

Equity-Weighted Allocation of Cooling and Weatherization Resources Across Census Tracts in Richmond, Virginia: Integrating Community Heat-Mapping and Modeled Energy Burden Under Uncertainty

Abstract

Extreme heat is the deadliest weather-related hazard in the United States, yet its burden is distributed unevenly within cities and is compounded by high household energy costs that restrict access to protective cooling. We develop and demonstrate a transparent, reproducible framework for allocating a fixed pool of 1,000 cooling and weatherization program slots across the 126 populated 2022 census tracts intersecting the Richmond, Virginia heat-mapping footprint. We summarize afternoon air temperature, afternoon heat index, and evening air temperature from the 2021 National Integrated Heat Health Information System (NIHHIS)–CAPA Strategies community-science campaign by tract using zonal statistics, and we combine these with counts of low-income (0–80% of Area Median Income, AMI) households and modeled energy burden from the U.S. Department of Energy Low-Income Energy Affordability Data (LEAD) tool. Three priority indices—Heat-Only, Energy-Only, and a balanced Joint-Priority index—drive a constrained proportional allocation algorithm that distributes exactly 1,000 slots among selected tracts subject to a floor of 25 and a ceiling of 200 slots per tract. We propagate zonal sampling uncertainty and LEAD estimation uncertainty through a 1,000-run Monte Carlo simulation. The three strategies select overlapping but materially different tract portfolios: Heat-Only and Joint-Priority allocations are strongly concordant (Pearson $r = 0.89$; selection Jaccard = 0.83), whereas Heat-Only and Energy-Only diverge most (Pearson $r = 0.61$; Jaccard = 0.47). The Joint-Priority strategy selects 27 tracts (20 in the City of Richmond and 7 in Henrico County), concentrating resources in a stable inner-city core while capturing roughly 36% of the study area’s low-income households. Monte Carlo analysis shows that only a small core of tracts is robustly selected across all trials, while many mid-ranked tracts are conditionally selected, underscoring the value of reporting selection probabilities alongside point allocations. We deliver tract GEOIDs, final slot counts, three allocation maps, and a full sensitivity analysis. We stress that the heat surface is a campaign-day snapshot of ambient conditions and not a measure of household-level risk, and we discuss the policy trade-offs between targeting heat exposure and targeting energy burden.




Figure 1: Graphical abstract. Workflow linking the 2021 NIHHS-CAPA Richmond community heat-mapping rasters and the 2022 DOE LEAD energy-burden dataset through tract-level zonal statistics and Monte Carlo uncertainty propagation to three priority indices (Heat-Only, Energy-Only, Joint-Priority) and a constrained allocation of 1,000 program slots.

1 Introduction

Extreme heat is now the leading cause of weather-related mortality in the United States and much of the world, and its toll is rising as the climate warms [1–3]. A multicountry analysis attributed roughly 0.9% of all deaths in the study populations to heat, and the warm-season mortality burden of recent decades is increasingly traceable to anthropogenic climate change [2, 4]. Country-level temperature–mortality damage functions further indicate that this burden will grow substantially under continued warming absent adaptation [5]. The danger is not distributed evenly across the calendar or the map. Heat waves kill through cumulative physiological strain, and elevated nighttime and evening temperatures are especially hazardous because they deny the body the overnight recovery on which thermoregulation depends [3, 6]. The canonical epidemiology of the 1995 Chicago heat wave established that mortality clustered among socially isolated, chronically ill, and low-income residents, and that access to a working air conditioner reduced the odds of death by approximately 80% [7, 8]. Air conditioning and building weatherization are therefore not luxuries but life-saving public-health infrastructure, and their uneven availability is a central axis of heat inequity.

The physical driver of intra-urban heat disparities is the urban heat island (UHI), the tendency of built-up areas to retain and re-radiate heat because of their surface materials, canyon geometry, reduced vegetation, and anthropogenic heat release [9]. Crucially for equity, the UHI is not uniform across a city. Within a single metropolitan area, near-surface air temperatures can vary by several degrees Celsius over short distances as a function of tree canopy, impervious surface fraction, and land-cover configuration [10–12]. The standardization of intra-urban thermal contrasts through

the Local Climate Zone framework has made these differences measurable and comparable across cities [12, 13]. A growing body of evidence demonstrates that these fine-scale thermal patterns track historical and contemporary patterns of disadvantage. Neighborhoods subjected to mid-twentieth-century mortgage “redlining” by the Home Owners’ Loan Corporation are today systematically hotter—by as much as 2–7°C of land-surface temperature—than their non-redlined counterparts, a legacy of decades of disinvestment in trees, parks, and cooling infrastructure [14, 15]. Nationally, low-income residents and communities of color are exposed to disproportionately intense surface UHIs [16], and neighborhood microclimate has been shown to interact with social vulnerability to elevate heat-stress risk [17, 18]. Richmond, Virginia is a paradigmatic case: it was among the cities in which the redlining–heat association was first quantified, and its stark thermal gradients motivated one of the earliest and most influential community heat-mapping campaigns [14].

Community heat-mapping campaigns coordinated through the National Integrated Heat Health Information System (NIHHIS) and its implementing partner CAPA Strategies have transformed the empirical basis for local heat policy. Trained volunteers drive pre-planned routes at fixed morning, afternoon, and evening windows on a hot day, logging ambient temperature and humidity with GPS-referenced sensors; machine-learning models then interpolate the traverse observations, together with satellite and land-cover predictors, into wall-to-wall air-temperature and heat-index surfaces at fine spatial resolution [19, 20]. These campaign products capture the ambient thermal environment far more faithfully than satellite land-surface temperature alone, which systematically overstates daytime and understates canopy-layer nighttime heat [10]. It is essential, however, to interpret these surfaces correctly: they represent the ambient conditions along the routes on the specific campaign day, not the heat actually experienced inside any dwelling, and certainly not household-level health risk, which additionally depends on building envelope, air-conditioning access, occupant health, and behavior [19]. Throughout this study we treat the Richmond heat surface as a campaign-day exposure snapshot and not as a household risk metric.

The second dimension of heat vulnerability is economic. Even where cooling equipment is physically present, high energy costs cause households to ration or forgo it. Energy burden—the share of household income spent on home energy—is the standard metric of this stress, with 6% and 10% commonly used as thresholds for “high” and “severe” burden [21, 22]. Low-income households, and disproportionately Black, Hispanic, and Native American households, face energy burdens several times the median, a disparity that persists after controlling for income [22–24]. Energy insecurity—the inability to meet basic household energy needs—is now recognized as a distinct social and environmental determinant of health, associated with adverse respiratory, cardiovascular, and mental-health outcomes and with unsafe coping behaviors such as leaving a home dangerously hot [23, 25]. Weatherization—air sealing, insulation, and efficiency upgrades—directly reduces energy burden and improves thermal safety, though program evaluations caution that realized savings can fall short of engineering projections and that program reach is uneven [26]. The U.S. Department of Energy’s Low-Income Energy Affordability Data (LEAD) tool provides the most granular publicly available estimates of household energy expenditure, income, and burden, disaggregated by AMI bracket, tenure, building vintage, and fuel at the census-tract scale [27].

Heat exposure and energy burden are conceptually distinct and only partially co-located. A tract may sit in a scorching, low-canopy corridor yet house relatively affluent residents, or it may be comparatively cool yet contain many cost-burdened low-income families. Programs that deliver both cooling assistance (for example, efficient air conditioners, heat-pump installation, or cooling-center access) and weatherization must therefore decide how to weight physical heat against economic vulnerability when targeting finite resources. This is fundamentally a problem of distributional environmental justice [28, 29]. The field has produced a rich toolkit of screening instruments—

CalEnviroScreen, EJSCREEN, and the federal Justice40 framework among them—that combine environmental and social indicators into composite disadvantage scores to guide investment [28]. Yet composite indices embed consequential and often opaque choices about normalization, weighting, and aggregation, and their rankings can be sensitive to those choices and to the spatial units used [29–31]. In the heat-adaptation domain specifically, analyses of cooling-center siting reveal that access is frequently misaligned with the populations at highest risk, motivating explicit, auditable allocation methods [32].

Two methodological hazards attend any tract-level allocation built on raster and modeled inputs. The first is the modifiable areal unit problem (MAUP): statistics aggregated to census tracts depend on the size and shape of those tracts, so associations and rankings can shift under alternative zonations [33–35]. The second is input uncertainty: zonal summaries of a heat raster carry sampling error, and LEAD figures are model estimates rather than measurements. Point-estimate rankings can convey a false precision that obscures how sensitive the resulting allocation is to these uncertainties. Best practice in composite-indicator construction is therefore to accompany rankings with formal uncertainty and sensitivity analysis, typically via Monte Carlo simulation [30, 31, 36].

In this paper we integrate these threads into a single, transparent, and reproducible allocation framework for Richmond, Virginia, and we use it to interrogate the policy trade-offs among competing definitions of priority. Our specific objectives are to: (i) summarize the 2021 NIHHS–CAPA afternoon temperature, afternoon heat index, and evening temperature surfaces at the 2022 census-tract scale by zonal statistics; (ii) integrate tract-level counts of 0–80% AMI households and modeled low-income energy burden from the 2022 DOE LEAD Virginia dataset; (iii) construct Heat-Only, Energy-Only, and Joint-Priority indices and use them to drive a constrained proportional allocation of exactly 1,000 program slots with a per-tract floor of 25 and ceiling of 200; (iv) propagate zonal and LEAD uncertainty through a 1,000-run Monte Carlo simulation to quantify allocation stability and selection probability; and (v) compare the three strategies, delivering tract GEOIDs, final slot counts, three allocation maps, and sensitivity results. Our contribution is not a claim that any single weighting is optimal, but rather a defensible, auditable procedure that makes the consequences of the heat-versus-burden trade-off explicit and that reports the uncertainty in its own recommendations.

2 Methods

2.1 Study area and data sources

The study area comprises the 2022 vintage U.S. Census Bureau TIGER/Line census tracts that intersect the spatial footprint of the 2021 Richmond community heat-mapping campaign [37, 38]. After spatial filtering and removal of tracts with no valid heat-raster coverage, 127 tracts remained; one of these (GEOID 51087980100, a water/institutional tract) contains no LEAD household records and was carried through the spatial layers but excluded from the allocation pool, leaving 126 populated tracts eligible to receive slots. The pool spans three jurisdictions: the independent City of Richmond (75 tracts), Henrico County (41 tracts), and Chesterfield County (10 tracts).

Three primary datasets were used. First, the NIHHS–CAPA Richmond 2021 heat rasters (Open Science Framework project `3xvmg`, archive file `kw5bc`) comprise six single-band GeoTIFFs on a shared 10 m grid in the UTM Zone 18N projection (EPSG:32618), giving morning, afternoon, and evening surfaces of both air temperature and heat index in degrees Fahrenheit [38]. Consistent with the campaign design, we used the afternoon air temperature (`af_t`), afternoon heat index (`af_hi`),

and evening air temperature (`pm_t`) layers. Second, the 2022 DOE LEAD Virginia dataset (Open Energy Data Initiative submission 6219; DOI 10.25984/2504170) provides modeled household counts and energy expenditures disaggregated by AMI bracket, tenure, building vintage, and heating fuel [27]. Third, the 2022 TIGER/Line Virginia tract shapefile supplied the tract geometries and 11-digit GEOID keys [37].

2.2 Spatial processing and zonal statistics

Tract geometries were reprojected from their native NAD83 geographic coordinates (EPSG:4269) to UTM Zone 18N (EPSG:32618) to match the rasters and to permit metrically accurate area and distance operations. For each tract and each of the three heat layers we computed zonal statistics—the mean and standard deviation of all valid (non-nodata) raster pixels falling within the tract polygon—using a per-tract raster window and a rasterized geometry mask. Tracts with no valid pixels across all layers were dropped. The resulting tract means span 89.5–92.5°F for afternoon air temperature (study-area mean 90.8°F), 91.9–98.5°F for afternoon heat index (mean 95.7°F), and 85.7–90.0°F for evening air temperature (mean 87.9°F). The per-tract pixel standard deviations were retained to characterize within-tract thermal heterogeneity and to parameterize the uncertainty analysis. We emphasize again that these summaries describe ambient campaign-day conditions and not household-level exposure or risk.

2.3 Socioeconomic integration

LEAD records were aggregated to the tract level using the 11-digit GEOID (LEAD field `FIP`). The low-income population was defined as households in the 0–80% AMI band, formed by summing the 0–30%, 30–60%, and 60–80% AMI brackets. Tract-level energy burden was computed using the unbiased ratio-of-sums estimator rather than the mean of household-level ratios, i.e.

$$\text{Energy burden} = \frac{\sum_i (\text{ELEP}_i + \text{GASP}_i + \text{FULP}_i) \cdot \text{UNITS}_i}{\sum_i \text{HINCP}_i \cdot \text{UNITS}_i}, \quad (1)$$

where the sums run over LEAD rows i within a tract, `ELEP`, `GASP`, and `FULP` are annual electricity, gas, and other-fuel costs, `HINCP` is household income, and `UNITS` is the modeled household count. This estimator, computed overall and for the 0–80% AMI subset, correctly weights households by number and avoids the small-denominator instability of averaging ratios; the statewide implied burden of 1.94% matched an independent sanity check. Among the 126 populated tracts, modeled low-income energy burden ranges from 3.2% to 9.8% (mean 6.2%), and 0–80% AMI household counts range from 49 to 1,648 per tract, totaling 89,726 low-income households across the study area. The socioeconomic table was left-joined to the tract geometries, preserving all boundaries.

2.4 Priority indices

Five tract-level variables—afternoon air temperature, afternoon heat index, evening air temperature, low-income energy burden, and low-income household count—were each min–max normalized to the interval $[0, 1]$ across the 126-tract pool:

$$\tilde{x}_i = \frac{x_i - \min_j x_j}{\max_j x_j - \min_j x_j}. \quad (2)$$

From these we defined three composite priority indices. The Heat-Only index is the unweighted mean of the three normalized thermal variables,

$$I_i^{\text{heat}} = \frac{1}{3} \left(\widetilde{\text{af_t}_i} + \widetilde{\text{af_hi}_i} + \widetilde{\text{pm_t}_i} \right), \quad (3)$$

capturing the intensity of ambient daytime and evening heat. The Energy-Only index is multiplicative,

$$I_i^{\text{energy}} = \widetilde{\text{burden}_i} \times \tilde{N}_i^{\text{LI}}, \quad (4)$$

so that a tract scores highly only when it combines a high low-income energy burden *and* a large number of low-income households; this favors places where weatherization would relieve both intense per-household stress and a large affected population, rather than tracts that are burdened but sparsely populated or populous but only mildly burdened. The Joint-Priority index is the equal-weight average of the two,

$$I_i^{\text{joint}} = 0.5 I_i^{\text{heat}} + 0.5 I_i^{\text{energy}}, \quad (5)$$

representing a balanced program that values physical heat exposure and economic vulnerability equally. The equal weighting is a transparent default; the framework accommodates alternative weights, and the Monte Carlo analysis below characterizes the stability of the resulting allocation.

2.5 Constrained allocation algorithm

Each strategy distributes exactly $S = 1,000$ slots among a selected set of tracts subject to a per-tract floor of $L = 25$ and ceiling of $U = 200$ slots. We first defined a scale-aware allocation weight $P_i = I_i \cdot N_i^{\text{LI}}$ that multiplies a tract’s index by its low-income household count, so that resources track the number of households a program could plausibly serve. For a candidate portfolio size K , the top- K tracts by P_i receive a proportional pre-allocation $s_i = S \cdot P_i / \sum_{j \in \text{top-}K} P_j$; an iterative clip-and-redistribute loop then enforces the $[L, U]$ bounds by capping tracts above U and raising tracts below L , redistributing the surplus or deficit proportionally among the unclipped tracts until the bounds are satisfied and the counts sum to S . To avoid the two degenerate regimes—too few tracts, which forces every tract to the ceiling, or too many, which forces every tract to the floor—we selected the optimal K from the “valid range,” the set of K for which the pure proportional split already lies within $[L, U]$ without clipping, choosing the K that maximizes the total priority $\sum_i P_i$ of the selected tracts. This procedure preserves proportionality while respecting the operational bounds. The selected portfolio sizes were $K = 28$ for Heat-Only, $K = 22$ for Energy-Only, and $K = 27$ for Joint-Priority, with realized per-tract allocations spanning 25–53, 25–93, and 26–69 slots, respectively; all three sum exactly to 1,000.

2.6 Monte Carlo uncertainty propagation

To characterize how the allocation responds to input uncertainty, we ran a 1,000-trial Monte Carlo simulation (fixed seed for reproducibility). In each trial we perturbed the inputs and recomputed the indices and the constrained allocation at the baseline optimal K . Zonal (heat) uncertainty was modeled by drawing each tract’s three temperature means from a normal distribution centered on the observed mean with standard deviation equal to the standard error of the mean, $\text{SEM} = \sigma_{\text{pixel}} / \sqrt{n_{\text{pixel}}}$, where σ_{pixel} and n_{pixel} are the within-tract pixel standard deviation and valid-pixel count. LEAD estimation uncertainty was modeled by perturbing low-income energy

burden with a coefficient of variation of 10% and low-income household counts with a coefficient of variation of 15%, using a median-preserving log-normal distribution truncated at zero to respect the non-negativity of both quantities. For every tract and strategy we recorded, across the 1,000 trials, the mean and standard deviation of allocated slots, the 2.5th and 97.5th percentiles (a 95% Monte Carlo interval), and the selection probability—the fraction of trials in which the tract received any allocation. This choice of noise model reflects the two dominant, independent error sources: finite-sample zonal averaging of the raster and the modeled nature of LEAD estimates [30, 36]. The framework is deliberately agnostic to the MAUP: because it operates on fixed 2022 tracts, its rankings are conditional on that zonation, a limitation we return to in the Discussion [33, 34].

2.7 Strategy comparison

We compared the three strategies pairwise using four metrics: the fraction of tracts whose slot counts differ (slot mismatch), the fraction whose binary selection status differs (selection mismatch), the Pearson and Spearman correlations of the per-tract slot vectors, and the Jaccard similarity of the selected tract sets. All analyses were implemented in Python using `geopandas`, `rasterio`, `numpy`, and `pandas`; maps were produced in `matplotlib` with layers reprojected to EPSG:32618 for metrically accurate scale bars.

3 Results

3.1 Tract-level heat and energy-burden landscape

The 126 populated tracts exhibit a coherent but heterogeneous heat-and-burden landscape. Afternoon heat index, the most policy-relevant thermal exposure, varies by 6.6°F across tracts (91.9–98.5°F), and evening air temperature—an indicator of the overnight relief that governs physiological recovery—varies by 4.3°F (85.7–90.0°F). Low-income energy burden spans a threefold range (3.2–9.8%), with a study-area mean of 6.2% that sits at the conventional “high-burden” threshold, and more than half of the tracts exceed it. Critically, the thermal and economic dimensions are only partially aligned: the multiplicative Energy-Only index rewards tracts that pair severe burden with large low-income populations, which need not be the hottest tracts. This partial decoupling is what makes the choice of priority definition consequential, as the allocations below make explicit.

3.2 Heat-Only allocation

The Heat-Only strategy selects 28 tracts and distributes 1,000 slots across them, with per-tract allocations from 25 to 53 (Figure 2). The highest allocations fall in a compact set of central corridors combining high afternoon heat index with substantial low-income populations: Census Tract 607 (53 slots; afternoon heat index 96.2°F), Tract 204 (48 slots), Tract 610.02 and Tract 305.01 (44 slots each; heat index 96.0–96.2°F), and Tract 706.01 (43 slots). Because the Heat-Only index ignores economic vulnerability, it draws in a distinct set of tracts—twelve tracts selected under Heat-Only are *not* selected under Energy-Only—that are thermally extreme but not necessarily among the most energy-burdened, including several in northern and eastern Henrico County. The spatial signature (Figure 2) is a concentrated hot core near the city center with satellite selections along peripheral heat pockets.




Figure 2: Heat-Only allocation. Per-tract cooling/weatherization slots under the Heat-Only priority index (28 selected tracts; 1,000 slots; 25–53 slots per tract). Unselected tracts are shown in neutral grey; selected tracts are shaded by slot count (truncated YlOrRd palette). The allocation concentrates in central high-heat-index corridors, with satellite selections in peripheral heat pockets. Colormaps differ across Figures 2–4 and encode strategy-specific slot ranges; they are not directly comparable between panels.

3.3 Energy-Only allocation

The Energy-Only strategy is the most concentrated of the three, selecting only 22 tracts but permitting individual allocations up to 93 slots—nearly the operational ceiling—because the multiplicative index produces a small number of very high-priority tracts (Figure 3). The largest allocations go to Census Tract 607 (93 slots; low-income burden 7.3% and 1,648 low-income households, the largest such population in the study area), Tract 708.03 (80 slots; burden 7.7%, 1,466 households), Tract 204 (64 slots), Tract 604 (60 slots; burden 8.0%), and Tract 2010.03 in Henrico County (59 slots). Six tracts are selected under Energy-Only but not under Heat-Only, reflecting places where economic burden and population scale, rather than ambient heat, drive priority. The concentration of the Energy-Only allocation means that fewer neighborhoods are served, but those that are receive

substantially deeper investment—a direct expression of the trade-off between breadth and depth of program reach.




Figure 3: Energy-Only allocation. Per-tract slots under the Energy-Only priority index (22 selected tracts; 1,000 slots; 25–93 slots per tract; truncated Purples palette). The multiplicative burden $\times$ population index yields a small number of deeply funded tracts, concentrating resources where both energy burden and the low-income population are large.

3.4 Joint-Priority allocation

The balanced Joint-Priority strategy selects 27 tracts with allocations from 26 to 69 slots (Figure 4), occupying an intermediate position between the breadth of Heat-Only and the concentration of Energy-Only. Twenty of the 27 selected tracts lie in the City of Richmond and seven in Henrico County; no Chesterfield County tract is selected. Together the 27 Joint-Priority tracts contain 32,112 low-income households, roughly 36% of the study area’s low-income households, indicating that the allocation reaches a substantial share of the population it is intended to serve while remaining geographically focused. The five largest allocations—Census Tract 607 (69 slots), Tract 204 (55),

Tract 604 (50), Tract 706.01 (47), and Tract 708.03 (46)—are tracts that rank highly on *both* dimensions, combining afternoon heat indices near 95–96°F with low-income burdens of 5.7–8.0% and low-income populations exceeding 1,200 households. Table 1 lists all 27 tracts with their GEOIDs, jurisdictions, final slot counts, key heat and burden metrics, and Monte Carlo selection statistics. The spatial pattern (Figure 4) is a stable inner-city core centered on the highest-priority tracts, with a ring of moderately funded tracts extending into inner Henrico County.




Figure 4: Joint-Priority allocation. Per-tract slots under the balanced Joint-Priority index ($0.5 I^{\text{heat}} + 0.5 I^{\text{energy}}$; 27 selected tracts; 1,000 slots; 26–69 slots per tract; viridis palette). The allocation concentrates in a central core (peak in Census Tract 607) and extends into inner Henrico County, balancing heat exposure and energy burden.

Table 1: Joint-Priority allocation: all 27 selected tracts. Tracts are ordered by final slot count. Heat index and evening temperature are tract-level zonal means ($^{\circ}\text{F}$); low-income (LI) energy burden and household count are from the 2022 DOE LEAD 0–80% AMI aggregation. Selection probability (p_{sel}) and the 95% Monte Carlo (MC) slot interval are from the 1,000-run simulation. A lower MC interval bound of 0 indicates trials in which the tract fell below the 25-slot floor and was therefore assigned no slots. RC = City of Richmond; HN = Henrico County.

GEOID	Tract	Juris.	Slots	Heat idx. ($^{\circ}\text{F}$)	Eve. T ($^{\circ}\text{F}$)	LI burden (%)	LI h/holds (n)	p_{sel}	95% MC CI (slots)
51760060700	607	RC	69	96.2	88.3	7.31	1648	1.00	42–96
51760020400	204	RC	55	95.4	89.4	6.98	1422	1.00	35–77
51760060400	604	RC	50	95.7	88.7	8.04	1225	1.00	32–70
51760070601	706.01	RC	47	95.1	88.1	5.70	1563	1.00	31–67
51760070803	708.03	RC	46	94.7	87.0	7.67	1466	0.97	0–69
51087201003	2010.03	HN	45	95.4	88.2	7.09	1352	0.98	27–66
51760020200	202	RC	42	94.6	89.7	7.76	1205	0.97	0–60
51760040400	404	RC	42	94.9	88.8	7.80	1192	0.98	25–59
51760071100	711	RC	40	96.9	86.3	5.55	1591	0.96	0–59
51760010800	108	RC	37	96.0	88.7	7.64	1058	0.93	0–57
51087200807	2008.07	HN	37	95.5	87.0	6.79	1282	0.90	0–54
51087201504	2015.04	HN	36	96.0	88.1	7.49	1106	0.92	0–54
51760030100	301	RC	36	96.4	89.3	6.04	1066	0.96	0–50
51760030501	305.01	RC	35	96.0	88.9	4.66	1185	0.98	25–51
51087200805	2008.05	HN	33	95.9	88.2	7.30	1023	0.87	0–49
51087201406	2014.06	HN	33	95.2	88.5	5.71	1209	0.87	0–49
51760070700	707	RC	32	94.9	87.0	5.80	1389	0.73	0–49
51760061002	610.02	RC	31	96.2	88.7	3.52	1392	0.94	0–45
51760041100	411	RC	30	95.7	89.1	6.20	931	0.76	0–42
51760061001	610.01	RC	29	96.4	88.4	5.11	1002	0.72	0–42
51087200600	2006	HN	29	93.9	87.7	6.82	1103	0.59	0–42
51760060800	608	RC	29	96.3	88.5	7.23	882	0.65	0–41
51760040202	402.02	RC	28	97.2	89.8	4.24	850	0.77	0–41
51760020502	205.02	RC	28	95.9	89.8	3.41	966	0.77	0–39
51760030200	302	RC	28	96.1	89.4	5.08	911	0.71	0–39
51760020501	205.01	RC	27	96.0	89.1	3.30	1190	0.73	0–39
51087200305	2003.05	HN	26	95.7	88.7	5.19	903	0.53	0–39

3.5 Strategy comparison

The three strategies overlap substantially but are far from interchangeable (Table 2). Heat-Only and Joint-Priority are the most concordant pair: their per-tract slot vectors correlate at Pearson $r = 0.89$ (Spearman $\rho = 0.90$), they disagree on the selection status of only 4% of tracts, and their selected sets have a Jaccard similarity of 0.83. This tight coupling arises because the low-income household count enters both the Heat-Only allocation weight ($P_i = I_i^{\text{heat}} \cdot N_i^{\text{LI}}$) and the Joint-Priority index, so both strategies favor populous, hot tracts. Energy-Only and Joint-Priority are moderately

concordant ($r = 0.84$, $\rho = 0.76$, Jaccard = 0.58). The Heat-Only and Energy-Only pair diverges the most, with the weakest correlation ($r = 0.61$, $\rho = 0.57$), a selection mismatch of 14%, and a Jaccard of only 0.47—that is, the two “pure” strategies agree on fewer than half of their selected tracts. This confirms quantitatively that targeting physical heat and targeting energy burden are genuinely different policies, and that the Joint-Priority strategy operates as a compromise that leans slightly toward the heat dimension because of the shared population weighting.

Table 2: Pairwise comparison of allocation strategies. Slot and selection mismatch are the fractions of the 126 pool tracts that differ in slot count and in binary selection, respectively. Pearson r and Spearman ρ compare the per-tract slot vectors; Jaccard is the similarity of the selected tract sets.

Strategy pair	Slot mismatch	Selection mismatch	Pearson r	Spearman ρ	Selection Jaccard
Heat vs. Energy	0.27	0.14	0.61	0.57	0.47
Heat vs. Joint	0.22	0.04	0.89	0.90	0.83
Energy vs. Joint	0.25	0.10	0.84	0.76	0.58

3.6 Uncertainty and selection robustness

The Monte Carlo analysis reveals that the Joint-Priority point allocation, though it identifies a clear priority ordering, rests on a spectrum of selection stability rather than a set of unambiguous choices (Figure 5). Across the 1,000 trials, only two tracts—Census Tract 607 and Tract 204—are selected in essentially every trial ($p_{\text{sel}} \geq 0.999$), constituting a robustly prioritized core that any reasonable perturbation of the inputs would retain. A further 13 tracts are selected with high probability ($0.90 \leq p_{\text{sel}} < 0.999$), so that 15 of the 27 baseline tracts have $p_{\text{sel}} \geq 0.90$ and 20 have $p_{\text{sel}} \geq 0.75$. The remaining seven baseline selections are more marginal ($p_{\text{sel}} < 0.75$, the lowest being Tract 2003.05 at 0.53), meaning they are chosen at the baseline but are displaced by higher-scoring competitors in a substantial minority of trials. Fifty-five tracts are never selected in any trial. The mean selection probability among the 27 baseline tracts is 0.86. The selection-robustness map (Figure 5) shows this structure geographically: a dark-green stable core near the city center, surrounded by lighter-green and yellow tracts whose inclusion is contingent, and a large field of red tracts that are consistently outside the priority set.




Figure 5: Joint-Priority selection robustness. Monte Carlo selection probability (fraction of 1,000 trials in which a tract received any allocation) under the Joint-Priority strategy, binned into discrete classes. Only a small core (dark green, $p_{\text{sel}} \geq 0.95$) is robustly selected; many tracts are marginal (yellow/light green), and the majority are never selected (red).

The per-tract error-bar analysis (Figure 6) complements the map by showing the magnitude as well as the existence of allocation uncertainty. For the top 30 tracts by baseline Joint-Priority allocation, the highest-ranked tracts have Monte Carlo intervals that remain well above the 25-slot floor—Census Tract 607, for example, has a 95% interval of 42–96 slots that never approaches zero—indicating both certain selection and a stable, if wide, funding level. Mid-ranked tracts, by contrast, have intervals whose lower bound reaches zero (see the $p_{\text{sel}} < 1.0$ rows of Table 1), the visual signature of conditional selection: in some trials they are funded near their baseline level, in others they drop out entirely. The width of the intervals is driven primarily by the Energy-Only component, because the log-normal perturbation of burden and household counts (coefficients of variation of 10% and 15%) introduces more variance than the tight zonal standard errors of the temperature means, whose large pixel counts make the SEM small. The practical implication is that point allocations for the top core are reliable operational targets, whereas allocations for marginal tracts should be treated as provisional and revisited as better data become available.




Figure 6: Allocation sensitivity for the top 30 Joint-Priority tracts. Points show baseline slot allocations; horizontal bars span the 95% Monte Carlo interval over 1,000 trials. The 25-slot floor and 200-slot ceiling are marked. Top-ranked tracts have intervals bounded well above zero (robust selection), whereas mid-ranked tracts have intervals reaching zero (conditional selection).

4 Discussion

4.1 Principal findings

This study operationalizes a transparent, reproducible, and uncertainty-aware framework for allocating a fixed pool of cooling and weatherization resources across census tracts, and applies it to Richmond, Virginia using community heat-mapping and modeled energy-burden data. Three

findings stand out. First, the choice of priority definition materially changes which neighborhoods are served: the “pure” heat and energy strategies agree on fewer than half of their selected tracts (Jaccard = 0.47), so an implicit or unexamined weighting decision would silently determine winners and losers. Second, the balanced Joint-Priority strategy behaves as a coherent compromise, concentrating 1,000 slots in 27 tracts—a stable inner-city core in the City of Richmond extending into inner Henrico County—that together house about 36% of the study area’s low-income households while ranking highly on both heat and burden. Third, and most consequentially for practice, the Monte Carlo analysis demonstrates that only a small core of tracts is robustly prioritized, while a majority of the baseline selections are conditional on inputs that carry real uncertainty; reporting a bare ranked list would therefore overstate the confidence of many individual recommendations.

4.2 Policy implications of the heat-versus-burden trade-off

The divergence between the Heat-Only and Energy-Only allocations is not a defect of the method but a faithful representation of a genuine policy fork. A Heat-Only strategy directs resources to where the ambient thermal hazard is greatest and, because heat is the proximate driver of heat mortality and morbidity, aligns most directly with acute public-health protection during heat events [1, 3, 7]. It is the natural logic for interventions whose benefit scales with heat exposure—emergency cooling, heat-pump deployment in the hottest corridors, or shade and canopy investment in low-canopy blocks [11, 14]. An Energy-Only strategy, by contrast, directs resources to where the economic barrier to cooling is highest and the affected population largest; it aligns with the logic of weatherization and bill relief, which reduce energy burden and the energy-insecurity coping behaviors—rationing air conditioning, leaving homes dangerously hot—that convert affordability stress into health risk [23, 25, 26]. The two logics are complementary rather than contradictory: heat determines *where* protective cooling is most needed, while burden determines *whether* residents can afford to use it. The multiplicative construction of our Energy-Only index further embeds an equity-of-scale principle, favoring tracts where relief reaches many households, consistent with distributional-justice framings that weight both the intensity and the reach of benefits [28]. For a program that funds both cooling and weatherization, the Joint-Priority strategy is the defensible default precisely because it refuses to resolve this trade-off by omission, making the equal weighting explicit and auditable while leaving the weights open to deliberation with affected communities [29].

4.3 Stability and robustness of the joint-priority selections

The robustness results carry a direct operational lesson. The two-tract robust core (Census Tracts 607 and 204) and the broader set of 15 high-probability tracts ($p_{\text{sel}} \geq 0.90$) can be treated as firm commitments: they combine top-tier heat and burden with large low-income populations, and their prioritization survives essentially any plausible perturbation of the inputs. The seven marginal tracts ($p_{\text{sel}} < 0.75$), however, are better understood as members of a “bubble” of near-equivalent candidates competing for the last few portfolio positions; small changes in modeled burden or household counts reshuffle them. Rather than a weakness, this is actionable intelligence. It suggests a tiered implementation—funding the robust core with confidence, treating marginal tracts as a flexible expansion set, and prioritizing ground-truthing (for example, targeted energy-audit sampling or supplementary heat measurement) in exactly those tracts where the decision is most sensitive to data quality [30, 36]. That the interval widths are dominated by LEAD’s modeled uncertainty rather than by zonal sampling error also identifies where investment in better data would most improve allocation confidence: refining household-level energy estimates would sharpen the allocation more

than denser heat measurement, given the already large pixel counts underlying the tract heat means.

4.4 Interpreting the heat surface correctly

A central caveat, which we have maintained throughout, is that the NIHHS–CAPA heat surface is a campaign-day snapshot of ambient outdoor conditions, not a measure of household-level heat risk [19, 20]. It reflects the temperatures encountered along volunteer traverse routes on one hot day, interpolated across space; it does not encode indoor temperatures, building envelope quality, air-conditioning access, or the health status of residents, all of which mediate the translation from ambient heat to human harm [17, 18]. Our indices should accordingly be read as prioritization signals for where ambient exposure and economic vulnerability coincide, not as risk predictions for individual households. This distinction matters for equity: two tracts with identical ambient heat may harbor very different true risk depending on housing and health, which is precisely why we pair the heat surface with energy burden and why we recommend community engagement and ground-truthing before final program siting.

4.5 Limitations

Several limitations bound the interpretation of our results. First, all tract-level statistics are conditional on the 2022 census-tract zonation and are therefore subject to the modifiable areal unit problem; a different partition (for example, block groups) could shift both the zonal heat summaries and the rankings [33–35]. Second, the heat data derive from a single campaign day and season, so they capture the spatial pattern of a representative hot afternoon and evening rather than the full climatological distribution of heat or its interannual variability [3]. Third, the LEAD figures are model-based estimates rather than measured expenditures; although we propagated their uncertainty, the assumed coefficients of variation and log-normal form are approximations, and any systematic bias in LEAD would propagate to the allocation [24, 27]. Fourth, the equal weighting of the Joint-Priority index, the multiplicative form of the Energy-Only index, the 25/200 slot bounds, and the 1,000-slot budget are policy parameters rather than empirical quantities; we have shown the framework’s behavior at defensible defaults, but different values would yield different portfolios, and the framework is explicitly designed to be re-run under alternatives. Fifth, our priority indices omit dimensions of vulnerability—age structure, disability, chronic disease prevalence, social isolation, and tenure—that are known determinants of heat risk and that richer screening tools incorporate [8, 18, 29]. Finally, the analysis addresses distributional targeting only; it does not model program uptake, cost-effectiveness, or the procedural- and recognition-justice dimensions that a complete equity assessment would require [28].

4.6 Future directions

Several extensions follow naturally. The framework could be re-run across multiple zonations and heat-campaign days to characterize MAUP and temporal sensitivity directly, converting the present single-zonation analysis into an ensemble [35]. The priority indices could be broadened to a fuller heat-vulnerability specification incorporating demographic and health covariates, ideally validated against observed heat-morbidity outcomes [18, 29]. The allocation objective could be generalized from proportional targeting to an explicit optimization—for instance, maximizing covered low-income households subject to budget and equity constraints, drawing on the maximal-covering location literature [39]—and could incorporate access to existing cooling infrastructure so

as to complement rather than duplicate current resources [32]. Most importantly, the transparent, uncertainty-annotated outputs produced here are intended as an input to, not a substitute for, deliberation with the communities they concern, whose lived knowledge of indoor heat and energy hardship the ambient surface and modeled burden cannot fully capture.

5 Conclusion

We have presented and demonstrated a reproducible framework that integrates community heat-mapping and modeled energy-burden data to allocate a fixed pool of cooling and weatherization resources across census tracts, with formal propagation of the underlying uncertainties. Applied to Richmond, Virginia, the framework distributes 1,000 program slots across 27 tracts under a balanced Joint-Priority strategy, concentrating resources in a stable inner-city core while capturing about 36% of the study area’s low-income households and honoring per-tract floors and ceilings. By comparing Heat-Only, Energy-Only, and Joint-Priority strategies, we make the trade-off between targeting physical heat and targeting economic burden explicit and quantifiable rather than implicit, and by reporting Monte Carlo selection probabilities alongside point allocations, we distinguish firm commitments from provisional ones. The method is deliberately transparent and parameter-driven so that program administrators and communities can adjust weights, bounds, and budgets and immediately see the distributional consequences. We underscore that the heat surface is a campaign-day exposure snapshot and not a household-risk metric, and that the framework’s outputs are decision-support inputs meant to be refined through ground-truthing and community engagement. Within those bounds, the approach offers a defensible, auditable path toward equitable heat-and-energy resource allocation that other cities conducting NHHIS–CAPA campaigns could readily adopt.

References

- [1] Antonio Gasparrini, Yuming Guo, Masahiro Hashizume, Eric Lavigne, Antonella Zanobetti, Joel Schwartz, et al. Mortality risk attributable to high and low ambient temperature: a multicountry observational study. *The Lancet*, 386(9991):369–375, 2015. doi: 10.1016/S0140-6736(14)62114-0.
- [2] Qi Zhao, Yuming Guo, Tingting Ye, Antonio Gasparrini, Shilu Tong, et al. Global, regional, and national burden of mortality associated with non-optimal ambient temperatures from 2000 to 2019: a three-stage modelling study. *The Lancet Planetary Health*, 5(7):e415–e425, 2021. doi: 10.1016/S2542-5196(21)00081-4.
- [3] Kristie L. Ebi, Anthony Capon, Peter Berry, Carolyn Broderick, Richard de Dear, George Havenith, Yasushi Honda, R. Sari Kovats, Wei Ma, Arunima Malik, Nathan B. Morris, Lars Nybo, Sonia I. Seneviratne, Jennifer Vanos, and Ollie Jay. Hot weather and heat extremes: health risks. *The Lancet*, 398(10301):698–708, 2021. doi: 10.1016/S0140-6736(21)01208-3.
- [4] Ana M. Vicedo-Cabrera, Noah Scovronick, Francesco Sera, Dominic Royé, Rochelle Schneider, et al. The burden of heat-related mortality attributable to recent human-induced climate change. *Nature Climate Change*, 11:658–667, 2021. doi: 10.1038/s41558-021-01058-x.
- [5] Kevin R. Cromar, Susan C. Anenberg, John R. Balme, Allen A. Fawcett, Marya Ghazipura, et al. Estimates of country level temperature-related mortality damage functions. *Scientific Reports*, 11:20282, 2021. doi: 10.1038/s41598-021-99156-5.

- [6] G. Brooke Anderson and Michelle L. Bell. Heat waves in the united states: Mortality risk during heat waves and effect modification by heat wave characteristics in 43 u.s. communities. *Environmental Health Perspectives*, 119(2):210–218, 2011. doi: 10.1289/ehp.1002313.
- [7] Jan C. Semenza, Carol H. Rubin, Kenneth H. Falter, Joel D. Selanikio, W. Dana Flanders, Holly L. Howe, and John L. Wilhelm. Heat-related deaths during the july 1995 heat wave in chicago. *New England Journal of Medicine*, 335(2):84–90, 1996. doi: 10.1056/NEJM199607113350203.
- [8] Eric Klinenberg. *Heat Wave: A Social Autopsy of Disaster in Chicago*. University of Chicago Press, Chicago, IL, 2002.
- [9] T. R. Oke. The energetic basis of the urban heat island. *Quarterly Journal of the Royal Meteorological Society*, 108(455):1–24, 1982. doi: 10.1002/qj.49710845502.
- [10] Qihao Weng, Dengsheng Lu, and Jacquelyn Schubring. Estimation of land surface temperature–vegetation abundance relationship for urban heat island studies. *Remote Sensing of Environment*, 89(4):467–483, 2004. doi: 10.1016/j.rse.2003.11.005.
- [11] Carly D. Ziter, Eric J. Pedersen, Christopher J. Kucharik, and Monica G. Turner. Scale-dependent interactions between tree canopy cover and impervious surfaces reduce daytime urban heat during summer. *Proceedings of the National Academy of Sciences*, 116(15):7575–7580, 2019. doi: 10.1073/pnas.1817561116.
- [12] Iain D. Stewart and Tim R. Oke. Local climate zones for urban temperature studies. *Bulletin of the American Meteorological Society*, 93(12):1879–1900, 2012. doi: 10.1175/BAMS-D-11-00019.1.
- [13] Matthias Demuzere, Jonas Kittner, Alberto Martilli, Gerald Mills, Christian Moede, Iain D. Stewart, Jasper van Vliet, and Benjamin Bechtel. A global map of local climate zones to support earth system modelling and urban-scale environmental science. *Earth System Science Data*, 14(8):3835–3873, 2022. doi: 10.5194/essd-14-3835-2022.
- [14] Jeremy S. Hoffman, Vivek Shandas, and Nicholas Pendleton. The effects of historical housing policies on resident exposure to intra-urban heat: A study of 108 us urban areas. *Climate*, 8(1):12, 2020. doi: 10.3390/cli8010012.
- [15] Bev Wilson. Urban heat management and the legacy of redlining. *Journal of the American Planning Association*, 86(4):443–457, 2020. doi: 10.1080/01944363.2020.1759127.
- [16] Angel Hsu, Glenn Sheriff, Tirthankar Chakraborty, and Diego Manya. Disproportionate exposure to urban heat island intensity across major us cities. *Nature Communications*, 12(1):4104, 2021. doi: 10.1038/s41467-021-22799-5.
- [17] Sharon L. Harlan, Anthony J. Brazel, Lela Prashad, William L. Stefanov, and Larissa Larsen. Neighborhood microclimates and vulnerability to heat stress. *Social Science & Medicine*, 63(11):2847–2863, 2006. doi: 10.1016/j.socscimed.2006.07.030.
- [18] Colleen E. Reid, Marie S. O’Neill, Carina J. Gronlund, Shannon J. Brines, Daniel G. Brown, Ana V. Diez-Roux, and Joel Schwartz. Mapping community determinants of heat vulnerability. *Environmental Health Perspectives*, 117(11):1730–1736, 2009. doi: 10.1289/ehp.0900683.

- [19] Vivek Shandas, Jackson Voelkel, Joseph Williams, and Jeremy Hoffman. Integrating satellite and ground measurements for predicting locations of extreme urban heat. *Climate*, 7(1):5, 2019. doi: 10.3390/cli7010005.
- [20] Jackson Voelkel and Vivek Shandas. Towards systematic prediction of urban heat islands: Grounding measurements, assessing modeling techniques. *Climate*, 5(2):41, 2017. doi: 10.3390/cli5020041.
- [21] Ariel Dreihobl and Lauren Ross. Lifting the high energy burden in america’s largest cities: How energy efficiency can improve low-income and underserved communities. Technical report, American Council for an Energy-Efficient Economy (ACEEE), Washington, DC, 2016.
- [22] Ariel Dreihobl, Lauren Ross, and Roxana Ayala. How high are household energy burdens? an assessment of national and metropolitan energy burden across the united states. Technical Report U2006, American Council for an Energy-Efficient Economy (ACEEE), Washington, DC, 2020.
- [23] Trevor Memmott, Sanya Carley, Michelle Graff, and David M. Konisky. Sociodemographic disparities in energy insecurity among low-income households before and during the covid-19 pandemic. *Nature Energy*, 6(2):186–193, 2021. doi: 10.1038/s41560-020-00763-9.
- [24] Eric Scheier and Noah Kittner. A measurement strategy to address disparities across household energy burdens. *Nature Communications*, 13(1):288, 2022. doi: 10.1038/s41467-021-27673-y.
- [25] Diana Hernández. Understanding ‘energy insecurity’ and why it matters to health. *Social Science & Medicine*, 167:1–10, 2016. doi: 10.1016/j.socscimed.2016.08.029.
- [26] Meredith Fowlie, Michael Greenstone, and Catherine Wolfram. Do energy efficiency investments deliver? evidence from the weatherization assistance program. *The Quarterly Journal of Economics*, 133(3):1597–1644, 2018. doi: 10.1093/qje/qjy005.
- [27] U.S. Department of Energy and National Renewable Energy Laboratory. Low-income energy affordability data (lead) tool: Virginia 2022 area median income census tracts dataset. OpenEI / Open Energy Data Initiative, submission 6219, 2022.
- [28] Lara Cushing, John Faust, Laura Meehan August, Rose Cendak, Walker Wieland, and George Alexeeff. Racial/ethnic disparities in cumulative environmental health impacts in california: Evidence from a statewide environmental justice screening tool (calenviroscreen 1.1). *American Journal of Public Health*, 105(11):2341–2348, 2015. doi: 10.2105/AJPH.2015.302643.
- [29] Yang Niu, Zhaowu Li, Yang Gao, Xin Liu, Lei Xu, Sotiris Vardoulakis, Yujuan Yue, Jun Wang, and Qiyong Liu. A systematic review of the development and validation of the heat vulnerability index: Major factors, methods, and spatial units. *Current Climate Change Reports*, 7(3):87–97, 2021. doi: 10.1007/s40641-021-00173-3.
- [30] Salvatore Greco, Alessio Ishizaka, Menelaos Tasiou, and Gianpiero Torrissi. On the methodological framework of composite indices: A review of the issues of weighting, aggregation, and robustness. *Social Indicators Research*, 141(1):61–94, 2019. doi: 10.1007/s11205-017-1832-9.
- [31] OECD and European Commission, Joint Research Centre. *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris, 2008. doi: 10.1787/9789264043466-en.

- [32] Quinn H. Adams, Elana M. G. Chan, Keith R. Spangler, Kate R. Weinberger, Kevin J. Lane, Nicole A. Errett, Jeremy J. Hess, Yuantong Sun, Gregory A. Wellenius, and Amruta Nori-Sarma. Examining the optimal placement of cooling centers to serve populations at high risk of extreme heat exposure in 81 us cities. *Public Health Reports*, 138(6):955–962, 2023. doi: 10.1177/00333549221148174.
- [33] Stan Openshaw. *The Modifiable Areal Unit Problem*. Number 38 in Concepts and Techniques in Modern Geography (CATMOG). Geo Books, Norwich, UK, 1984.
- [34] Jue Wang, Mei-Po Kwan, and Yanwei Chai. Incorporating geography into a new generalized theoretical and statistical framework addressing the modifiable areal unit problem. *International Journal of Health Geographics*, 18(1):7, 2019. doi: 10.1186/s12942-019-0170-3.
- [35] Jeremiah J. Nieves and Andrea E. Gaughan. A simulated ‘sandbox’ for exploring the modifiable areal unit problem in aggregation and disaggregation. *Scientific Data*, 11:239, 2024. doi: 10.1038/s41597-024-03061-1.
- [36] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global Sensitivity Analysis: The Primer*. John Wiley & Sons, Chichester, UK, 2008. doi: 10.1002/9780470725184.
- [37] U.S. Census Bureau. 2022 tiger/line shapefiles: Census tracts, virginia. U.S. Department of Commerce, 2022. URL <https://www2.census.gov/geo/tiger/TIGER2022/TRACT/>.
- [38] CAPA Strategies, National Integrated Heat Health Information System (NIHHIS), and Science Museum of Virginia. Richmond, virginia urban heat island mapping campaign 2021: Community-science heat and heat-index rasters. Open Science Framework, project 3xvmg, 2021. URL <https://osf.io/3xvmg/>.
- [39] Richard Church and Charles ReVelle. The maximal covering location problem. *Papers in Regional Science*, 32(1):101–118, 1974. doi: 10.1007/BF01942293.