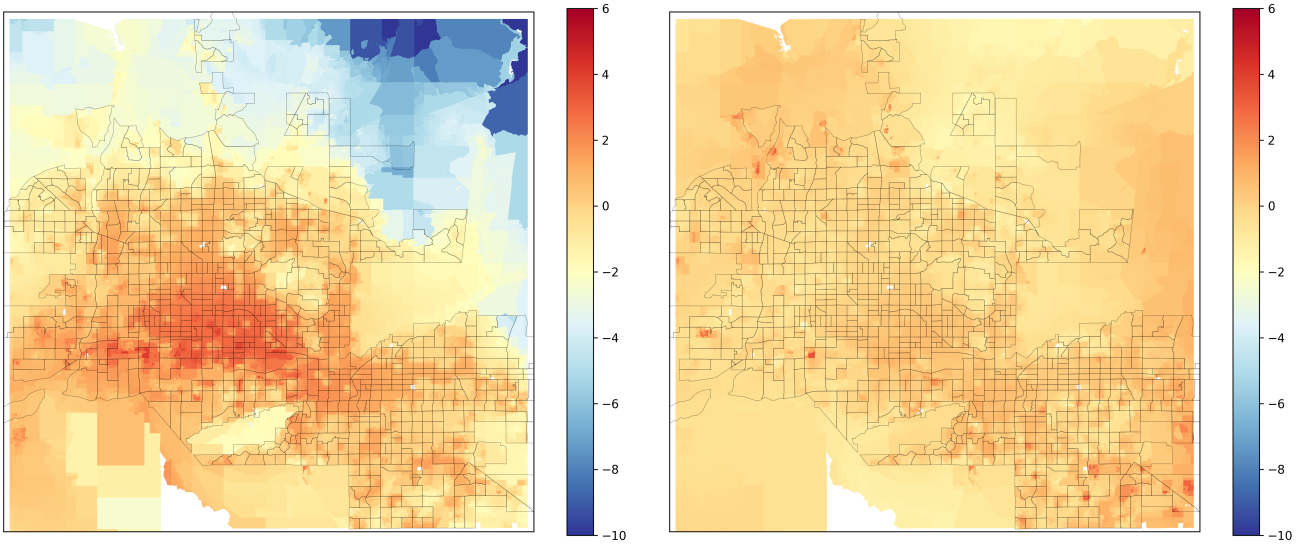


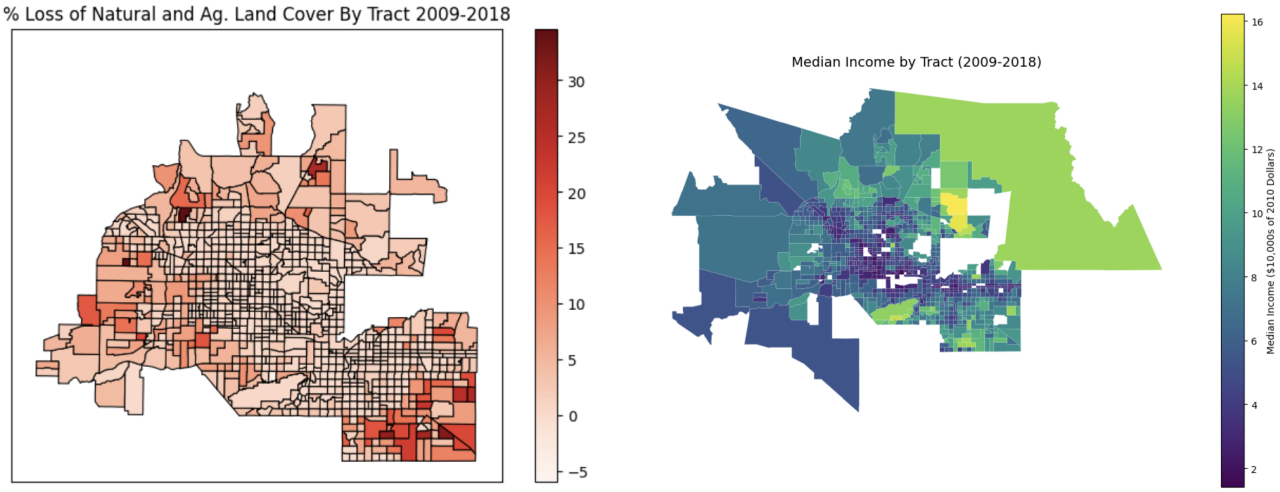
Figure 1: Phoenix Cool Corridors program: completed, in-progress, and planned corridor segments. The black outline represents the extent of Phoenix's city boundary.



(i) Spatial variation

(ii) Temporal variation

(a)



(i) Development change over time

(ii) Median tract income

(b)

Figure 2: Urban heat variability and correlates in Maricopa County. Panel **(a)** displays the high degree of variability in urban heat over **(i)** space and **(ii)** time. Panel **(b)** motivates the likelihood that urban heat is endogenous to other amenities by illustrating its apparent correlation with **(i)** development changes over time and **(ii)** median tract income. The urban heat measure used in this figure is average July nighttime minimum air temperature (1000m), accessed from the National Center for Atmospheric Research's (NCAR) High-resolution Urban Meteorology for Impacts Dataset, (HUMID).

4 Data

This section contains a summary of all data sources used in this paper. Further detail on these sources is available in Appendix B.

4.1 Hedonic Analysis Data

I assemble a parcel-level dataset of non-foreclosure, arms-length single-family, townhome, and apartment home sales in Maricopa County, Arizona from 2009-2018 using Maricopa County Tax Assessor records, which provide transaction prices and housing characteristics used to construct the sample used for the hedonic regressions. Homes that sold below and above the 5th and 95th percentiles of home price for the total sample were excluded from the sample, along with homes greater than the 90th percentile of distance away from Phoenix's city center. Furthermore, homes that underwent renovation projects (i.e change in square footage, or addition of pool, garage, patio, bathroom) were excluded from the sample. I annualize home prices using a 5 percent discount rate (in perpetuity), following (Himmelberg, 2005). In addition to sales price, this dataset contains a rich set of covariates specific to each property including structure type (single family, apartment, condo, or townhome), square footage, number of bathroom fixtures, date of construction, and whether or not the property has a pool and garage. I tie sales data to GIS boundaries for each residential parcel using the Parcel Shape File for all of Maricopa County available from Maricopa County GIS Open Data portal.

To measure urban heat exposure, I use NCAR's 1000m resolution Urban Meteorology for Impacts Dataset (HUMID), using average July nighttime minimum (2m) air temperatures. I focus on July minimum nighttime temperatures because July is typically the hottest month of the year in Maricopa county, and because nighttime min temps are a conventional metric used to assess the severity of the urban heat island effect (Klaiber, Abbott, and Smith 2017; Oke 1982). In order to construct the airflow trajectories used to instrument for July nighttime minimum temperature, I use near-surface hourly atmospheric condition data (wind speed, direction, air pressure, non-urban resolving temperature data) from NOAA's High-Resolution Rapid Refresh (HRRR) dataset. This is the input dataset for NOAA's air parcel trajectory modelling tool (HYSPLIT) discussed in more detail in Section 4.

In order to run robustness checks for my main specification I combine several datasets. To characterize development intensity changes by tract over time, I use USGS National Land Cover Database (NLCD) annual 30 m land cover classifications for years 2009-2018, reclassified to capture development intensity. I reclassify the existing NLCD categories to reflect two categories, intense and non-intense

land development. Intense development land area per tract reflects the annual percentage of each tract that corresponds to NLCD development categories medium and high, and non-intense development reflects the annual percentage of each tract that corresponds to open space and low intensity development. The distance that each property is located relative to urban greenspace is classified using to nearest park/open space is inferred using 2010 and 2015 1m resolution land cover classification datasets from the Central Arizona-Phoenix Long Term Ecological Research database. "Greenspace" is defined as any contiguous 30mx30m or larger area of turf grass and/or tree and shrub land use classifications, to calculate parcel distance to greenspace I generate a shape file from the land cover areas within Maricopa county that meet this requirement, and calculate each property centroid's distance to the nearest greenspace polygon in both 2010 and 2015. The distance that each property is located relative to a major roadway is calculated using the 2019 TIGER/line shapefile of Arizona primary and secondary state roads from the US census bureau. I then calculate the distance between each property's centroid and the closest major road.

Annual tract level median income from 2009 to 2018 is taken from ACS 5-year times-series estimates of median income in previous year by tract, converted to real 2025 dollars, accessed via from IPUMS NHGIS. Median income for each year is calculated from the ACS 5-year estimate with end year prior to a given year. My measure for annual air quality uses the Daily Census Tract Level PM2.5 Concentration, 2011-2015 and 2016-2020 datasets available from the U.S. Department of Health and Human Services. I use annual mean June, July and August PM2.5 concentrations for each tract in Maricopa county. For sample years 2009 and 2010, which occur prior to the creation of these datasets, I assign the PM2.5 from the closest year, 2011. Distance to Phoenix's city center is defined as the distance between each parcel's centroid and the coordinates -112.074, 33.45, which represent the center of Phoenix's central business district.

4.2 Cool Corridors Data

I obtain GIS data on the City of Phoenix Cool Corridors program from the City of Phoenix Open Data portal. The 2025 Cool Corridors layer represents street segments that are completed or planned

as part of the program. For completed segments, the data includes completion date and number of trees planted. I limit the dataset to corridor lengths greater than 400m, and for completed corridors, corridors with least 62 trees per km density. This corresponds to 50% of the corridor's initial stated goal of 200 trees per corridor mile. To avoid confounding temperature changes in the control "planned and scheduled" corridors, I remove all control corridors within 100m of an already completed corridor. To avoid confounding from Phoenix's partially simultaneous cool pavement program, I remove corridors within 100m of a completed cool pavement projects. Completed Cool Pavement projects GIS outlines are also obtained from the City of Phoenix Open Data portal. Furthermore, in order to account for small areas of cooling surrounding small urban areas of dense tree cover (Soltanifard and Amani-Beni 2017), I extend Cool Corridor boundaries to extend to a 100m buffer surrounding each corridor.

Because the last year of air temperature observation in the HUMID temperature dataset is 2018, which is prior to the first year of cool corridor implementation 2022, I must estimate surface air temperatures in Cool Corridor Neighborhoods using a combination of annual (2019-2025) 30m Landsat July land surface temperature imagery and an estimate of the relationship between surface and air temperature in Phoenix from Song et al. (2017). The authors estimate the average ratio between air and surface temperature changes to be 0.63:1 over developed areas in Phoenix, Arizona using meteorological and land surface data from 2013. Land surface temperature data is obtained from Landsat 8 and 9 land surface temperature image collections for Maricopa county on days with $< 10\%$ cloud cover in July of each year between 2019 and 2025. July mean surface temperature for each year for each corridor reflect the mean of land surface temperature pixels which overlap corridor boundaries (including a 100m buffer). In order to calculate the contribution of corridor cooling to 500m, 1000m ring buffers surrounding each residential property, and tract zones I join corridor boundaries with each property's 500m radius, 1000m radius and tract boundaries and calculate percent area overlap. Cooling from Corridor overlap in each property's 500m radius, 1000m radius and tract boundaries is the calculated using percent corridor overlap, and degree of cooling inferred from the results of the differences in differences analysis.

5 Airflow Instrumentation Strategy

I instrument for local air temperature surrounding each property using the degree to which hot or cool air nearby contributes to its microclimate via advection- the horizontal transport of air via surface winds. The instrument exploits differences in a property's *orientation* relative to nearby temperature gradients. I refer to this measure as each home's "gradient alignment score": an index that measures how closely the direction of surface airflow lines up with a property's orientation toward the hottest and coolest surrounding it. The score therefore captures how strongly a property is influenced by neighboring relatively hot or cool air while remaining independent of the local temperature level.

For an instrument to work it must (i) be plausibly unrelated to unobserved drivers of housing demand (exclusion) and (ii) generate meaningful variation in local temperature (relevance). In this context, I argue that exclusion holds because a property's orientation toward the hottest and coolest pockets of urban land surrounding it is plausibly random with respect to other price drivers at the small spatial scales examined in this paper. I support this assertion by documenting significant variation in property orientation relative to its surrounding maximum and minimum temperature zones within tracts (Figure 5), suggesting that this measure is not correlated with other unobserved neighborhood attributes that typically vary less considerably within urban tracts.

Relevance follows from the use of high-resolution airflow trajectories that-when paired with a property's orientation relative to nearby temperature gradients- show how much that property's microclimate is influenced by outside pockets of hot and cool air. This measure produces substantial, measurable temperature differences between blocks and neighborhoods that the instrument exploits. Below I describe the meteorological process involved (advection), discuss in more detail how the instrument isolates exogenous temperature variation, and describe how the gradient-alignment score is constructed.

Temperature transfer from urban "hot" and "cool" spots to surrounding areas via surface airflow is well established in climate and atmospheric science. Local airflow patterns redistribute hot and cool air in cities, creating temperature gradients at neighborhood scales, a process referred to as urban heat advection (Heaviside et al. 2015; Bassett et al. 2016; Oke et al. 2017). In Birmingham, UK, for example,

a meteorological model of the August 2003 heatwave attributes up to 2.5°C of intra-urban temperature variation to advection, and Bassett et al. (2016) use in situ measurements to show that hot air from impervious areas within cities "exported" to nearby districts, raising temperatures downwind by 1.2°F on average. Moustaoui and Georgescu (2025) similarly document a significant annual warming contribution of advection from upstream urban development in Phoenix, Arizona. Working at a much finer spatial scale, Middel et al. (2014) find that surface air flow plays a significant role in distributing warm and cool air over neighborhoods in simulated experiments in Phoenix, Arizona. Figure 3 uses an adapted figure from Middel et al.(2014) to illustrate the role of urban air advection via surface winds on air temperatures across a simulated city block in Phoenix[3], the largest city in my study area.

Advection also transmits cooling. Urban vegetation patches generate "park cool islands"–zones of lower temperature inside and around parks (Declet-Barreto et al. 2013; Rasul et al. 2017). Observational studies document that the cooling influence of parks extends farther downwind than upwind because cooler air produced over vegetation is carried by prevailing winds (Norouzi et al. 2024). In many cases, airflow over cool islands transports cooler air hundreds of meters downwind of park boundaries (Sugawara 2016; Motazedian 2020), and similar patterns arise around non-vegetative cooling interventions such as cool pavements (Ko et al. 2020). Taken together, this literature implies that an individual parcel's microclimate depends not only on its own and adjacent land cover, but also on where it sits relative to surrounding hot and cool areas and the directions from which air tends to arrive.

I leverage this phenomenon to construct an instrument for neighborhood temperature that both varies at fine spatial scales, and satisfies the exclusion restriction. For each property parcel, I first characterize the spatial pattern of heat in an "outlying" ring around its 500m immediate surroundings, identifying directional slices that are relatively hotter or cooler than the ring average. Importantly, this calculation does not contain any information regarding the level value of temperature surrounding each property parcel because the warmer and cooler areas are defined relative to each property. I then use high resolution air flow trajectories which I generate from NOAA's Hybrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model-a widely used atmospheric modelling system capable of simulating hourly trajectories of surface air originating over a single spatial point- to recover the frequency

with which nighttime surface air passes over each each residential property from each compass direction (further detail in Appendix section C). Combining these two inputs I define a property-level "gradient alignment score" that measures how often surface air winds blow from relatively hot versus relatively cool parts of each home's surrounding neighborhood.

Intuitively, the gradient alignment score is high and positive when prevailing winds align with hotter slices of the outlying buffer ring, implying that advection tends to warm the parcel's microclimate. It is negative when winds predominantly arrive from cooler slices, implying net advective cooling, and near zero when winds traverse over relatively hot and cool slices symmetrically. Critically, the instrument is constructed from *relative* temperature anomalies in the outlying ring and from the angular relationship between those anomalies and local airflow; it does not use the level of the property's own microclimate or the mean temperature of the surrounding microclimate. As a result, the gradient alignment score shifts local temperatures by altering the contribution of advected hot and cool air, while remaining orthogonal by construction to the level of nearby temperatures. Because parcels that are only a block apart can face very different orientations toward hot and cool slices, and because these micro-scale orientation differences are unrelated to unobserved neighborhood quality, the gradient alignment score generates plausibly exogenous variation in neighborhood heat exposure at fine and varying spatial scales relative to each residential property's location.

5.1 Gradient Alignment Score

To implement the airflow-based instrument at the parcel level, I create a "gradient alignment" score that captures whether nighttime winds tend to bring hot or cool air toward each property. In the baseline specification, I define for each property a 500 m buffer as its immediate surroundings and an outlying ring that extends 500 m beyond that buffer. I divide the outlying ring into eight directional slices (N, NE, ..., NW), and for each residential property i in year t I compute July nighttime minimum air temperatures in each slice d and form a temperature anomaly α_{itd} that measures how much hotter or cooler the slice is relative to the ring's mean. I normalize these anomalies by the range of temperatures in the ring so they represent the relative "heating" or "cooling" potential of each slice. This normaliza-

tion also ensures the instrument is orthogonal to overall temperature variation in the outlying buffer, which might otherwise correlate with urban-form features.

I use near-surface meteorological data from NOAA's High Resolution Rapid Refresh (HRRR) dataset as the primary input into NOAA's HYSPLIT model, which is used to model the frequency of airflow over each property from each cardinal direction. HYSPLIT's air parcel trajectory modeling feature models air "parcel" paths through space and time by simulating the movement of surface air through time using gridded meteorological fields (wind speed, direction and related variables), producing backward paths that show where surface air above a given property has traveled from. I use this feature to generate backward air parcel trajectories (denoting origin of incoming air one hour prior) for each residential property in Maricopa county, in each of three sample hours for two sample days in July for each year of historic HRRR data (2015-2019), at 10m above surface level. I run HYSPLIT's air parcel trajectory module for sample periods ending at 7a, 1pm, and 7pm for each sample day in each year. I then recover, for each parcel and year, the frequency with which nighttime air arrives from each compass direction and the corresponding mean bearing. θ_{it}. Let θ_d denote the central angle of slice d. I define an alignment weight:

$$A_{itd} = 1 - \frac{\min\left(|\theta_d - \theta_{it}|,\, 360^\circ - |\theta_d - \theta_{it}|\right)}{180^\circ}, \tag{1}$$

which equals 1 when the wind bearing and slice direction coincide and declines linearly to 0 as the angular difference reaches 180°. Let ω_{id} denote the frequency weight for trajectories from direction d in each year, and let α_{itd} represent how much hotter or cooler slice d is relative to the ring average, normalized by the range of temperatures in the ring. The gradient alignment score for parcel i in year t is then

$$\Gamma_{it} = \sum_{d \in \mathcal{D}} \omega_{id}\, A_{itd}\, \alpha_{itd}, \tag{2}$$

where $\mathcal{D}$ indexes the eight directional slices. Intuitively, Γ_{it} is positive when prevailing winds align with hotter slices of the outlying ring (net advective warming), negative when winds align with cooler slices (net advective cooling), and close to zero when airflow is aligned with hot and cool slices symmetrically. Because the score is constructed from *relative* temperature anomalies in the outlying ring and from the angular relationship between those anomalies and local airflow–rather than from the level of local temperatures–it shifts microclimates by altering the contribution of advected hot and cool air while

remaining orthogonal by construction to the underlying temperature field. A schematic illustration of this logic is provided in Figure 4, and full computational details are given in Appendix C.

Figure 5 visualizes the spatial variation in property-level bearings that underlie the gradient-alignment instrument. Panels (a) and (b) map each property's average bearing toward the coolest (a) and warmest (b) directional slice in the outlying ring, with census tract boundaries overlaid for reference. Even within the already fine spatial scale of tracts, the maps exhibit substantial within-tract heterogeneity in the direction of relatively hot and cool surroundings. This pattern supports the identifying intuition that parcel orientation relative to nearby hot and cool zones varies at a much finer spatial scale than tract-level characteristics that influence housing prices. Panel (c) maps each parcel's frequency-weighted mean bearing of incoming airflow trajectories, summarizing the predominant direction from which air arrives at the parcel across the HYSPLIT simulations.

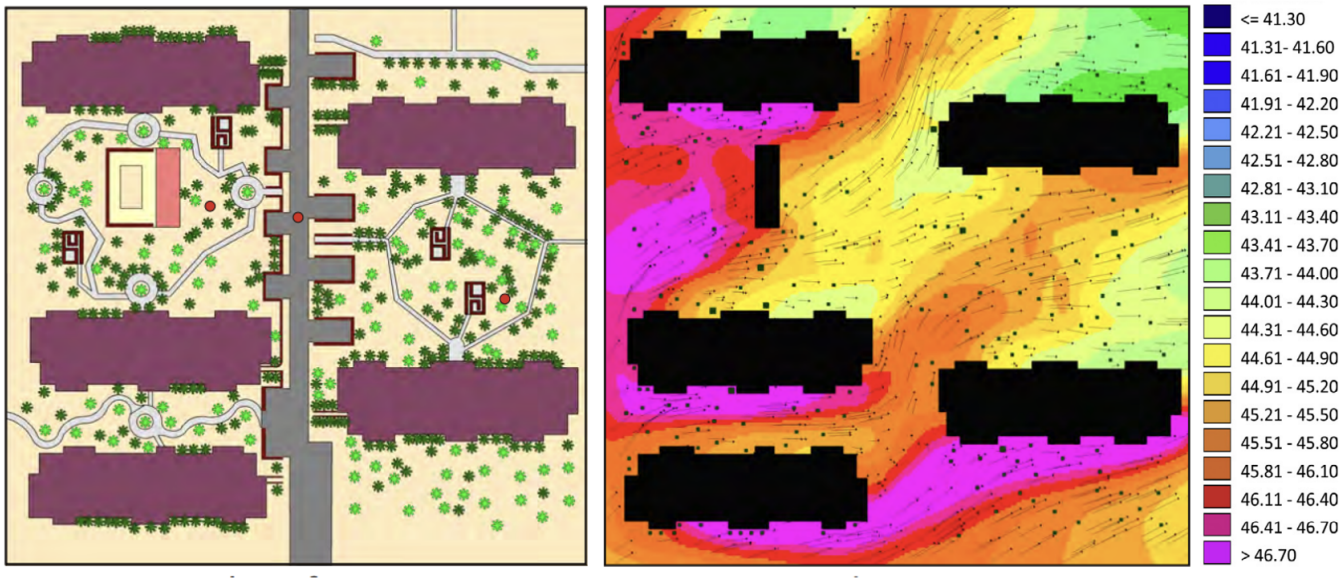


Figure 3: Patterns of surface air temperature distribution in a simulated city block in Phoenix, Arizona. The left panel represents the location of buildings and vegetation within the residential block. The right panel represents 2m surface air temperature with wind vectors overlaid. Wind vectors display a point at the downwind end of each vector to signal the direction of airflow. Simulated temperature patterns (°C) show transfer of warmed air from buildings and cooled air from natural land cover with the direction of near surface winds. Figure adapted from Middel et al. (2014)

[3]Adapted under the terms of the original publication's license. License can be accessed here: https://creativecommons.org/licenses/by-nc-sa/3.0/

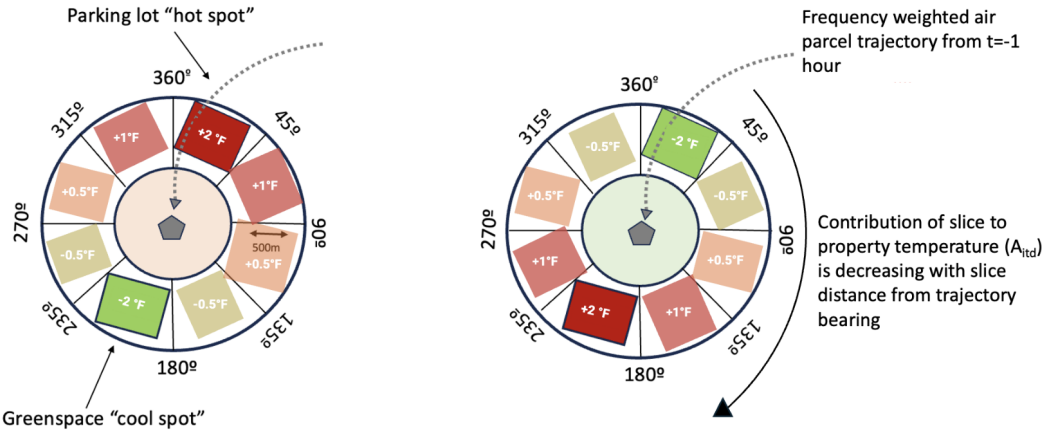


Figure 4: Illustration of the gradient alignment instrument. Conditional on mean buffer temperature, the orientation of neighborhood "hot" and "cool" spots relative to prevailing winds isolates exogenous variation in 2m air temperature within 500 m of each parcel.

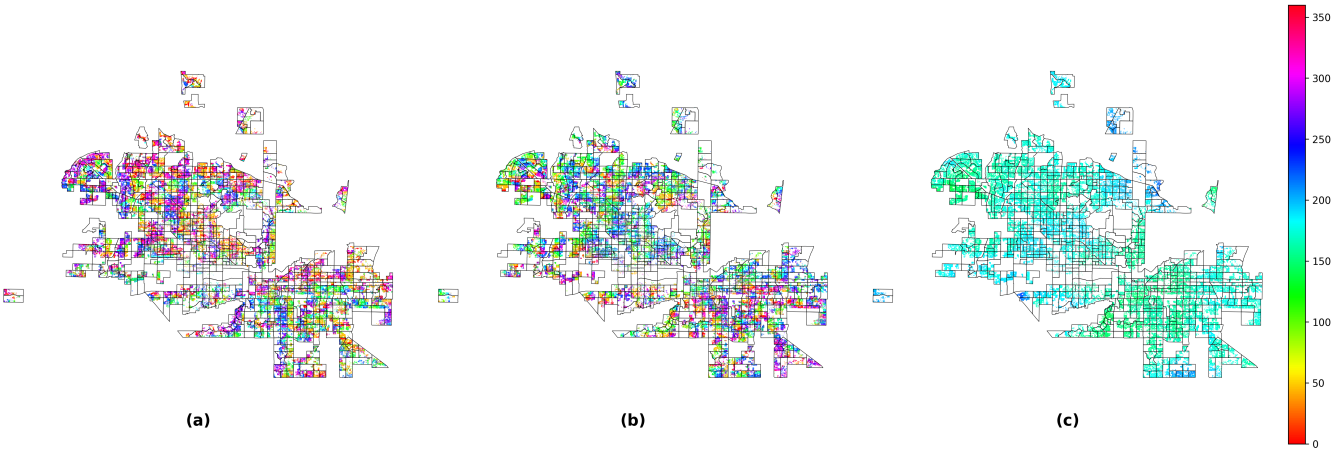


Figure 5: Parcel orientation and airflow bearings. Panels (a) and (b) plot heatmaps of each property's average bearing relative to its minimum (a) and maximum (b) temperature buffer-ring "slice," with 2010 census tract boundaries overlaid. The maps show substantial variation even within tract boundaries, revealing that parcel orientation relative to hotter and cooler surrounding zones varies at a finer spatial scale than tract-level characteristics that influence home prices. Panel (c) plots the frequency-weighted airflow trajectory bearing for each property across the sample.

6 Empirical Strategy

6.1 Econometric Models

In this section I outline the instrumental variable (IV) models I use to identify household MWTP for urban cooling at both the immediate (500m) scale in my baseline specification and at outlying neighborhood scales using a spatial lag model. I then employ a differences in differences (DiD) model in

order to estimate cooling benefits from Phoenix's Cool Corridors Program. This framework allows me to calculate benefits from Phoenix's Cool Corridor Program under two assumptions, i) households value temperature reductions from cool corridors uniformly and at the tract level and ii) household value changes heterogeneously over space and are not necessarily limited to or uniform across administrative boundaries.

6.1.1 Baseline IV Hedonic Model

My baseline specification is a two–stage least squares (2SLS) hedonic model in which neighborhood temperature within 500 meters of each property is instrumented using the parcel–level gradient alignment score described in Section 5. The first–stage equation is

$$T_{it}^{UHI_{500}} = \alpha_0 + \alpha_\Gamma \, \Gamma_{it} + \boldsymbol{\alpha}_z \mathbf{z}_i' + \theta_{zip} + \phi_{puma \times t} + v_{it}, \tag{3}$$

where $T_{it}^{UHI_{500}}$ denotes the mean July nighttime minimum temperature in the 500 meter buffer around parcel i in year t, Γ_{it} is the gradient alignment score, $\mathbf{z}_i$ is a vector of time invariant parcel–level controls (home characteristics, elevation, distance to Phoenix's Central Business District), θ_{zip} denotes ZIP–code fixed effects, and $\phi_{puma(i)\times t}$ denotes PUMA–by–year fixed effects, allowing annual housing market trends to differ at sub-regional scales cross the county.

Using the predicted values $\widehat{T}_{it}^{UHI_{500}}$ from (3), the second–stage equation for annualized log house prices is

$$\ln p_{it}^{ann} = \beta_0 + \beta_\Gamma \, \widehat{T}_{it}^{UHI_{500}} + \boldsymbol{\beta}_z \mathbf{z}_i' + \theta_{zip} + \phi_{puma(i)\times t} + u_{it}. \tag{4}$$

The coefficient β_Γ captures the semi-elasticity of housing prices with respect to temperature in the 500 meters surrounding each home in the sample, and represents the estimated percentage change in $\ln p_{it}^{ann}$ associated with a 1°C increase in July temperatures. In the results section, Table 1 reports first–stage estimates for the 500 meter buffer and tract–level temperatures (tract level estimates are obtained by replacing $T_{it}^{UHI_{500}}$ and $\widehat{T}_{it}^{UHI_{500}}$ in the equations above with 500 meter buffer temperature averaged over the tract scale, which are instrumented using the tract average gradient alignment score). In both 500 meter buffer and tract-level analyses, the gradient alignment score enters with the expected positive sign, and the associated first–stage F–statistics are large, indicating a relevant instrument.

Second stage results are shown in Table 2 in Section 7.

6.1.2 Spatial Lag Model

Because the airflow gradient IV method does not rely on granular fixed effects to improve causal inference, I am able to test the hypothesis that households not only value urban heat reduction within their immediate (500 meter) locale, but in outlying neighborhood zones as well. Specifically, I estimate the following spatial lag model:

$$\ln p_{it}^{ann} = \beta_0 + \sum_{b \in \mathcal{B}} \beta_{\Gamma b} \widehat{T}_{it}^{UHI_b} + \boldsymbol{\beta}_z \mathbf{z}_i' + \theta_{zip} + \phi_{puma(i)\times t} + u_{it}, \tag{5}$$

In equation (5), the index b denotes the spatial field (distance band) around parcel i, with $b \in \{0\text{–}500,\ 500\text{–}1500,\ 1500\text{–}2500,\ 2500\text{–}3500\}$ meters. For each band b, $\widehat{T}_{it}^{UHI_b}$ is the instrumented measure of local heat exposure averaged over that ring. For all buffer rings beyond 500 meters, I instrument temperature using the gradient alignment score calculated over a 1000m buffer surrounding each ring, excluding the buffer itself. The coefficient $\beta_{\Gamma b}$ captures the semi-elasticity of housing prices with respect to temperature in band b, holding constant temperatures in the other bands. Results from the spatial lag model are shown in Table 5 in Section 7

6.2 Measuring Cooling Benefits from the Cool Corridors Program

To estimate total benefits from Phoenix's Cool Corridors program I first need the program's effect on local temperatures, I then covert those temperature changes into dollar benefits using my MWTP estimates from the hedonic regressions. Because the program was implemented after my NCAR high-resolution air temperature sample, I cannot include Cool Corridors directly in the hedonic regressions. Instead I proceed in two steps: I first estimate cooling from the Cool Corridors using high resolution (30m) annual land-surface temperature from USGS as a proxy for air temperature with a differences-in-differences design, and (2) apply the MWTP estimates from the hedonic analysis to those temperature changes to get household benefits. Subsection 6.2.1 describes the DiD approach for estimating cooling; subsection 6.2.2 explains how I combine those cooling estimates with the hedonic results to calculate the average value generated per Cool Corridor.

6.2.1 DiD Specification

Urban temperatures change for many reasons besides Cool Corridors (citywide trends, weather shocks, neighborhood change), so simple before–after comparisons would conflate program effects with background variation. A difference-in-differences (DiD) design compares post-installation temperature changes on treated corridors to contemporaneous changes on similar untreated corridors to isolate cooling impacts. But with staggered adoption, conventional two-way FE DiD can implicitly compare early and late adopters and be biased when effects vary over event time (Goodman 2021; Chaisemartin and D'Haultfoeuille 2020). To prevent those problematic comparisons, I use a stacked DiD that limits controls to the appropriate cohort (Butters et al. 2020) and apply corrective stack weights following Wing et al. (2024) so control observations are reweighted consistently across stacks, removing bias from varying treatment shares across cohorts.

To implement this design, I compare temperatures of each treated cohort (corridors first installed in 2022, 2023, or 2024) to a control group of corridors that were selected for the Cool Corridors program but were not yet treated during the cohort-specific estimation window. Planned/scheduled corridors are a natural counterfactual because they were identified using the same siting criteria and planning process as completed corridors, and the timing of installation is largely driven by administrative rollout (funding cycles, contracting capacity, and other institutional constraints) rather than by short-run changes in local surface temperature, so treatment timing is plausibly orthogonal to the cooling potential being measured.

My stacked DiD event-study model is expressed as:

$$Y_{its} = \alpha_{is} + \gamma_{ts} + \sum_{k \in \mathcal{K}, k \neq -1} \beta_k \, \mathbf{1}\{\text{rel}_{its} = k\} \times \text{Treat}_{is} + \varepsilon_{its}, \tag{6}$$

where $\mathcal{K} = \{-3, -2, -1, 0, 1, 2, 3\}$, $k = -1$ is omitted and serves as the reference period. Each β_k represents the difference between treated and control groups at event time k, relative to the treated-minus-control difference at event time -1.

Let s index the sub-experiment (cohort) and denote the stacked outcome as Y_{its}. I define Treat_{is} as an indicator equal to 1 if corridor i is treated in sub-experiment s (i.e. $G_i = s$), and 0 if corridor i is a clean control in sub-experiment s, $\text{rel}_{its} = t - s$ represents event time measured relative to the sub-experiment cohort year s.

α_{is} represents unit by sub-experiment fixed effects, which absorb all time-invariant differences for corridor i within each sub-experiment s. γ_{ts}, represents all shocks common to all corridors in calendar year t within each sub-experiment s.

The weights used to adjust for differences in control observations across sub-experiments are expressed as:

$$Q_{i,s} = \begin{cases} 1, & \text{if } D_{i,s} = 1, \\ \dfrac{\left(N_s^T / N_\Omega^T\right)}{\left(N_s^C / N_\Omega^C\right)}, & \text{if } D_{i,s} = 0, \end{cases} \tag{7}$$

where $D_{i,s}$ indicates whether unit i is treated in sub-experiment s, N_s^T is the number of treated units in sub-experiment s, N_Ω^T is the total number of treated units in the stacked sample. Similarly, N_s^C represents the number of control units in sub-experiment s and N_Ω^C represents the total number of control units in the sample.

I implement these $Q_{i,s}$ weights as analytic weights in the stacked event-study regression. Intuitively, the weighting prevents sub-experiments with disproportionately many control observations from dominating the counterfactual trend, ensuring that the control trend is aggregated in proportion to the treated cohort's contribution to the stacked sample.

6.2.2 Estimating the Value of Benefits

To evaluate the program's cooling benefits I combine MWTP from the hedonic analysis with estimated temperature reductions during implemented corridor years. I compute benefits under two capitalization assumptions: (1) cooling values accrue only at the tract scale, and (2) cooling values accrue heterogeneously across space and may extend beyond the tract. I test sensitivity to the beneficiary population by calculating results for (a) all residential parcels in Phoenix's parcel database and (b) only parcels

that sold during the hedonic sample period. I also report results with and without a 100 meter cooling buffer around each corridor.

For tract-scale accrual I use MWTP from the tract-level IV; for spatially heterogeneous accrual I use MWTP from the spatial-lag model across all buffer rings where MWTP is significant. For each property I compute the annual area overlap between its tract (or ring buffers) and the Cool Corridors in the DiD sample. I assign the corridor's total land-surface cooling uniformly across the corridor region (with and without a 100 m buffer) and scale that cooling by the share of the tract or ring area that overlaps the corridor. I convert corridor LST cooling to air-temperature cooling using the 0.63 Phoenix LST to air factor from Song et al. (2017). Finally, I translate temperature changes into dollar values by multiplying the cooling estimate by either (a) the median annualized housing value in the spatial zone or (b) each property's observed annualized home value from the hedonic sample.

7 Results

In this section I discuss results from my baseline IV specification using each property's heat gradient alignment score as an instrument for local (500m) temperature. I then compare the estimates from this specification to estimates from several other econometric specifications, and perform a series of robustness checks for my baseline specification. I then report results from my spatial lag model and compare projected total cooling benefits from Phoenix's Cool Corridor program across a range of assumptions.

7.1 Baseline IV Results

Table 2 shows estimated household MWTP for urban heat using the gradient alignment IV. All hedonic regression coefficients (Tables 2-5) are expressed as a percentage of annualized home price. Column (1) reflects a specification in which the urban heat measure is taken within the 500m radius immediately surrounding each home in the sample, and (2) reflects urban heat aggregated to the tract scale. Both specifications show a negative and significant MWTP to live in hotter locales, -8.3% ($85.86) and -10.7% ($110.68) of annualized home value, respectively, and calculated at the mean of sample annualized home value in 2025 dollars ($12,413). Table 1 shows results from the first stage for each

IV regression in table 2. Both regressions have large F-statistics, indicating a relevant instrument, and reflect an intuitive relationship between gradient alignment score temperature; the gradient alignment score is highest when the path of air transfer passes over the hottest micro-climates directly surrounding it, thus the instrument should impact temperature positively.

Alternate hedonic specifications reflect identification challenges associated with MWTP for urban heat. Further detail on these specifications are available in Appendix section D. In Table 3, columns (1) and (2) reflect estimate of MWTP for urban heat from a Taylor-Hausman specification which instruments for urban heat at the tract scale using tract elevation. Because this method relies on between tract comparisons of elevation for identifying variation, estimates are sensitive to the sample selected. Column (1) uses the central sample, in which properties situated in tracts comprised of more than the 95th percentile area coverage of farmland or bare soil are eliminated, reflecting properties whose micro-climates are largely unaffected by development. Using this sample yields an insignificant and positive MWTP for urban heat, by contrast (2) includes these properties and reflects the opposite, a relatively high magnitude, negative and significant MWTP for urban heat (-12.6%). The difference in result sensitivity to sample reflects the between-tract comparisons that the Taylor-Hausman specification relies on– most of the identifying variation in (2) comes from the difference in elevation between rural-urban fringe tracts at much higher elevations and those in the city center, suggesting that unobserved factors associated with living in urban-adjacent undeveloped zones may be influencing estimates, (i.e. views, flood risk, luxury homes).

Columns (3) and (4) of Table 3 show estimates from a two-way fixed effects specification using granular spatial fixed effects at the tract level, and from a repeat sales model, which relies on property level fixed effects to purge the model of endogeneity from time-invariant property characteristics. While the MWTP estimate from the two-way fixed effects model has the expected sign, the estimate is quite low in magnitude compared to the IV result at the same 500m buffer aggregation (2.1%). This likely reflects upward bias from valuable, time-variant amenities such the degree of commercial development that likely positively impact both home values and temperatures. The repeat sales model in column (4) estimates a positive and significant MWTP. Notably the sample used in this specification is much

smaller, because repeat sales models are limited to identifying variation from homes that sold at least twice over the sample period. The MWTP estimate for this specification is slightly positive, 1.4% indicating that homes prefer hotter summer microclimates. This estimate is likely biased from time-varying unobservables; for example, if property values rose the most for homes that experienced a relatively large amount of urban development in their surrounding area, estimates would be biased upward.

In Table 4 I demonstrate the validity of the baseline IV approach by running systematic robustness checks of my main gradient alignment score specification. I test sensitivity of MWTP estimates to clustering at the block and tract-year level (as opposed to clustering at the block-year level in the main specification), adding additional time-varying amenity controls (PM 2.5 air pollution, median tract income, land cover, and feature distances), and winsorizing both temperature and gradient alignment score at the 95% and 90% level. Table 4 shows that my estimate of MWTP stays significant, negative and stable across all checks, suggesting that the gradient alignment score calculation alters variation in temperature throughout the sample, producing estimates of the true average treatment effect of temperature on home prices. Furthermore, MWTP robustness to income and amenity controls support that the exclusion restriction holds for the gradient alignment score.

7.2 Spatial Lag Results

Table 5 reports the results of a spatial lag model in which I estimate the effect of urban heat at the 500m, 1500m, 2500m and 3500m microclimate scales using the gradient alignment score IV specification. Determining impacts separately at these levels would be infeasible using tract fixed effects, due to negligible levels of variation in larger buffer temperatures within the tract scale.

Results from the spatial lag model in Table 5 indicate a negative and significant MWTP for heat at the 1500m scale (15.67%) as well as within the 500m buffer zone (5.37%) used in the baseline specification. Adjusting for spatial scale between buffer size, these estimates correspond to a MWTP of $676 annually for a 1°C decrease in summer air temperatures in immediate surroundings, and $245 annually for equivalent cooling in the broader neighborhood. [4] The estimate for the 500m buffer zone is likely

[4]Because the temperature regressors are spatial averages over different exposure areas, I scale the implied MWTP to

smaller in the spatial lag specification relative to the baseline IV specification due to explicit representation of the microclimate in the 1500m buffer ring . Since urban heat is a spatially correlated amenity, the larger MWTP in the baseline specification was likely picking up on some value from outlying zones being cooler. The positive, insignificant coefficients for urban heat at larger spatial scales likely reflects neighborhood substitution effects between nearby districts. For example, holding urban heat constant at the 500m and 1500m scale, hotter alternative options for housing within a 1500m to 2500m radius may result in households willing to pay more for the same amount of local heat. The positive MWTP estimate for urban heat from 500m to 1500m in the surrounding neighborhood provides evidence that households consider both direct energy cost reduction benefits of urban heat at the hyper local level, as well as the increased complementarity between neighborhood amenities such as parks, shopping districts, and schools, and reduced urban heat at the broader neighborhood scale.

Furthermore, because the 1500m buffer extends outside of most tract delineations, the positive and significant estimate for MWTP at the 1500m suggests that evaluating benefits limited to the tract scale may result in an underestimation of cooling infrastructure benefits, as these benefits analyses exclude homes within this buffer range but outside of the tract scale. I provide concrete evidence of this accounting effect in section 6.4, In which I exhibit my results from Phoenix's Cool Corridors benefits analysis. Furthermore, the results of the spatial lag model suggest that urban heat reduction exhibits significant spillover effects, calling into question the efficiency of market forces alone in allocating heat over urban landscapes. For example, because households value urban heat outside of their immediate microclimate, private decision-makers do not fully internalize how landscaping changes, private investment in cool roofs, or commercial development may impact other households nearby. The presence of these spatial externalities justifies public provision of green and cool infrastructure that accounts for these effects.

Table 6 reports robustness checks for the spatial lag specification, similar to those performed for the baseline IV specification (Table 4) Panel A adds time-varying controls for air quality (tract-level summer PM 2.5 concentration) and median tract income, respectively. In both cases, the MWTP estimates

reflect the relative area represented by each buffer. The 500m effect corresponds to cooling over a disk of area $A_{500} = \pi(500)^2$. The 1500m regressor is constructed as a ring that excludes the inner 500m radius, so its relevant exposure area is $A_{1500\setminus 500} = \pi(1500)^2 - \pi(500)^2$. To express the valuation of a 1°C decrease over the full 1500m disk in a way that is comparable to the 500m estimate, I convert the ring-based MWTP into a per-unit-area value and then rescale by the target exposure area, using the ratio $A_{500}/A_{1500\setminus 500}$ to map the 1500m ring estimate onto an equivalent area basis.

at the 500m and 1500m scales remain negative and significant, and their magnitudes are close to the baseline spatial lag estimates reported in Table 5. The 500m coefficient ranges from −0.051 to −0.054 across these specifications, and the 1500m coefficient ranges from −0.113 to −0.158, compared to −0.054 and −0.157 in the baseline. Robustness to income and air quality controls supports the plausibility of the exclusion restriction for the gradient alignment score, as these are among the most likely sources of time-varying confounding in this setting. Panel B adds controls for neighborhood land cover (tract-level shares of high- and low-intensity development) and feature distances (distance to nearest major road and nearest greenspace). Results are again stable: the 500m estimate ranges from −0.044 to −0.057 and the 1500m estimate from −0.143 to −0.145, both remaining significant at conventional levels. The consistency of the 1500m estimate across all four additional control specifications in Panels A and B reinforces the finding that households value cooling beyond their immediate surroundings.

Panel C tests sensitivity to outlier treatment by winsorizing both temperature and the gradient alignment score at the 95th and 90th percentiles. The 500m and 1500m coefficients remain negative, significant, and similar in magnitude to the baseline (−0.058 to −0.062 at 500m; −0.160 to −0.161 at 1500m). As in the baseline spatial lag model, coefficients at the 2500m and 3500m scales are not statistically significant in either specification, confirming that the detectable household valuation of cooling extends to approximately 1500m from the home. All specifications pass weak identification tests, with large F-statistics.

7.3 Benefits of Cool Corridors

Figure 6a and Table 6 show the results of the event study analysis (equation 6) I perform in order to determine the cooling impact of Phoenix's Cool Corridors between 2022 and 2025. I find significant, negative and decreasing temperature impacts of Cool Corridors with time from corridor implementation. This is intuitive, given the expectation that as trees establish their canopy grows and cooling impacts increase. The event study results imply a cooling impact of -1.6 - -2.4 °C from the first year of installation to the third year after the corridor is installed. [5]

The coefficients prior to corridor implementation in the event study plot in 6a detect no significant

difference in temperature change between corridors that are eventually treated and planned corridors, and Figure 6b suggests that both treatment and control groups were exposed to similar microclimate trends in the years leading up to corridor implementation. This provides evidence in support of the key identifying assumption– that treated and planned-but-not-yet-treated corridors would have followed the same temperature trend had the Cool Corridors program not been implemented.

To further validate the difference-in-differences design, I conduct a spatial displacement placebo test. For each treated corridor, I construct a placebo corridor of identical length on the same road, displaced 500 meters along the road axis in a randomly chosen direction. Placebo corridors that overlap any existing treated, control, or other placebo corridor are excluded, and the placebo corridors retain the treatment timing of their real counterparts. I then re-estimate the stacked event-study model on the displaced segments. Because each placebo corridor shares the same road, surrounding urban form, and satellite measurement environment as its real counterpart– differing only in that no trees were planted there, this test serves as a joint falsification exercise. If the observed cooling were driven by road-level confounders, neighborhood trends that coincided with program timing, or systematic land surface temperature measurement errors favoring certain urban zones during treatment years, the displaced segments would exhibit similar temperature effects.

Table 8 reports the results. All event-study coefficients—both pre- and post-treatment—are statistically insignificant, with point estimates that are considerably smaller in magnitude than the actual treatment effects reported in Table **??**. The null results support the assertion that the cooling effects are spatially localized to the segments that actually received tree canopy improvements.

Table 8 reports the results. All event-study coefficients—both pre- and post-treatment—are statistically insignificant, with point estimates that are considerably smaller in magnitude than the actual treatment effects reported in Table **??**. The absence of significant effects at displaced locations confirms that the cooling signal is spatially localized to corridors that actually received tree canopy improvements, rather than reflecting road-level confounders or broader neighborhood temperature trends that happen to coincide with program timing.

Tables 7 and 8 report the cooling value of cool corridors under several assumptions. Table 7 column (1) uses the MWTP estimate from column (2) of Table 2, in which the benefits of cooling infrastructure are assumed to accrue at the tract level, and columns (2) - (4) record benefits that reflect the results of the spatial lag model from equation (5). All results in Table 7 reflect the assumption that i) benefits accrue to all residential households that reside within the relevant buffer zone or tract containing a cool corridor, not only sold parcels and ii) that each cool corridor transmits cooling to a 100m area surrounding it, consistent with the cooling buffer effect for compact vegetation in urban zones recorded in (Soltanifard and Amani-Beni 2017). Under these assumptions, cooling benefits from cool corridors yield a $500k return per corridor at the tract scale and $987k return for corridor when the full extent of corridor value is included in the benefits calculation. Under both tract and spatial lag assumptions, the benefits of cool corridors significantly outweigh costs of the program, which are estimated to be $135,240 per corridor on average between 2022 and 2024, based on $1.4M yearly costs recorded in the City's 2024 Shade Phoenix document (City of Phoenix, 2024).[7] Critically, benefits calculated at the tract scale are $10.52 million dollars 48% less than those calculated in column (4), where the total scale over which households value urban heat improvements are accounted for. As Table 7 suggests, this is due to both 1) underestimation of relatively high benefits accruing to households abutting cool corridors, as well as the inclusion of many more benefiting households when the 500m-1500m ring buffer is included.

[5]Benefits estimated from the differences-in-differences estimator are within range of local cooling estimates based on tree canopy increases, assuming an initial tree canopy are of 8 m^2 and a canopy growth rate of 3 m^2 per year, based on growth trajectories that are consistent with region-specific published urban tree allometry guides. Urban tree growth models used in the Desert Southwest Community Tree Guide and i-Tree Streets documentation report early crown diameter growth rates of 0.5–0.7 m per year for open-grown street trees, which translate into increases of 2–4 m^2 per year in crown-projected area during the first decade after planting.[6]. Middel et al. (2015) find a 0.2°C decrease in air temperature per 1% increase in the urban tree canopy in areas with no pre-existing tree coverage, an appropriate comparison to Phoenix cool corridor sites pre-installation based on the city's priortization process (City of Phoenix 2024). Using the median corridor area of 0.78 km x 0.45 km, and median 76 trees per completed corridor, the back of the envelope calculation implies a 1.3 °C decrease in corridor temperature the first year trees are planted and a 2.1 °C decrease by the third year after planting, within range of cooling estimates for each year after planting obtained from the differences in differences analysis.

[7]To translate the program budget into a per corridor cost, I start from the City of Phoenix reported annual Cool Corridors program cost of $1.4 million and assume this annual cost applies to the 2022 to 2024 planting period. I convert the resulting three year total to 2025 dollars (total $4,508,000). My DiD analysis is restricted to corridors that are at least 400 meters long and that achieve a minimum planting density of 62 trees per kilometer, approximately half of the City's stated goal of 200 trees per mile, since shorter or sparser corridors are less likely to generate temperature signals detectable in Landsat land surface temperature data, which is sampled at 100 meter resolution and resampled to 30 meters. This restriction yields 20 corridors; corridors within 100 meters of cool pavement zones or within 100 meters of another corridor are excluded from the DiD sample to limit confounding, but are added back into the corridor inventory used for the benefits calculations. To align costs with the subset of corridors used for estimation, I assume program costs are proportional to trees planted, and the 20 corridors meeting the length and density criteria account for 60% of trees planted over 2022 to 2024. I therefore allocate 60% of the total three year program cost to this subset and divide by 20, implying an average cost of approximately $122,945 per corridor in 2025 dollars.

Using the results that value accrues heterogeneously over space, I transition to exhibiting benefit calculations over a range of assumptions regarding 1) the extent of the cooling buffer surrounding corridors and 2) the homes that benefits accrue to. In Table 8, I perform the same benefit calculation as expressed in column (4) of Table 7, but which tests the sensitivity of benefits to three further assumptions; i) benefits only accrue to homes that were observed as sold during the sample period reflected in the hedonic model, 39% of the total sample of Phoenix residential homes recorded in Maricopa County's parcel database, and ii) cooling effects of cool corridors do not extend past the boundary of the corridor itself, and iii) the most conservative scenario, in which assumptions (i) and (ii) are combined. I find that in all but scenario (iii) the estimated benefits of Phoenix's Cool Corridor program outweigh the costs. Notably, it is very unlikely that the assumptions in scenario 3 hold, as it is highly likely that benefits accrue to households that did not sell their homes in the period covered in the hedonic analysis, and that the cooling effects of the corridors are strictly limited to the corridor zone itself.[8]

[8] A large urban climate literature documents that vegetation generates cooling effects that extend beyond the vegetated footprint through shading and air movement. In Phoenix, microclimate simulations of an inner city park show a Park Cool Island effect that extends cooling to nearby non vegetated surfaces (Declet-Barreto et al. 2012), and neighborhood scale simulations indicate that increasing tree canopy cover measurably reduces near surface air temperatures at the local scale (()Middel, 2015). More generally, empirical and remote sensing studies find that cooling from parks and green spaces is detectable tens to hundreds of meters beyond park boundaries depending on park size, vegetation structure, and wind conditions (Aram et al. 2019).

Table 1: First Stage Estimation Results For Baseline Specification

July Nighttime Min. ($°C$)	(1) Gradient IV (500m)	(2) Gradient IV (Tract)
Urban Heat Measure		
Gradient Alignment Score (500m)	0.4347***	
	(0.0116)	
Gradient Alignment Score (Tract)		0.3381***
		(0.0154)
PUMA x Year Fixed Effects	Yes	Yes
Zip Code Fixed Effects	Yes	Yes
Observations	498,522	498,584
K-P Wald F-statistic	1396.69	481.72

Notes: This table reports first-stage estimates from the baseline IV specification. The dependent variable is average July nighttime minimum air temperature (°C), measured within a 500m buffer around each property in column (1) and averaged to the census tract level in column (2). The instrument is the parcel-level gradient alignment score, which measures the degree to which prevailing surface airflow aligns with relatively hot or cool areas in each property's outlying buffer ring. All specifications include ZIP code fixed effects and PUMA × year fixed effects. The sample consists of non-foreclosure, arms-length residential home sales in Maricopa County, AZ from 2009–2018. Standard errors clustered at the block level are in parentheses. The K-P Wald F-statistic is the Kleibergen-Paap rk Wald F statistic for weak identification. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Baseline IV Estimation Results (Selected Variables)

Log Annualized Home Price (2025 Dollars)	(1) Gradient IV (500m)	(2) Gradient IV (Tract)
Urban Heat Measure		
500m Buffer July Nighttime Min. (°C)	-0.0826***	
	(0.0154)	
Tract July Nighttime Min. (°C)		-0.1068***
		(0.0392)
Sale Year Fixed Effects	Yes	Yes
PUMA × Year Fixed Effects	Yes	Yes
Zip Code Fixed Effects	Yes	Yes
Observations	498,522	498,584

Standard errors in parentheses

Only selected variables are included in this table. All regressions are clustered at the block level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Estimation Results From Alternate Specifications (Selected Variables)

Log Annualized Home Price (2025 Dollars)	(1) Hausman-Taylor (Tract)	(2) Hausman-Taylor: Extended Sample	(3) Two-Way FE (500m)	(4) Repeat Sale (500m)
Urban Heat Measure				
500m Buffer July Nighttime Min. (°C)			-0.021***	0.014***
			(0.0043)	(0.0048)
Tract July Nighttime Min. (°C)	0.056	-0.126***		
	(0.04424)	(0.0249)		
Sale Year Fixed Effects	Yes	Yes	Yes	Yes
PUMA × Year Fixed Effects	No	No	Yes	Yes
Census Tract Fixed Effects	No	No	Yes	No
Parcel Fixed Effects	No	No	No	Yes
Observations	498,522	554,085	498,521	180,607

Notes: This table reports second-stage IV estimates of household marginal willingness to pay (MWTP) for summer temperature reductions. The dependent variable is log annualized home price in 2025 dollars. Column (1) instruments for average July nighttime minimum air temperature within a 500m buffer around each property using the parcel-level gradient alignment score. Column (2) instruments for the same temperature measure averaged to the census tract level using the tract-average gradient alignment score. Coefficients represent the semi-elasticity of annualized home price with respect to a 1°C increase in July nighttime minimum temperature. Only selected variables are shown; all specifications include parcel-level controls (square footage, structure type, bathrooms, pool, garage, year built, elevation, distance to CBD). The sample consists of non-foreclosure, arms-length residential home sales in Maricopa County, AZ from 2009–2018. Standard errors clustered at the block level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Estimation Results From Spatial Lag Model (Selected Variables)

	(1) IV Spatial Lag Model
July Nighttime Min. (°C)	
500m	-0.0537***
	(0.0192)
1500m	-0.1567**
	(0.0639)
2500m	0.1448
	(0.1083)
3500m	0.0086
	(0.1021)
Observations	498,522

Notes: This table reports second-stage IV estimates from a spatial lag model in which July nighttime minimum air temperature is measured separately in concentric distance bands around each property (Equation (5)). Each band's temperature is instrumented using the gradient alignment score calculated over a 1000m ring surrounding that band. The dependent variable is log annualized home price in 2025 dollars. Coefficients represent the semi-elasticity of annualized home price with respect to a 1°C increase in temperature within the indicated distance band, holding temperatures in all other bands constant. Only selected variables are shown; all specifications include parcel-level controls, ZIP code fixed effects, and PUMA × year fixed effects. The sample consists of non-foreclosure, arms-length residential home sales in Maricopa County, AZ from 2009–2018. Standard errors clustered at the block level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Robustness checks of Main IV Specification

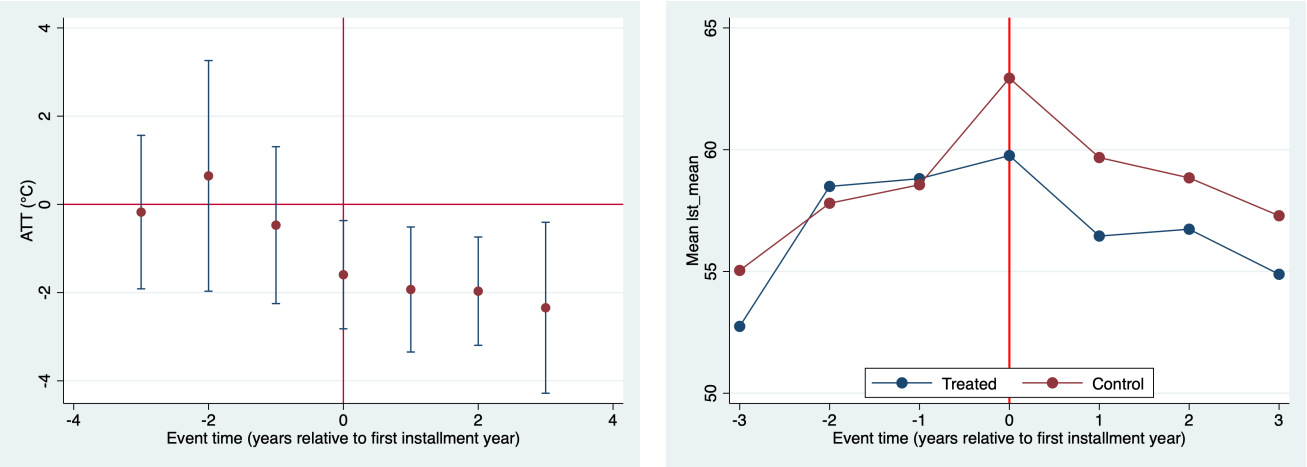
Log Annualized Home Price (2025 Dollars)	(1)	(2)
Panel A: Clustering		
Min. July Temp. (500m)	-0.083***	-0.083***
	(.0293)	(0.0312)
Clustering	Tract	Zip Code
Observations	498,522	498,522
F-statistic	415.12	378.18
Panel B: Additional Controls: Air Quality and Income		
Min. July Temp. (500m)	-0.074***	-0.072***
	(.0158)	(.0153)
Controls	Air Quality	Income
Observations	498,522	498,522
F-statistic	1769.88	1714.77
Panel C: Additional Controls: Land Cover and Feature Distances		
Min. July Temp. (500m)	-0.071***	-0.084***
	(.0161)	(.0153)
Controls	Land Cover	Feature Distances
Observations	498,522	498,522
F-statistic	1392.12	1425.41
Panel D: Outlier Treatment		
Min. July Temp. (Tract)	-.0869***	-.0904***
	(0.0088)	(0.0097)
Outlier Treatment	Winsorized: 95%	Winsorized: 90%
Observations	498,522	498,522
F-statistic	1691.09	1578.83

Notes: Each column of each column reports a robustness chech for the second-stage IV estimates from the baseline specification in which average July nighttime minimum air temperature within a 500m buffer is instrumented using the parcel-level gradient alignment score. The dependent variable is log annualized home price in 2025 dollars. Panel A varies the level at which standard errors are clustered. Panel B adds time-varying controls: tract-level summer $PM_{2.5}$ concentration (Air Quality) and real median tract income (Income). Panel C adds tract-level shares of high- and low-intensity development (Land Cover) and distance to the nearest major road and nearest greenspace (Feature Distances). Panel D winsorizes both temperature and the gradient alignment score at the indicated percentiles. All specifications include ZIP code fixed effects, PUMA × year fixed effects, and parcel-level controls. The sample consists of non-foreclosure, arms-length residential home sales in Maricopa County, AZ from 2009–2018. Standard errors clustered at the block level are in parentheses unless otherwise noted. The F-statistic is the Kleibergen-Paap Wald F statistic for weak identification. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Robustness checks of Spatial Lag Specification

Log Annualized Home Price (2025 Dollars)	(1)	(2)
Panel A: Additional Controls: Air Quality and Income		
Min. July Temp. (500m)	-0.054***	-0.051***
	(0.0192)	(0.0190)
Min. July Temp. (1500m)	-0.158**	-0.113*
	(0.0639)	(0.0637)
Min. July Temp. (2500m)	0.128	0.107
	(0.1090)	(0.1060)
Min. July Temp. (3500m)	0.017	0.009
	(0.1026)	(0.1001)
Controls	Air Quality	Income
Observations	498,522	498,496
K-P Wald F-statistic	126.07	121.84
Panel B: Additional Controls: Land Cover and Feature Distances		
Min. July Temp. (500m)	-0.044**	-0.057***
	(0.0194)	(0.0192)
Min. July Temp. (1500m)	-0.143**	-0.145**
	(0.0628)	(0.0644)
Min. July Temp. (2500m)	0.118	0.106
	(0.1063)	(0.1106)
Min. July Temp. (3500m)	0.059	0.039
	(0.1019)	(0.1053)
Controls	Land Cover	Feature Distances
Observations	498,522	498,522
K-P Wald F-statistic	116.40	117.08
Panel C: Outlier Treatment		
Min. July Temp. (500m)	-0.058***	-0.062***
	(0.0199)	(0.0211)
Min. July Temp. (1500m)	-0.161**	-0.160**
	(0.0660)	(0.0699)
Min. July Temp. (2500m)	0.180	0.213*
	(0.1118)	(0.1185)
Min. July Temp. (3500m)	-0.027	-0.061
	(0.1046)	(0.1119)
Outlier Treatment	Winsorized: 95%	Winsorized: 90%
Observations	498,522	498,522
K-P Wald F-statistic	116.41	105.75

Notes: Each column reports second-stage IV estimates from the spatial lag model (Equation 5) with the indicated additional controls or outlier treatment. The dependent variable is log annualized home price in 2025 dollars. All specifications include ZIP code and PUMA × year fixed effects. Standard errors clustered at the block level are in parentheses. The K-P Wald F-statistic is the Kleibergen-Paap rk Wald F statistic for weak identification. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.



(a) (b)

Figure 6: Stacked difference-in-differences analysis of Cool Corridors on July median land surface temperature (LST). Panel (a) shows the event-study estimates of the ATT (°C) by event time relative to the first installment year, with 95% confidence intervals. Panel (b) shows the parallel-trends diagnostic for treated and control corridors in the pre-treatment period, and after corridor installation.

Table 9: Cool Corridors benefits analysis. Total benefits reflect value of cooling benefits to all residential households in Phoenix who reside in tracts or buffer zones overlapping cool corridor segments, and under the assumption of a 100m cooling buffer zone surrounding each Cool Corridor (Soltanifard and Amani-Beni 2025)

	Tract (1)	0-500 m (2)	500-1500 m (3)	Spatial Lag (0-1500 m) (4)
Panel A: Value of cooling benefits 2022-2025				
Benefits 2022-2025 (Millions of 2025 $)	11.20	5.85	15.88	21.72
Benefits Per Corridor (2025 $)	509,000	265,900	721,800	987,300
Est. Return ($135k cost per corridor)	3.76	1.97	5.34	7.30
Benefiting HHs	63,637	22,750	87,231	109,981
Total Benefit per benefiting HH (2025 $)	178.78	257.04	182.01	197.53
Panel B: Cooling				
Median Area Corridor Overlap (% total area)	3.03	9.38	2.25	
Median Cooling Benefit by 2025 (°C)	0.05	0.22	0.06	
Panel C: Household value				
MWTP per °C red. in UHI (% ann. home value)	10.7	5.4	15.7	
Median Ann. Home Value (2025 $)	10,629.56	9867.82	10,100.11	

Table 7: Event-study coefficients centered at mean of pre-periods

	(1) Corridor LST
Event time -3	-0.175
	(0.867)
Event time -2	0.646
	(1.266)
Event time -1	-0.471
	(0.863)
Event time 0	-1.594***
	(0.589)
Event time +1	-1.930***
	(0.666)
Event time +2	-1.966***
	(0.586)
Event time +3	-2.348**
	(0.950)
N	394

Notes: This table reports event-study estimates from a stacked difference-in-differences model (Equation (6)) of the effect of Phoenix's Cool Corridors tree-planting program on July median land surface temperature (°C). Treated corridors are segments that received tree plantings in 2022, 2023, or 2024; controls are corridors selected for the program but not yet treated during each cohort's estimation window. Each coefficient represents the difference in temperature between treated and control corridors at event time k, centered at the mean of the three pre-treatment periods ($k = -3, -2, -1$). Observations are corridor-by-year cells. The stacked design uses cohort-specific sub-experiments with corrective weights following Wing et al. (2024) to address bias from staggered adoption. Standard errors are cluster-robust at the corridor level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Spatial Displacement Placebo Test: Stacked Difference-in-Differences Event Study

	(1) LST Mean (°C)
Event time −3	0.171
	(0.672)
Event time −2	0.427
	(0.859)
Event time −1	−0.598
	(0.669)
Event time 0	−0.987
	(0.612)
Event time +1	−0.728
	(0.623)
Event time +2	−0.755
	(0.673)
Event time +3	−0.307
	(0.399)
Observations	393

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each treated corridor is replaced by a placebo corridor of identical length on the same road, displaced 500m along the road axis in a randomly chosen direction. Placebo corridors that overlap any existing treated, control, or other placebo corridor are excluded. The placebo corridors retain the treatment timing of their real counterparts. Coefficients are normalized relative to the mean of the three pre-treatment periods ($t = -3, -2, -1$). Standard errors are cluster-robust at the corridor level. Unit-by-stack and year-by-stack fixed effects are absorbed. All coefficients are statistically insignificant, consistent with the treatment effect being localized to corridors that received tree canopy improvements.

Table 10: Cool Corridors benefits analysis under alternate assumptions. All benefits reflect alternate assumptions with use of spatial lag estimates for valuation

	Sold Homes, 100m buffer (1)	All homes, no buffer (2)	Sold homes, no buffer (3)
Value of cooling benefits 2022-2025			
Benefits 2022-2025 (Millions of 2025 $)	8.38	3.41	1.33
Benefits Per Corridor (2025 $)	380,900	155,000	60,500
Est. Return (est. $135k cost per corridor)	2.82	1.15	0.45
Benefiting HHs	42,366	94,130	87,231
Total Benefit per benefiting HH (2025 $)	197.72	36.27	36.83

8 Conclusion

Extreme heat is increasingly shaping welfare and risk in U.S. cities, yet policymakers have little empirical guidance on how households value hyper local changes in heat exposure and how those values should enter benefit calculations for public urban cooling investments. This paper contributes new evidence on both fronts. I develop a novel instrumental variable strategy that leverages urban air advection and modeled wind trajectories to isolate exogenous variation in neighborhood summer temperatures, and I use this strategy to estimate household marginal willingness to pay for cooler microclimates within the Phoenix metropolitain area. Because the approach does not rely on tract fixed effects, between tract comparisons of temperature, or on cross metro equilibrium adjustments, it allows me to value microclimate improvements at spatial scales that match the physical scale at which urban heat varies.

My central results have three main conclusions. The first is that households are willing to pay to avoid hotter outdoor summer conditions, and the magnitude of the effect is economically significant. In the baseline IV specification, hotter local microclimates reduce housing values, consistent with mechanisms such as higher cooling expenditures, reduced outdoor comfort, and heightened heat stress. I show that this novel instrumental variable approach is robust to a number of tests, and demonstrate important differences between results from my instrumental variable approach and other econometric specifications. Second, the value of cooling to houeholds is not confined to a single administrative boundary or a single spatial scale. The results of my spatial lag analysis indicate that households value

cooling both in their immediate surroundings ($56.36 monthly per 1 °C reduction) and in the broader neighborhood ($20.45 monthly per 1 °C reduction, adjusted to reflect a comparable scale), and that the relative valuation differs across these zones. Third, using empirical evidence from a differences-in-differences analysis of the cooling impacts of Phoenix's Cool Corridors Program I show that this spatial heterogeneity matters significantly for how cities should measure benefits from cooling infrastructure investments. I also demonstrate that the values of cooling benefits for the Cool Corridors program outweighs costs in all but the most conservative of evaluation scenarios.

A key implication of this paper is that conventional approaches that value heat reductions only at the administrative boundary scales can understate benefits. In the context of Phoenix's Cool Corridors program, I show that accounting for the full spatial extent over which households value cooling leads to substantially larger ($10.52M) estimated benefits than a tract-level valuation. This difference is driven by two forces documented in the paper. Aggregation to the tract scale dilutes temperature variation that is most salient to households in their immediate vicinity, and it ignores spillovers that extend beyond tract boundaries when cooling affects nearby residents. The benefit calculations therefore provide concrete evidence that the choice of spatial scale is not a technical detail but a primary determinant of urban heat policy valuation.

These findings also contribute more broadly to the economics of urban heat adaptation. The evidence that households value microclimate changes outside their immediate surroundings suggests that cooling exhibits a local public good component. When cooling generates spatial spillovers, private investments in heat mitigation do not internalize benefits to nearby households, which can lead to under-provision relative to the socially efficient level. This strengthens the case for public investment in neighborhood cooling, particularly in dense urban settings where small scale spillovers can affect many residents, and where the distribution of heat risk often overlaps with historical patterns of disinvestment. At the same time, the housing market channel that makes hedonic valuation possible also implies that the benefits of cooling can be accompanied by higher housing costs in cooled areas. The results therefore underscore that urban heat policy is closely tied to housing policy, and that equitable adaptation requires attention to both the spatial incidence of cooling benefits and the local supply constraints that determine who

can remain in neighborhoods with improved microclimates.

Several limitations suggest productive directions for future work. While the airflow based instrument is designed to isolate exogenous variation in heat exposure, further research could explore whether similar identification is feasible in other climates and urban forms. Future work could also connect the housing market valuation results more directly to health outcomes, energy insecurity, and heat related mortality, which are increasingly central motivations for city heat policy. Finally, because the effects of cooling infrastructure likely evolve as local tree canopy grows and maintenance practices change, longer post periods and richer measures of vegetation quality would further improve inference about the durability and timing of cooling benefits.

Despite these limitations, the paper provides new evidence that the economic value of urban cooling infrastructure cannot be estimated reliably without careful attention to endogeneity and to spatial scale. By combining a novel advection based instrument with a direct evaluation of a city-wide cooling program, this paper offers a framework that cities can use to quantify the benefits of microclimate investments and to design policies that recognize spillovers and heterogeneous value across space. As public spending on urban heat mitigation accelerates, these considerations are likely to be central to determining where investments occur, how they are evaluated, and whether they reduce or reinforce existing urban inequities.

References

Abbott, J. K., & Klaiber, H. A. (2011). An Embarrassment of Riches: Confronting Omitted Variable Bias and Multi-Scale Capitalization in Hedonic Price Models. *The Review of Economics and Statistics*, *93*(4), 1331–1342. https://doi.org/10.1162/REST_a_00134

Albouy, D., Graf, W., Kellogg, R., & Wolff, H. (2016). Climate Amenities, Climate Change, and American Quality of Life. *Journal of the Association of Environmental and Resource Economists*, *3*(1), 205–246. Retrieved January 9, 2026, from https://www.jstor.org/stable/26544410

Aram, F., Higueras García, E., Solgi, E., & Mansournia, S. (2019). Urban green space cooling effect in cities. *Heliyon*, *5*(4), e01339. https://doi.org/10.1016/j.heliyon.2019.e01339

Bajari, P., Fruehwirth, J. C., Kim, K. i., & Timmins, C. (2012). A Rational Expectations Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution. *American Economic Review*, *102*(5), 1898–1926. https://doi.org/10.1257/aer.102.5.1898

Bakkensen, L. A., & Ma, L. (2020). Sorting over flood risk and implications for policy reform. *Journal of Environmental Economics and Management*, *104*, 102362. https://doi.org/10.1016/j.jeem.2020.102362

Basu, R., Colaninno, N., Alhassan, A., & Sevtsuk, A. (2024). Hot and bothered: Exploring the effect of heat on pedestrian route choice behavior and accessibility. *Cities*, *155*, 105435. https://doi.org/10.1016/j.cities.2024.105435

Batur, I., Alhassan, V. O., Chester, M. V., Polzin, S. E., Chen, C., Bhat, C. R., & Pendyala, R. M. (2024). Understanding how extreme heat impacts human activity-mobility and time use patterns. *Transportation Research Part D: Transport and Environment*, *136*, 104431. https://doi.org/10.1016/j.trd.2024.104431

Bayer, P., Keohane, N., & Timmins, C. (2009). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, *58*(1), 1–14. https://doi.org/10.1016/j.jeem.2008.08.004

Beltrán, A., Maddison, D., & Elliott, R. (2019). The impact of flooding on property prices: A repeat-sales approach. *Journal of Environmental Economics and Management*, *95*, 62–86. https://doi.org/10.1016/j.jeem.2019.02.006

Biden-Harris Administration invests $4.55 million for community heat resilience through Investing in America agenda — National Oceanic and Atmospheric Administration. (2024, May). Retrieved December 23, 2025, from https://www.noaa.gov/news-release/biden-harris-administration-invests-4-point-55-million-for-community-heat-resilience

Bishop, K. C., Kuminoff, N. V., Banzhaf, H. S., Boyle, K. J., Gravenitz, K. v., Pope, J. C., Smith, V. K., & Timmins, C. D. (2020). Best Practices for Using Hedonic Property Value Models to Measure Willingness to Pay for Environmental Quality. *Review of Environmental Economics and Policy*. https://doi.org/10.1093/reep/reaa001

Bishop, K. C., Kuminoff, N. V., Banzhaf, H. S., Boyle, K. J., von Gravenitz, K., Pope, J. C., Smith, V. K., & Timmins, C. D. (2020). Best Practices for Using Hedonic Property Value Models to Measure Willingness to Pay for Environmental Quality. *Review of Environmental Economics and Policy*, *14*(2), 260–281. https://doi.org/10.1093/reep/reaa001

Borsky, S., Fesselmeyer, E., & Vogelsang, L. (2024). Urban heat and within-city residential sorting. *Journal of Environmental Economics and Management*, *127*, 103014. https://doi.org/10.1016/j.jeem.2024.103014

Butters, R. A., Sacks, D. W., & Seo, B. (2022). How Do National Firms Respond to Local Cost Shocks? *American Economic Review*, *112*(5), 1737–1772. https://doi.org/10.1257/aer.20201524

Carbone, J. C., Lee, S.-K., & Shen, Y. (2021). U.s. household preferences for climate amenities: Demographic analysis and robustness testing. *Climate Change Economics*, *12*(01), 2050016. https://doi.org/10.1142/S2010007820500165

Chow, W. T. L., Brennan, D., & Brazel, A. J. (2012). Urban Heat Island Research in Phoenix, Arizona: Theoretical Contributions and Policy Applications. https://doi.org/10.1175/BAMS-D-11-00011.1

Cool Corridors Program. (n.d.). Retrieved January 6, 2026, from https://www.phoenix.gov/administration/departments/streets/initiatives/cool-corridors.html

Daily Census Tract-Level PM2.5 Concentrations, 2011-2015 [Type: dataset]. (2021, July). Retrieved January 12, 2026, from https://catalog.data.gov/dataset/daily-census-tract-level-pm2-5-concentrations-2011-2015

Daily Census Tract-Level PM2.5 Concentrations, 2016 - 2020 — Data — Centers for Disease Control and Prevention. (n.d.). Retrieved January 12, 2026, from https://data.cdc.gov/Environmental-Health-Toxicology/Daily-Census-Tract-Level-PM2-5-Concentrations-2016/96sd-hxdt/about_data

Daily Census Tract-Level PM2.5 Concentrations, 2016 - 2020 — HealthData.gov. (n.d.). Retrieved January 12, 2026, from https://healthdata.gov/CDC/Daily-Census-Tract-Level-PM2-5-Concentrations-2016/k9st-jhz8/about_data

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, *110*(9), 2964–2996. https://doi.org/10.1257/aer.20181169

Declet-Barreto, J., Brazel, A. J., Martin, C. A., Chow, W. T. L., & Harlan, S. L. (2013a). Creating the park cool island in an inner-city neighborhood: Heat mitigation strategy for Phoenix, AZ. *Urban Ecosystems*, *16*(3), 617–635. https://doi.org/10.1007/s11252-012-0278-8

Declet-Barreto, J., Brazel, A. J., Martin, C. A., Chow, W. T. L., & Harlan, S. L. (2013b). Creating the park cool island in an inner-city neighborhood: Heat mitigation strategy for Phoenix, AZ. *Urban Ecosystems*, *16*(3), 617–635. https://doi.org/10.1007/s11252-012-0278-8

Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., & Reif, J. (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, *109*(12), 4178–4219. https://doi.org/10.1257/aer.20180279

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, *106*(3), 479–524. https://doi.org/10.1257/aer.20131706

The Effect of Natural Disasters and Extreme Weather on Household Location Choice and Economic Welfare. (n.d.). Retrieved December 26, 2025, from https://ideas.repec.org/a/ucp/jaerec/doi10.1086-728887.html

Fan, Q., Fisher-Vanden, K., & Klaiber, H. A. (2018). Climate Change, Migration, and Regional Economic Impacts in the United States. *Journal of the Association of Environmental and Resource Economists*, *5*(3), 643–671. https://doi.org/10.1086/697168

Fan, Y., Wang, J., Obradovich, N., & Zheng, S. (2023). Intraday adaptation to extreme temperatures in outdoor activity. *Scientific Reports*, *13*(1), 473. https://doi.org/10.1038/s41598-022-26928-y

Freeman, R., Liang, W., Song, R., & Timmins, C. (2019). Willingness to pay for clean air in China. *Journal of Environmental Economics and Management*, *94*, 188–216. https://doi.org/10.1016/j.jeem.2019.01.005

Garnache, C. (2023, November). Does the Salience of Climate-Related Risk Affect Asset Prices? https://doi.org/10.2139/ssrn.3398404

Goodman-Bacon, A. (2021a). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

Goodman-Bacon, A. (2021b). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

hali. (2025, December). NJDEP— News Releases — DEP Awards Nearly $11 Million For Carbon Sequestration Projects Through Its Natural Climate Solutions Grant Program (25/P057). Retrieved December 23, 2025, from https://dep.nj.gov/newsrel/25_0057/

Han, L., Heblich, S., Timmins, C., & Zylberberg, Y. (2024, January). Cool cities: The value of urban trees. https://doi.org/10.3386/w32063

Haninger, K., Ma, L., & Timmins, C. (2017). The Value of Brownfield Remediation. *Journal of the Association of Environmental and Resource Economists*, *4*(1), 197–241. https://doi.org/10.1086/689743

Heat Island Reduction Solutions — US EPA. (n.d.). Retrieved December 22, 2025, from https://www.epa.gov/heatislands/heat-island-reduction-solutions

Himmelberg, C., Mayer, C., & Sinai, T. (2005a). Assessing High House Prices: Bubbles, Fundamentals and Misperceptions. *Journal of Economic Perspectives*, *19*(4), 67–92. https://doi.org/10.1257/089533005775196769

Himmelberg, C., Mayer, C., & Sinai, T. (2005b). Assessing High House Prices: Bubbles, Fundamentals and Misperceptions. *Journal of Economic Perspectives*, *19*(4), 67–92. https://doi.org/10.1257/089533005775196769

Hsu, A., Sheriff, G., Chakraborty, T., & Manya, D. (2021). Disproportionate exposure to urban heat island intensity across major US cities. *Nature Communications*, *12*(1), 2721. https://doi.org/10.1038/s41467-021-22799-5

Klaiber, H. A., Abbott, J., & Smith, V. K. (2015, March). Some Like it (Less) Hot: Extracting Tradeoff Measures for Physically Coupled Amenities. https://doi.org/10.3386/w21051

Kuminoff, N. V., Parmeter, C. F., & Pope, J. C. (2010). Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities? *Journal of Environmental Economics and Management*, *60*(3), 145–160. https://doi.org/10.1016/j.jeem.2010.06.001

Li, X. (2015). CAP LTER land cover classification using 2010 National Agriculture Imagery Program (NAIP) Imagery. https://doi.org/10.6073/PASTA/F4ACED7E801F1B5E14B43CF755199C04

Mayor Bass Announces New Momentum to Combat Climate Change Through City Services — Mayor Karen Bass. (2023, November). Retrieved January 10, 2026, from https://mayor.lacity.gov/news/mayor-bass-announces-new-momentum-combat-climate-change-through-city-services

Middel, A., Chhetri, N., & Quay, R. (2015). Urban forestry and cool roofs: Assessment of heat mitigation strategies in Phoenix residential neighborhoods. *Urban Forestry & Urban Greening*, *14*(1), 178–186. https://doi.org/10.1016/j.ufug.2014.09.010

Muehlenbachs, L., Spiller, E., & Timmins, C. (2015). The Housing Market Impacts of Shale Gas Development. *American Economic Review*, *105*(12), 3633–3659. https://doi.org/10.1257/aer.20140079

Schlenker, W., & Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, *83*(2), 768–809. https://doi.org/10.1093/restud/rdv043

Soltanifard, H., & Amani-Beni, M. (2025). The cooling effect of urban green spaces as nature-based solutions for mitigating urban heat: Insights from a decade-long systematic review. *Climate Risk Management*, *49*, 100731. https://doi.org/10.1016/j.crm.2025.100731

Song, J., Wang, Z.-H., Myint, S. W., & Wang, C. (2017). The hysteresis effect on surface-air temperature relationship and its implications to urban planning: An examination in Phoenix, Arizona, USA. *Landscape and Urban Planning*, *167*, 198–211. https://doi.org/10.1016/j.landurbplan.2017.06.024

TIGER/Line Shapefile, 2019, state, Arizona, Primary and Secondary Roads State-based Shapefile [Type: dataset]. (n.d.). Retrieved January 11, 2026, from https://catalog.data.gov/dataset/tiger-line-shapefile-2019-state-arizona-primary-and-secondary-roads-state-based-shapefile

Tree Shade Programs. (n.d.). Retrieved January 10, 2026, from https://www.phoenix.gov/administration/departments/heat/tree-shade-programs.html

US Department of Commerce, N. (n.d.). Annual and Monthly Record data for Phoenix. Retrieved December 25, 2025, from https://www.weather.gov/psr/PhoenixRecordData

US EPA, O. (2016, April). New York City Adapts To Deal with Projected Increase of Heat Waves. Retrieved December 23, 2025, from https://www.epa.gov/arc-x/new-york-city-adapts-deal-projected-increase-heat-waves

US EPA, O. (2014, June). Heat Island Reduction Solutions. Retrieved December 25, 2025, from https://www.epa.gov/heatislands/heat-island-reduction-solutions

US EPA, O. (2020, November). Local Heat Stories. Retrieved January 10, 2026, from https://www.epa.gov/heatislands/local-heat-stories

US EPA, O. (2024, August). Learn About Heat Island Effects. Retrieved December 25, 2025, from https://www.epa.gov/heatislands/learn-about-heat-island-effects

von Graevenitz, K. (2018). The amenity cost of road noise. *Journal of Environmental Economics and Management*, *90*, 1–22. https://doi.org/10.1016/j.jeem.2018.04.006

Winds of Change: The Role of Urban Expansion and Thermal Advection in Driving Phoenix's (AZ) Warming Trends - Moustaoui - 2025 - Journal of Geophysical Research: Atmospheres - Wiley Online Library. (n.d.). Retrieved December 25, 2025, from https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2024JD043166

Wing, C., Freedman, S. M., & Hollingsworth, A. (2024, January). Stacked Difference-in-Differences. https://doi.org/10.3386/w32054

Wrenn, D. H. (2024a). The Effect of Natural Disasters and Extreme Weather on Household Location Choice and Economic Welfare. *Journal of the Association of Environmental and Resource Economists*, *11*(5), 1101–1134. https://doi.org/10.1086/728887

Wrenn, D. H. (2024b). The Effect of Natural Disasters and Extreme Weather on Household Location Choice and Economic Welfare. *Journal of the Association of Environmental and Resource Economists*, *11*(5), 1101–1134. https://doi.org/10.1086/728887

Zabel, J., Phaneuf, D. J., & Krause, A. (2023). Introduction to the Special Issue Property Value Analysis using ZTRAX: Applications under the Approaching Sunset. *Journal of Housing Economics*, *62*, 101969. https://doi.org/10.1016/j.jhe.2023.101969

Zhang, Y., & Turner II, B. (2020). Land-cover mapping of the central Arizona region based on 2015 National Agriculture Imagery Program (NAIP) imagery. https://doi.org/10.6073/PASTA/E671ED549A55FDA3338B177A2AD54487

Zivin, J. G., & Neidell, M. (2014). Temperature and the Allocation of Time: Implications for Climate Change. *Journal of Labor Economics*. https://doi.org/10.1086/671766

Appendix

A Urban heat and residential electricity expenditure

A central mechanism through which households incur welfare losses from urban heat is via increased energy expenditures for residential cooling (Han et al. 2024). To illustrate this channel, and to motivate my analysis of household MWTP for urban heat reduction, I pair tract-level estimates of annual

household electricity expenditures from the Department of Energy's Low-Income Energy Affordability Dataset (LEAD, 2022) with minimum July nighttime temperatures from the NCAR high-resolution urban climate dataset.

Figure A1 presents average electricity expenditure per square foot in 2020 by decile of tract-level nighttime temperature using Low-Income Energy Affordability Data (LEAD) (DOE 2020). The relationship suggests that average tract-level annual household electricity expenditures are sensitive to changes in the urban heat island effect: households in the hottest decile face expenditures per square foot that are more than 12% higher ($0.15/sq ft) than those in the coolest decile, representing an average difference of $227 per year.To normalize electricity spending by unit size, I compute expenditures per square foot, using a tract-level median home size of 1,623 square feet derived from the Maricopa County parcel dataset. The median annual electricity expenditure across tracts is $1,714, yielding an average cost of approximately $1.06 per square foot per year. Because these data are limited to the tract level, true household savings from reduced temperature are likely higher. Nevertheless, these patterns underscore the potential for heat-induced financial strain, particularly for lower-income households disproportionately concentrated in hotter neighborhoods. This descriptive evidence provides conceptual motivation for estimating the extent to which urban heat exposure capitalizes into housing prices, which I pursue in the remainder of the paper. Importantly, the estimates I provide in this paper encapsulate benefits from cooling cost reductions as well as broader amenity effects, such as increased comfort, safety, and enjoyment of other neighborhood amenities due to a cooler microclimate.

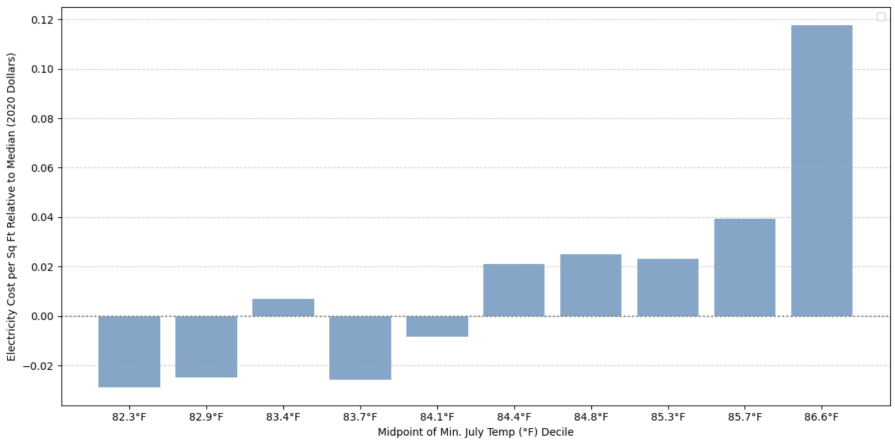


Figure A1: Average tract-level annual residential electricity costs per square foot, by tract-level urban heat decile. The x-axis labels reflect the midpoint temperature (°F) of each decile range for minimum July nighttime temperature from the NCAR urban heat database for 2018. The y-axis reflects average annual residential electricity expenditure per square foot.

B Data Appendix

B.1 Temperature data

In order to measure UHI I take advantage of high-resolution Urban Meteorology for Impacts Dataset (HUMID), developed by the National Center for Atmospheric Research (NCAR). I use gridded 1km resolution annual average July nighttime minimum near-surface (2m) temperatures as my measure for UHI, which is a common metric for assessing UHI severity. HUMID leverages the High-Resolution Land Data Assimilation System (HRLDAS), which allows temperatures to account for the impacts of spatiotemporally varying land surfaces on air temperature. The dataset is then bolstered by in-situ meteorological data (Newman et al. 2024). As a result, while temperature comparisons between HUMID and PRISM and DAYMET (two widely used gridded temperature datasets which do not explicitly account for changes in land cover and land surface temperature) show a high degree of correlation between estimates, HUMID is more accurate in measuring temperature changes across small spatial scales in urban regions (Newman et al. 2024).

While previous work at comparable spatial scales utilizes land surface temperature (LST), available at 30m resolution from satellite data (Han et al. 2024), studies in climate and urban sciences literature show that LST changes are a poor proxy for air temperature change (Goldblatt et al. 2021; Mutiibwa et al. 2015; Williamson et al. 2014), with LST varying at larger scales than air temperature, and with the relationship between LST and near-surface air temperature varying heterogeneously across land surface types (Stoll and Brazel, 1992). This threatens the validity of the scale of estimates obtained using land surface temperature, because results are likely sensitive to scale variations between land and air temperature and measurement error. On the other hand, gridded air temperature datasets such as PRISM are not explicitly urban-resolving and publicly-available versions do not reflect an appropriate spatial scale (4km) at which to evaluate urban microclimates. As a result, I argue that HUMID is best suited to investigate the impact of urban heat on home prices relative to both other air temperature datasets as well as LST satellite data.

B.2 Land cover data

To characterize vegetation and impervious surface coverage across space, I use data from the National Land Cover Database (NLCD), developed by the U.S. Geological Survey (USGS). The NLCD provides 30-meter resolution land cover classifications annually across the continental United States, including urban and vegetated land types. I extract annual classifications for Maricopa County and reclassify land cover codes to identify development intensity in each tract in the sample, as well as the presence of natural and agricultural land cover. :contentReference[oaicite:6]index=6

B.3 Demographic data

To characterize the social and economic context of neighborhoods, I assemble small-area demographic data from IPUMS National Historical Geographic Information System (NHGIS), which provides harmonized census and American Community Survey (ACS) tabulations for consistent geographic units over time (Ruggles et al. 2023). I use tract-level ACS 5-year estimates for Maricopa County over my study period to recover measures of median household income, household age, and family structure. I convert all income variables to real 2025 dollars. These demographic measures enter the analysis as covariates in my robustness checks. :contentReference[oaicite:7]index=7

B.4 Home sales data

I use home price and characteristic data from the Maricopa County Tax Assessor's Office records for all non-foreclosure, arms-length, single-family home sales transactions between 2009 and 2018. After removing properties that underwent significant changes in square footage, the addition of a pool, patio, or garage, or correspond to the bottom 5th percentile and top 95th percentile of home sales, and home sales that occurred in large, undeveloped ZIP codes greater than 30km from Phoenix's urban center, the number of home sales in the sample is 434,002. The median sale price in this sample is $191,830.6 in 2010 dollars, corresponding to an annualized sale price of $20,035. Summary statistics for observed home characteristics from this dataset are shown in Table 1. :contentReference[oaicite:8]index=8

B.5 Cool Corridors data

I obtain geospatial data on the City of Phoenix Cool Corridors program from the City of Phoenix open data portal. The 2025 Cool Corridors layer is a polyline feature class representing street segments that are completed, under construction, or planned as part of the program. Each segment includes an attribute for project status (for example, completed, in progress, or planned), which I use to distinguish between existing and future investments. I reproject the layer to a local projected coordinate system so that lengths are measured in meters, and intersect corridor segments with census tracts and parcel–centered buffers to calculate the total length of Cool Corridors within each tract and within 0–500 m and 500–1500 m rings around each parcel. These segment lengths are subsequently converted into implied numbers of street trees and associated canopy changes when I evaluate the spatial distribution of cooling benefits from the program. :contentReference[oaicite:9]index=9

B.6 Temperature data

In order to measure UHI I take advantage of high-resolution Urban Meteorology for Impacts Dataset (HUMID), developed by the National Center For for Atmospheric Research (NCAR). I use gridded 1km resolution annual average July nighttime minimum near-surface (2m) temperatures as my measure for UHI, which is a common metric for assessing UHI severity. HUMID leverages the High-Resolution Land Data Assimilation System (HRLDAS), which allows temperatures to account for the impacts of spatiotemporally varying land surfaces on air temperature. The dataset is then bolstered by in-situ meteorological data (Newman et. al 2024). As a result, while temperature comparisons between HUMID and PRISM and DAYMET (two widely used gridded temperature datasets which do not explicitly account for changes in land cover and land surface temperature) show a high degree of correlation between estimates, HUMID is more accurate in measuring temperature changes across small spatial scales in urban regions (Newman et al. 2024).

While previous work at comparable spatial scales utilizes land surface temperature (LST), available at 30m resolution from satellite data (Han et. al 2024), studies in climate and urban sciences literature show that LST changes are a poor proxy for air temperature change (Goldblatt et al. 2021, Mutiibwa et al. 2015, Williamson et. al 2014), with LST varying at larger scales than air temperature, and with

the relationship between LST and near-surface air temperature varying heterogeneously across land surface types (Stoll and Brazel, 1992). This threatens the validity of the scale of estimates obtained using land surface temperature, because results are likely sensitive to scale variations between land and air temperature and measurement error. On the other hand, gridded air temperature datasets such as PRISM are not explicitly urban-resolving and publicly-available versions do not reflect an appropriate spatial scale (4km) at which to evaulate urban microclimates. As a result, I argue that HUMID is best suited to investigate the impact of urban heat on home prices relative to both other air temperature datasets as well as LST satellite data.

B.7 Atmospheric data

To characterize vegetation and impervious surface coverage across space, I use data from the National Land Cover Database (NLCD), developed by the U.S. Geological Survey (USGS). The NLCD provides 30-meter resolution land cover classifications annually across the continental United States, including urban and vegetated land types. I extract annual classifications for Maricopa County and reclassify land cover codes to identify development intensity in each tract in the sample, as well as the presence of natural and agricultural land cover.

B.8 Demographic data

To characterize the social and economic context of neighborhoods, I assemble small-area demographic data from IPUMS National Historical Geographic Information System (NHGIS), which provides harmonized census and American Community Survey (ACS) tabulations for consistent geographic units over time.(Ruggles et al. (2023)). I use and tract-level ACS 5-year estimates for Maricopa County over my study period to recover measures of median household income, household age, and family structure. I convert all income variables to real 2025 dollars. These demographic measures enter the analysis as covariates in my robustness checks

B.9 Home sales data

I use home price and characteristic data from from the Maricopa County Tax assessors office records for all non-forclosure, arms length, single family home sales transactions between 2009 and 2018. After

removing properties that underwent significant changes in square footage, the addition of a pool, patio, or garage, or correspond to the bottom 5th percentile and top 95th percentile of home sales, and home sales that occurred in large, undeveloped zips greater than 30km from Phoenix's Urban center, the number of home sales in the sample is 434,002. The median sale price in this sample is $191,830.6 in 2010 dollars, corresponding to an annualized sale price of $20,035. Summary statistics for observed home characteristics from the this dataset are shown in Table 1.

B.10 Land cover data

To characterize vegetation and impervious surface coverage across space, I use data from the National Land Cover Database (NLCD), developed by the U.S. Geological Survey (USGS). The NLCD provides 30-meter resolution land cover classifications annually across the continental United States, including urban and vegetated land types. I extract annual classifications for Maricopa County and reclassify land cover codes to identify development intensity in each tract in the sample, as well as the presence of natural and agricultural land cover.

B.11 Cool Corridors data

I obtain geospatial data on the City of Phoenix Cool Corridors program from the City of Phoenix open data portal. The 2025 Cool Corridors layer is a polyline feature class representing street segments that are completed, under construction, or planned as part of the program. Each segment includes an attribute for project status (for example, completed, in progress, or planned), which I use to distinguish between existing and future investments. I reproject the layer to a local projected coordinate system so that lengths are measured in meters, and intersect corridor segments with census tracts and parcel–centered buffers to calculate the total length of Cool Corridors within each tract and within 0–500 m and 500–1500 m rings around each parcel. These segment lengths are subsequently converted into implied numbers of street trees and associated canopy changes when I evaluate the spatial distribution of cooling benefits from the program.

C Construction of the gradient alignment instrument

C.1 Gradient alignment score calculation

To determine how much each property's outlying buffer zone transfers cool or hot air to its immediate surroundings, I define a 500 m buffer around each parcel as its "immediate" neighborhood and construct an outlying ring that extends 500 m beyond this buffer (i.e., from 500 m to 1 km). I divide the outlying ring into eight directional slices, evenly spaced at 45 degrees (N, NE, E, SE, S, SW, W, NW). For each parcel i, year t, and direction d, I compute the July nighttime temperature in slice d and form a temperature anomaly

$$\alpha_{itd} = \frac{T_{itd} - \bar{T}_{it}^{\text{ring}}}{\max_{d'} T_{itd'} - \min_{d'} T_{itd'}},$$

where T_{itd} is the slice temperature and $\bar{T}_{it}^{\text{ring}}$ is the mean nighttime temperature in the outlying ring. The anomaly α_{itd} measures how much hotter or cooler slice d is relative to the ring average, normalized by the range of temperatures in the ring. Positive values indicate relatively hot slices, negative values indicate relatively cool slices, and values near zero indicate slices close to the ring mean. These anomalies can be interpreted as the relative "heating" or "cooling" potential of each directional slice if air flows from that slice into the inner 500 m buffer.

Figure 4 illustrates the logic in a simple case with a single air parcel trajectory arriving from the north. The innermost circle represents the parcel's 500 m immediate surroundings, while the outer circle represents its outlying ring. Numbers in each slice represent deviations from the ring mean temperature. In the left panel, prevailing flow from the north aligns with a relatively hot slice, implying net advective warming; in the right panel, the same northerly flow aligns with a relatively cool slice, implying net advective cooling, even though the mean temperature of the outer ring is the same in both cases.

Formally, let θ_{it} denote the mean bearing of nighttime air parcel trajectories arriving at parcel i in year t, and let θ_d denote the central angle of directional slice d. Let ω_{id} be the frequency weight for parcel i associated with trajectories from direction d, and let $\mathcal{D}$ be the set of eight directional slices. I define an angular alignment weight

$$A_{itd} = 1 - \frac{\min\left(|\theta_d - \theta_{it}|,\, 360^\circ - |\theta_d - \theta_{it}|\right)}{180^\circ},$$

which attains a maximum value of 1 when the airflow bearing perfectly aligns with slice d ($\theta_{it} = \theta_d$)

and diminishes linearly to 0 as the angular difference approaches 180°. The gradient alignment score for parcel i in year t is then

$$\Gamma_{it} = \sum_{d \in \mathcal{D}} \omega_{id} \, A_{itd} \, \alpha_{itd}. \tag{8}$$

The score Γ_{it} can be interpreted as a weighted average across all directional slices, where for each direction the temperature anomaly α_{itd} is multiplied by (i) how frequently winds arrive from that direction (ω_{id}) and (ii) how closely aligned the wind bearing is with that slice (A_{itd}). Parcels for which winds frequently arrive from relatively hot slices have large positive scores; parcels for which winds arrive from relatively cool slices have large negative scores; parcels for which winds "see" hot and cool slices symmetrically have scores near zero.

C.2 HYSPLIT air parcel trajectories

To assign frequency-weighted air parcel trajectory flows to each 500 m property buffer, I use air parcel trajectory simulations from NOAA's HYSPLIT model (Hybrid Single-Particle Lagrangian Integrated Trajectory). HYSPLIT models the movement of near-surface air parcels over time by simulating backward trajectories that trace air parcels back in time from a receptor location (here, the parcel centroid) to identify upwind source regions. The model integrates wind fields and meteorological conditions to estimate parcel paths at designated atmospheric elevations.

Trajectories are driven using the High-Resolution Rapid Refresh (HRRR) dataset, which provides hourly, 3 km resolution atmospheric data for the continental U.S. HRRR assimilates radar, satellite, and surface observations to generate high-frequency snapshots of atmospheric dynamics, making it well suited for modeling local airflows in urban environments.

Backward air parcel trajectories are produced by solving the Lagrangian transport equation, which describes the evolution of an air parcel's position under the influence of the local wind field. Starting from a known receptor location $\mathbf{x}_0$ at time t_0, the model integrates

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{u}(\mathbf{x}(t), t),$$

where $\mathbf{x}(t)$ is the air parcel position at time t and $\mathbf{u}(\mathbf{x}(t), t)$ is the wind velocity vector at that location and time. Using a numerical integration scheme (e.g., an Euler backward step over a small time interval

Δt),

$$\mathbf{x}(t - \Delta t) \approx \mathbf{x}(t) - \mathbf{u}(\mathbf{x}(t), t)\, \Delta t,$$

HYSPLIT iteratively traces each parcel's path backward in time to estimate its origin.

Due to the computational intensity and large file sizes associated with HRRR data, I generate trajectories for a subset of five evenly spaced days in July for each year between 2015 and 2018, the period over which HRRR overlaps with my study window. For each day, I simulate backward trajectories from each parcel centroid for representative nighttime hours (6 pm, 12 am, and 6 am), and record the bearing from which air parcels arriving one hour earlier originated. For each parcel, I then compute the frequency with which trajectories pass through its 500 m buffer from each of the eight compass directions, aggregating across sampled days, hours, and years to obtain a period-average distribution of nighttime airflow directions.

I assume that this period-average pattern of nighttime airflow is representative for earlier years in the sample and apply the same directional frequencies to pre-2015 sales. Figure A2 displays the distribution of frequency-weighted air parcel trajectory bearings across the study area, and Figure A3 displays the mean gradient alignment score across census blocks, computed by averaging equation (8) over all sample years. Positive values indicate a high degree of alignment between warmer outlying slices and prevailing air parcel trajectories, while negative values reflect alignment with cooler slices. As expected, there is substantial variation in the gradient alignment score within census tracts. By differencing off the local temperature gradient and interacting it with parcel-specific airflow orientation, the instrument removes much of the broad spatial correlation in the airflow field and yields a source of variation in neighborhood temperature that is unlikely to be correlated with other spatially correlated unobservables such as air quality, other climatic factors, or flood risk.

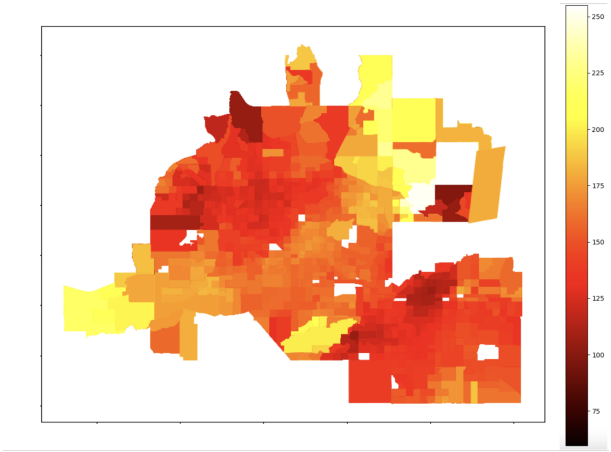


Figure A2: Frequency-weighted nighttime air parcel trajectory bearings relative to property centroids by census tract.

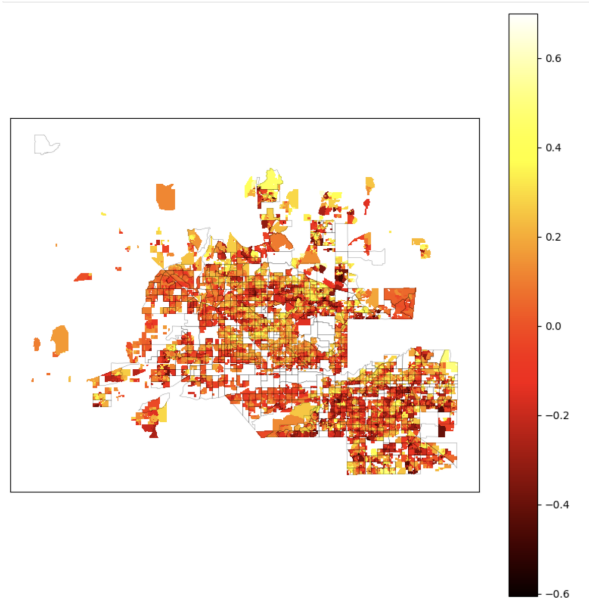


Figure A3: Mean gradient alignment score across census blocks, 2009–2018.

D Additional Econometric Specifications

D.1 Hausman–Taylor replication

This appendix section summarizes the implementation of the Hausman–Taylor (HT) model used to benchmark my instrumented estimates, corresponding to equation (12) in the main text.

As in Klaiber, Abbott, and Smith (2017), I recast the cross–section of sales as a panel of tract–year cells. Let j index census tracts, t years, and i transactions within a tract–year cell. The semi–log

hedonic price function is

$$\ln p_{ijt}^{ann} = \delta_0 + \delta_T T_{jt}^{UHI} + \mathbf{X}_{ijt}^{1\prime}\boldsymbol{\delta}_1 + \mathbf{X}_{ijt}^{2\prime}\boldsymbol{\delta}_2 + \mathbf{Z}_{jt}^{1\prime}\boldsymbol{\delta}_3 + \mathbf{Z}_{jt}^{2\prime}\boldsymbol{\delta}_4 + h_{jt} + \varepsilon_{ijt}, \tag{9}$$

where T_{jt}^{UHI} is tract–level mean July nighttime temperature, $\mathbf{X}_{ijt}^{1}$ collects parcel–level regressors that are potentially correlated with the tract–year random effect h_{jt} (e.g., lot size and parcel greenspace), and $\mathbf{X}_{ijt}^{2}$ collects parcel–level regressors that are assumed exogenous (e.g., structural characteristics). The vectors $\mathbf{Z}_{jt}^{1}$ and $\mathbf{Z}_{jt}^{2}$ similarly partition tract–level regressors into endogenous and exogenous components (e.g., land use shares versus exogenous demographic controls).

The key identifying assumption of the HT estimator is that any correlation between $\mathbf{X}_{ijt}^{1}$, $\mathbf{Z}_{jt}^{1}$ and h_{jt} can be captured by tract–year means of exogenous regressors. Instruments are constructed using (i) within–tract deviations of all regressors that vary at the parcel level and (ii) tract–year means of exogenous regressors, following Abbott and Klaiber (2011). I apply the same partition of regressors as in Klaiber, Abbott, and Smith (2017) and cluster standard errors at the tract level.

D.2 Repeat–sales hedonic model

The repeat–sales sample is constructed by restricting the data to parcels that transact at least twice during the study period and dropping any sales with missing key covariates. For parcel i in year t, the repeat–sales model is

$$\ln p_{it}^{ann} = \lambda_0 + \lambda_T T_{it}^{UHI_{500}} + \mathbf{z}_{it}^{\prime}\boldsymbol{\lambda}_z + \mu_i + \tau_t + \varepsilon_{it}, \tag{10}$$

where μ_i is a parcel fixed effect and τ_t is a full set of year dummies. The vector $\mathbf{z}_{it}$ includes time–varying parcel characteristics (e.g., renovations, changes in assessed square footage when observed) and time–varying neighborhood controls. Estimation proceeds by within–parcel demeaning. Because μ_i absorbs all time–invariant parcel heterogeneity, identification comes solely from changes in temperature and controls between successive sales for the same parcel.

D.3 Spatial lag model and benefit aggregation

This section provides additional details on the spatial lag specification in equation (5).

Ring construction. For each parcel i and year t, I construct a series of concentric annuli centered on the parcel centroid with outer radii 500, 1000, 2000, and 3000 meters. Let $\mathcal{D} = \{1, 2, 3, 4\}$ index the distance bands corresponding to 0–500 m, 500–1000 m, 1000–2000 m, and 2000–3000 m, respectively. I then compute the mean July nighttime temperature in each band, denoted T_{itd}^{UHI}, using the high–resolution temperature rasters described in the data section.

First stage for spatial lags. For each band $d \in \mathcal{D}$, I construct a distance–specific gradient alignment score Γ_{itd} using the same routine as in Appendix C , but applied to an "outlying" 100m buffer ring surrounding band d. The first–stage equations are

$$T_{itd}^{UHI} = a_{0d} + a_{\Gamma d}\,\Gamma_{itd} + \mathbf{z}_i'\boldsymbol{a}_{zd} + \mathbf{z}_{tract,t}'\boldsymbol{a}_{tract,d} + \theta_{zip}^{iv} + \phi_{puma(i)\times t}^{iv} + \nu_{itd}, \qquad d \in \mathcal{D}, \tag{11}$$

with the same set of controls and fixed effects used in the baseline 500m specification. Predicted temperatures $\widehat{T}_{itd}^{UHI}$ from these first stages enter the second–stage spatial lag model in equation (5).

E Alternative panel specifications

E.1 Hausman–Taylor replication

To benchmark my IV estimates against the quasi–experimental approach in Klaiber, Abbott, and Smith (2017), I implement a Hausman–Taylor (HT) model that treats tract–level temperature as an endogenous regressor in a panel of tract–year cells. Let j index tracts and t years, and let i index houses transacting within a tract–year cell. The semi–log hedonic model can be written as

$$\ln p_{ijt}^{ann} = \delta_0 + \delta_T\, T_{jt}^{UHI} + \mathbf{X}_{ijt}^{1\prime}\boldsymbol{\delta}_1 + \mathbf{X}_{ijt}^{2\prime}\boldsymbol{\delta}_2 + \mathbf{Z}_{jt}^{1\prime}\boldsymbol{\delta}_3 + \mathbf{Z}_{jt}^{2\prime}\boldsymbol{\delta}_4 + h_{jt} + \varepsilon_{ijt}, \tag{12}$$

where T_{jt}^{UHI} is tract–level temperature, $\mathbf{X}_{ijt}^{1}$ and $\mathbf{X}_{ijt}^{2}$ are parcel–level characteristics classified as endogenous and exogenous, respectively, and $\mathbf{Z}_{jt}^{1}$ and $\mathbf{Z}_{jt}^{2}$ are tract–level variables that are endogenous and exogenous. The tract–year random effect h_{jt} captures unobserved neighborhood–time heterogeneity. The HT estimator uses within–tract deviations and tract–level means of exogenous regressors as instruments for the endogenous components. Appendix D.1 provides the precise partition of regressors and instrument set used in my replication of Klaiber, Abbott, and Smith (2017).

E.2 Repeat–sales specification

As an additional comparison, I estimate a repeat–sales model that differences out time–invariant parcel heterogeneity. Restricting the sample to parcels that sell at least twice, I estimate

$$\ln p_{it}^{ann} = \lambda_0 + \lambda_T T_{it}^{UHI_{500}} + \mathbf{z}_{it}'\boldsymbol{\lambda}_z + \mu_i + \tau_t + \varepsilon_{it}, \tag{13}$$

where μ_i is a parcel fixed effect and τ_t are year fixed effects. Because μ_i absorbs all time–invariant parcel–level attributes (including permanent differences in lot size, structure quality, or unobserved local amenities), identification comes from within–parcel changes in temperature and controls between sales. In practice, I report the repeat–sales estimates as a comparison to my baseline IV specification rather than as my main source of MWTP estimates; additional details on sample construction and weighting are provided in Appendix D.2.

E.3 Tract fixed–effects model

Finally, I estimate a tract fixed–effects model as a simple panel demonstration of how within–tract variation in temperature is related to housing prices. Letting j index tracts and t years as above, I estimate

$$\ln p_{ijt}^{ann} = \kappa_0 + \kappa_T T_{jt}^{UHI} + \mathbf{x}_{ijt}'\boldsymbol{\kappa}_x + \gamma_j + \tau_t + \varepsilon_{ijt}, \tag{14}$$

where T_{jt}^{UHI} denotes tract–level temperature (measured as the tract–level mean of my 500 m UHI field), $\mathbf{x}_{ijt}$ collects parcel–level controls, γ_j are tract fixed effects, and τ_t are year fixed effects. The tract fixed effects γ_j absorb all time–invariant neighborhood characteristics, so κ_T is identified from within–tract changes in temperature and controls over time. I treat this specification as a descriptive illustration of how much of the cross–sectional correlation between prices and heat remains after differencing out permanent tract–level heterogeneity, rather than as a fully credible causal estimate, given that T_{jt}^{UHI} may still be correlated with time–varying unobservables that are not captured by τ_t or $\mathbf{x}_{ijt}$.

F Cool Corridor DiD Analysis

This appendix documents the stacked difference-in-differences(DiD) design I use to estimate the causal effect of Cool Corridors installations on July median land surface temperature (LST) around corridor

segments.

F.1 Difference-in-Differences model

The analysis unit is a corridor segment i observed annually in calendar year t. The outcome is Y_{it}, defined as the July median LST (in °C) within a 100-meter buffer around each corridor segment. A buffer of 100m is chosen to reflect the highly localized cooling effect of street level tree planting recorded for small compact urban vegetation spaces Soltanifard and Amani-Beni (2025). For each year, corridor-level LST is computed by overlaying corridor buffers (100m) on annual LST rasters and taking the mean temperature. Phoenix's Cool corridors program began installations in 2022 and continued up until present (2025).The LST panel covers calendar years 2018–2025.

Treated Corridors with an installed tree-density >= 62 trees/km (>= 50% of Phoenix's stated goal of 200 trees per mile, (>= 124 trees/km)) are assigned a first-treatment year based on year of recorded installation; $G_i \in \{2022, 2023, 2024\}$, $n = 13$. Corridors with sparse tree installation are removed from the analysis <= 62 trees/km. Corridors in the sample that are scheduled/planned and not yet installed comprise the "never-treated" control group, $n = 21$. The event time for corridor i in year t is

$$\text{rel}_{it} \equiv t - G_i.$$

F.2 Motivation for stacked DiD

Standard two-way fixed effects (TWFE) DiD models can produce "bad comparisons" in staggered adoption settings when already-treated units serve as controls for later-treated units. To avoid these comparisons, I implement a stacked DiD design that constructs separate cohort-specific sub-experiments (one per adoption year) and then stacks them into a single estimation sample. By doing this I ensure that, within each cohort-specific sub-experiment, treated units are compared only to clean controls that are never treated not-yet treated.

F.3 Cohort-specific sub-experiments and trimming

For each adoption cohort, $g \in \{2022, 2023, 2024\}$, I construct a sub-experiment $s = g$ that includes the land within a 100m buffer surrounding treated corridors with $G_i = g$, and land within 100m surrounding clean controls, i.e. corridors with $G_i > g + \kappa_{\text{post}}$ or never treated corridors that remain in

"scheduled" or "planned" status as of 2025. I then trim the calendar time dimension to an event window $\text{rel}_{it} \in [-\kappa_{\text{pre}}, \kappa_{\text{post}}]$, using $\kappa_{\text{pre}} = 3$ and $\kappa_{\text{post}} = 3$. This produces a balanced event-time window within each cohort-specific sub-experiment. (In practice, estimates for the farthest post period rel = +3 is only contributed by the earliest cohort ($g = 2022$) because later cohorts do not have enough post years within the data.). All sub-experiments are then appended into one input dataset for DiD estimation.

F.4 DiD specification

The stacked DiD event-study model is:

$$Y_{its} = \alpha_{is} + \gamma_{ts} + \sum_{k \in \mathcal{K}, k \neq -1} \beta_k \, \mathbf{1}\{\text{rel}_{its} = k\} \times \text{Treat}_{is} + \varepsilon_{its},$$

where $\mathcal{K} = \{-3, -2, -1, 0, 1, 2, 3\}$, $k = -1$ is omitted and serves as the reference period. Each β_k represents the difference between treated and control groups at event time k, relative to the treated-minus-control difference at event time -1.

Let s index the sub-experiment (cohort) and denote the stacked outcome as Y_{its}. I define Treat_{is} as an indicator equal to 1 if corridor i is treated in sub-experiment s (i.e. $G_i = s$), and 0 if corridor i is a clean control in sub-experiment s, $\text{rel}_{its} = t - s$ represents event time measured relative to the sub-experiment cohort year s.

α_{is} represents unit by sub-experiment fixed effects, which absorb all time-invariant differences for corridor i within each sub-experiment s. γ_{ts}, represents all shocks common to all corridors in calendar year t within each sub-experiment s.

F.5 Corrective stacking weights (Wing et al. 2024)

Stacking changes the relative contribution of control trends in the sample. For example, later cohorts can appear more often as controls in earlier treated cohorts. To recover an ATT that is not mechanically driven by how many controls appear for a given cohort. I apply corrective weights following the methods outlined in Wing et al. (2024).

For each event time k, I define N_{sk}^{treat} and N_{sk}^{ctrl} as the number of treated and control observations in sub-experiment s at event time k, and define N_{k}^{treat} and N_{k}^{ctrl} as the totals across all sub-experiments at event time k. I define shares:

$$\pi_{sk}^{\text{treat}} = \frac{N_{sk}^{\text{treat}}}{N_k^{\text{treat}}}, \qquad \pi_{sk}^{\text{ctrl}} = \frac{N_{sk}^{\text{ctrl}}}{N_k^{\text{ctrl}}}.$$

Weights are assigned as:

$$w_{its} = \begin{cases} 1 & \text{if Treat}_{is} = 1, \\ \frac{\pi_{sk}^{\text{treat}}}{\pi_{sk}^{\text{ctrl}}} & \text{if Treat}_{is} = 0, \end{cases}$$

so that, within each event time k, the weighted control group matches the sub-experiment composition of the treated group. Estimation is performed using weighted least squares with the defined weights. ATT estimates relative to $k = -1$ are reported in the below event study plot.

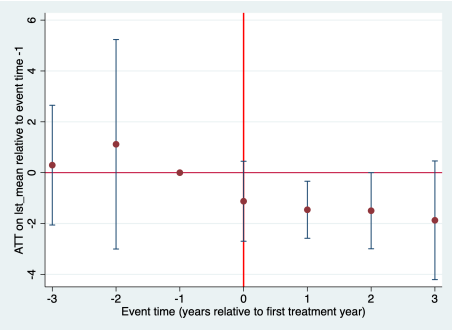


Figure A4: ATT by event time for estimates relative to $k = -1$, plotted with 95% confidence bands

F.6 Re-normalization to address raster-level differences

Because year-to-year raster products can differ in level (e.g., processing changes), I also report a re-centered event-study series where the average of pre-treatment coefficients is normalized to zero:

$$\tilde{\beta}_k \equiv \beta_k - \frac{1}{\mathcal{K}_{\text{pre}}} \sum_{j \in \mathcal{K}_{\text{pre}}} \beta_j,$$

with $\mathcal{K}_{\text{pre}} = \{-3, -2, -1\}$ (and $\beta_{-1} = 0$ by construction under the omitted-category normalization). Standard errors for $\tilde{\beta}_k$ are computed using linear combinations of estimated coefficients.

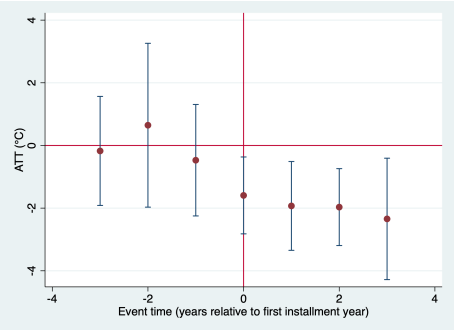


Figure A5: ATT by event time for estimates relative to the normalized mean of treatment periods , plotted with 95% confidence bands, estimates reflect heteroskedasticity-robust standard errors clustered at the corridor level.

Table 11: Event-study coefficients centered at mean of pre-periods

	(1) ATT of Cool Corridor on LST
Year Relative to Cool Corridor Installation	
Event time -3	-0.175
	(0.867)
Event time -2	0.646
	(1.266)
Event time -1	-0.471
	(0.863)
Event time 0	-1.594***
	(0.589)
Event time +1	-1.930***
	(0.666)
Event time +2	-1.966***
	(0.586)
Event time +3	-2.348**
	(0.950)
Observations	394.000

Standard errors in parentheses

Each coefficient is (treated-control at event time k) minus mean(pre) where pre = -3,-2,-1.

,

Standard errors are cluster-robust and computed using the regression VCV via nlcom (exact for linear combinations).

,

Fixed effects absorbed: unit_stack year_stack.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

F.7 Parallel trends diagnostics and interpretation

I plot mean LST for treated and control corridors by event time within the stacked sample over rel $\in [-3, 3]$. Results show that control and treatment corridors remained very close in LST level prior to treatment, with the first divergence recorded at $k = 0$

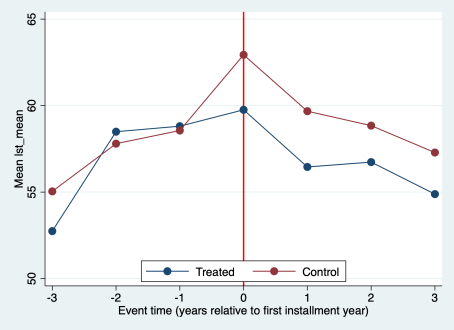


Figure A6: ATT by event time for estimates relative to the normalized mean of treatment periods , plotted with 95% confidence bands