https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

URBADAPT-HEAT v1.0: a scalable geospatial framework for city-level public-private adaptation infrastructure cost-benefit analysis and its urban heat risk implementation

Armande Aboudrar-Méda[1,2,★] and Giacomo Falchetta[1,2,3,★]

[1]Centro Euro-Mediterraneo sui Cambiamenti Climatici, Venice, Italy
[2]RFF-CMCC European Institute on Economics and the Environment, Venice, Italy
[3]International Institute for Applied Systems Analysis, Schlossplatz 1, Laxenburg, A-2361, Austria
★These authors contributed equally to this work.

Correspondence: Armande Aboudrar-Méda (armande.aboudrar-meda@cmcc.it) and Giacomo Falchetta (giacomo.falchetta@cmcc.it)

Abstract. Urban populations experience increasing heat-attributable mortality under climate change and several modelling frameworks for assessing risk across the hazard, exposure, and vulnerability components have been developed in recent years. Yet, few explicitly represent and evaluate public and private adaptation options to tackle urban heat risk at a within-city, spatially-explicit scale, as well as their effectiveness, costs, interactions, and equity implications. We introduce URBADAPT-HEAT v1.0, a modular, city-agnostic geospatial framework for urban heat risk assessment and adaptation cost-benefit analysis at the urban scale. URBADAPT-HEAT couples four risk components - hazard, exposure, vulnerability, and impact - tailored to outdoor ambient heat and heat-attributable mortality. It combines high-resolution temperature fields with scenario-consistent climate projections to 2050, demographic projections, and a composite social vulnerability index that evolves under socioeconomic futures. Age-specific heat-mortality impact functions bridge epidemiological evidence and risk modelling within a unified impact-function framework. Three adaptation policy modules - spatially allocated air conditioning and street trees, together with a citywide, age-specific early warning system - are evaluated within a common adaptation module that represents key biophysical and economic feedbacks, including air-conditioning-induced urban heat-island intensification and vegetation-driven reductions in cooling demand. A cost-benefit analysis module integrates discounted costs and mortality benefits across a 25-year policy horizon, supports budget-constrained portfolio optimisation and Pareto-frontier analysis, and reports distributional diagnostics alongside the aggregate portfolio search. Selected structural and parametric uncertainty dimensions are propagated through Latin Hypercube sampling and evaluated using PAWN global sensitivity analysis. The framework is configured through city-specific YAML files and implemented as a sequential pipeline of modular notebooks with documented input-output interfaces, enabling consistent application across cities. We demonstrate URBADAPT-HEAT through an implementation across four European cities (Rome, Athens, Lisbon, and Copenhagen) to characterise framework behaviour under diverse climatic and socioeco-

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

nomic conditions. The framework has additionally been implemented across the 40-city European PROVIDE sample; the resulting comparative analysis is reserved for a separate application study.

1 Introduction

Urban heat is one of the most rapidly escalating climate-related risks for human health Kriit et al. (2026); Matthews et al. (2025); Dodman et al. (2022). Rising global temperatures are amplified through the urban heat island (UHI) effect produces localised warming of several degrees above rural ambient levels in densely built city environments (Yang et al., 2024; Huang et al., 2019). More than half of the global population now lives in urban areas, a share projected to reach nearly 70 % by 2050, concentrating heat exposure in environments whose built form, social inequality, and infrastructure legacy create strong within-city heterogeneity in both hazard intensity and adaptive capacity (Creutzig et al., 2025). Heat-related mortality is already substantial and growing: epidemiological analyses attribute approximately 20,000 premature deaths per year to heat across 854 European cities alone, with the burden concentrated in southern Europe and in socially disadvantaged urban populations (Masselot et al., 2023). The excess warming produced by the UHI itself accounts for an estimated 4.3 % of all summer deaths across 93 European cities (Iungman et al., 2023), imposing economic costs of around €192 per adult urban resident per year (Huang et al., 2023). Under continued climate change and demographic ageing, heat-related mortality in European cities is projected to increase several-fold by mid-century in the absence of targeted adaptation (Masselot et al., 2025; Gallo et al., 2024; Stalhandske et al., 2022). These trends elevate urban heat adaptation from a locally desirable to a strategically necessary element of climate policy.

Translating this evidence base into adaptation planning requires tools that go substantially beyond ex-ante risk assessment or ex-post epidemiological studies. City governments, infrastructure planners, and health authorities need frameworks capable of (i) characterising and projecting heat hazard at the intra-city spatial resolution relevant to adaptation decisions; (ii) quantifying heat-attributable mortality under alternative climate, demographic, and socioeconomic scenarios; (iii) evaluating the relative costs and benefits of public and private adaptation measures on a common accounting basis; and (iv) identifying cost-effective portfolios of interventions while tracking distributional outcomes across socially vulnerable populations. Despite significant scientific advances on each of these components individually, few open-source, city-agnostic framework currently integrate all four elements in a single reproducible workflow Liu et al. (2026). Multi-city epidemiological studies provide robust attributable burden estimates and future projections but are not designed for comparative adaptation appraisal (Masselot et al., 2023, 2025). High-resolution urban climate models characterise intra-city temperature gradients but typically stop short of mortality impact quantification and cost-benefit integration (De Ridder et al., 2015). General-purpose probabilistic risk platforms such as CLIMADA offer flexible adaptation appraisal frameworks but lack the heat-specific pre-processing, dose-response calibration, and public-private policy structure required for urban heat implementation (Aznar-Siguan and

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Bresch, 2019; Bresch and Aznar-Siguan, 2021). Existing spatially explicit cost-benefit analyses of specific measures - urban trees, air conditioning, heat-health warning systems - are typically confined to a single city and a single policy option, limiting cross-city comparability and precluding portfolio optimisation (Johnson et al., 2021; Chiabai et al., 2018).

Urban heat adaptation is further complicated by the distinct mechanisms, cost structures, and interaction effects of the principal policy options. Urban vegetation primarily reduces the outdoor heat hazard through evapotranspiration and shading (with co-benefits for air quality, biodiversity, urban amenity, mental health and other sub-clinical outcomes) (Bowler et al., 2010; Li et al., 2024; Sharifi et al., 2021). Air conditioning provides individual-level indoor protection but generates electricity demand, can amplify the UHI through waste-heat rejection, and raises concerns about equity and affordability because the households most exposed to heat are often least able to access and operate cooling systems (Salamanca et al., 2014; Mazzone et al., 2023; Falchetta et al., 2024). Heat-health early warning systems (EWS) are institutionally managed, relatively low-cost public interventions that reduce mortality through behavioural and health-service response on high-risk days Pavanello and Valenti (2025); Urban et al. (2025), but their marginal benefit depends on existing adaptation infrastructure and the baseline level of air conditioning coverage (Toloo et al., 2013; Casanueva et al., 2019). These pathways are not independent: vegetation cooling reduces air-conditioning energy demand; widespread AC attenuates the residual mortality pool addressable by warning systems; and AC waste-heat feedbacks partially offset hazard reductions from urban greening. Because AC is predominantly a private investment while greening and EWS are publicly financed, the distribution of costs and benefits across households and public budgets also diverges systematically (Tompkins and Eakin, 2012). Capturing these interactions and distributional patterns requires a framework that treats adaptation options as pathways operating on hazard, exposure, or impact functions, rather than as interchangeable percentage reductions in a single risk metric.

Here we present URBADAPT-HEAT v1.0, an open-source geospatial framework for city-level urban heat risk assessment and public-private adaptation cost-benefit analysis. The framework builds on the CLIMADA probabilistic risk engine (Aznar-Siguan and Bresch, 2019; Bresch and Aznar-Siguan, 2021) and extends it with purpose-built modules for high-resolution heat hazard construction from the UrbClim urban climate dataset (Souverijns et al., 2026; De Ridder et al., 2015), age-stratified population exposure derived from WorldPop Bondarenko et al. (2025) and projected with Wittgenstein Centre demographics KC et al. (2024), a dynamic composite social vulnerability index Huisman et al. (2025), age-specific heat-mortality impact functions calibrated from published epidemiological evidence Masselot et al. (2025, 2023), and a three-pathway adaptation module covering AC, urban street trees, and EWS. Policy scenarios are evaluated against a realistic current-AC baseline rather than a hypothetical no-adaptation counterfactual, and adaptation options are represented through their physical mechanisms - hazard modification for trees, impact-function attenuation for AC, and day-specific mortality reduction for EWS - preserving the causal structure of their interactions and externalities. A 25-year discounted cost-benefit analysis integrates mortality benefits,

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

capital and operating costs, externalities such as the AC waste-heat penalty, co-benefits such as vegetation-driven reductions in AC electricity demand, and distributional outcomes stratified by a composite social vulnerability surface. Configuration is externalised to city-specific YAML files, allowing the same analytical code to operate across cities without modification and making the framework straightforwardly extensible to new cities and future hazards. URBADAPT-HEAT is the heat-specific implementation of the broader URBADAPT modular architecture, designed to progressively encompass additional urban climate hazards; this paper documents the heat implementation in full as version 1.0 and introduces the overall URBADAPT framework.

The paper is structured as follows. Section 2 provides background on urban climate risk modelling, urban heat adaptation assessment, cost-benefit analysis tools, and the public-private adaptation interface. Section 3 describes the URBADAPT-HEAT modelling framework in detail, covering the hazard, exposure, vulnerability, heat-mortality impact functions, adaptation pathways, cost-benefit analysis, and uncertainty quantification components. Section 4 presents the implementation and multi-city results for four European cities spanning contrasting climatic, demographic, and morphological contexts. Section 5 discusses implications, limitations, and the planned extension of URBADAPT to additional urban climate hazards.

2 Background

2.1 Urban climate risk modelling

Urban climate risk arises from the interaction between climate-related hazards, exposed people and assets, and the vulnerability of those exposed systems, and cities concentrate all three elements. They host large and growing populations, dense infrastructure, and strong social inequalities, while their built form modifies local climate through altered surface energy balance, reduced ventilation, heat storage, and anthropogenic heat emissions. The IPCC AR6 assessment Dodman et al. (2022) identifies cities, settlements, and key infrastructure as central sites of climate risk, and notes that urbanisation can both increase exposure and create adaptation opportunities when planning, infrastructure, and governance are aligned. For heat risk this implies that city-scale modelling cannot rely on regional mean warming alone; it must represent spatial variation in heat hazard, age and socioeconomic structure, baseline mortality, and the location of adaptation-relevant infrastructure.

Recent urban heat-health assessments have quantified present and future mortality burdens at the city scale. Multi-city European analyses show that heat-attributable mortality is already substantial and is expected to increase under climate change unless mitigation and adaptation reduce exposure and vulnerability (Masselot et al., 2023, 2025). Counterfactual analysis of the 2023 European summer suggests that past adaptation reduced the burden relative to a less adapted population facing similar temperatures (Gallo et al., 2024). These studies are valuable benchmarks, but their primary objective is usually epidemiological attribution or projection rather than the spatially

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

explicit appraisal of alternative adaptation portfolios.

A second strand has developed the physical and geospatial basis for urban risk modelling. Urban climate datasets such as UrbClim provide high-resolution temperature fields that resolve within-city hazard gradients (De Ridder et al., 2015), and the Local Climate Zone framework offers a standardised classification of urban form and land cover, allowing the thermal behaviour of compact, open, industrial, vegetated, and water-dominated areas to be compared across cities (Stewart and Oke, 2012). Together these make it possible to model risk on a common grid that preserves neighbourhood-scale heterogeneity. Probabilistic platforms then combine hazard, exposure, vulnerability, and impact functions: CLIMADA is an open-source, event-based risk modelling platform for weather and climate hazards, later extended to adaptation appraisal through avoided-damage calculations (Aznar-Siguan and Bresch, 2019; Bresch and Aznar-Siguan, 2021). Its generic structure is well suited to transparent and reproducible risk accounting, but urban heat adaptation raises specific needs, including daily gridded urban heat hazards, age-specific population exposure, heat-mortality response functions, social vulnerability surfaces, and policy modules that modify the hazard, the exposure-response relationship, or both.

2.2 Urban adaptation assessment

Urban heat adaptation spans measures that operate through different mechanisms. Some reduce outdoor exposure by changing the urban environment, for example through trees, parks, permeable surfaces, shading, cool roofs, or urban design Han et al. (2023). Others reduce indoor exposure or physiological vulnerability through air conditioning, passive cooling, building retrofits, cooling centres (Adams et al., 2023; Santamouris and Vasilakopoulou, 2026), and targeted social care. A third class acts through information, preparedness, and behavioural response, including heat-health warning systems and emergency outreach Pavanello and Valenti (2025); Urban et al. (2025); Johar et al. (2025). Treating these as interchangeable "adaptation measures" can mislead, because their spatial scale, timing, cost bearer, equity implications, and physical feedbacks differ. A substantial literature accordingly develops structured approaches to assess, prioritise, and plan urban heat adaptation, from vulnerability and impact-assessment models for complex urban systems (Apreda et al., 2019), to multi-outcome prioritisation of adaptation infrastructure (Jiang et al., 2025), to city-scale resilient development and local adaptation planning (Simpson et al., 2023; Buzási et al., 2024) and adaptive decision-making under deep uncertainty (Haasnoot et al., 2013).

Urban greening is among the most studied options. Systematic reviews show that parks and vegetated areas are generally cooler than their surroundings, though the magnitude depends on vegetation structure, urban morphology, background climate, water availability, and measurement scale (Bowler et al., 2010; Norton et al., 2015). Recent global evidence confirms substantial cooling from green spaces while showing that the benefits are unequally distributed within and across cities (Li et al., 2024; Falchetta et al., 2026b). Street-level greenness indicators add a

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

complementary perspective by capturing the greenery visible to pedestrians rather than only the horizontal vegetation cover seen from above (Li et al., 2015; Falchetta and Hammad, 2025); this matters because pedestrian-level shade, canopy geometry, and street-canyon conditions affect human thermal exposure in ways conventional vegetation indices miss.

Air conditioning is a major but more ambivalent pathway. Epidemiological evidence suggests AC can reduce heat-related mortality risk, with the size of the protection varying across countries and populations (Sera et al., 2020), yet it increases electricity demand, can stress power systems during heat waves, and rejects waste heat outdoors, potentially intensifying nocturnal urban heat islands (Salamanca et al., 2014). Access is also socially unequal: the concept of cooling poverty highlights that vulnerability to heat depends not only on climate exposure but on income, housing quality, energy affordability, health, and access to cooling technologies (Mazzone et al., 2023; Falchetta et al., 2024). Sustainable cooling therefore requires more than AC diffusion; it must consider passive cooling, urban design, behavioural practices, electricity systems, affordability, and emissions (Khosla et al., 2021). Heat-health warning systems, by contrast, are relatively low-cost institutional measures that can reduce mortality by triggering communication, outreach, behavioural change, and service mobilisation. Reviews find evidence of effectiveness but stress that causal evaluation is difficult, because warning systems are embedded in broader health and emergency-management responses (Toloo et al., 2013; Urban et al., 2025; Pavanello and Valenti, 2025). In Europe such systems are already widespread, but their thresholds, target groups, and communication pathways differ across countries and cities (Casanueva et al., 2019); modelling their benefits against a no-warning baseline can therefore overstate the marginal benefit where mature systems exist, and the more policy-relevant quantity is often the incremental benefit of improving, targeting, or expanding them.

2.3 Cost-benefit analysis tools for adaptation

Cost-benefit analysis is widely used to compare adaptation options because it provides a common accounting framework for present-value costs and avoided impacts. The Economics of Climate Adaptation framework formalised this around a sequence of questions: what climate losses are expected, how much can be averted, which measures avert them, and whether their benefits justify their costs (Economics of Climate Adaptation Working Group, 2009). CLIMADA implements this avoided-damage logic probabilistically, allowing measures to be compared through reduced impacts, implementation costs, and benefit-cost ratios (Bresch and Aznar-Siguan, 2021).

CBA of urban adaptation nonetheless has important limitations. A review highlights large differences in valuation methods, discount rates, time horizons, treatment of uncertainty, and inclusion of non-market and distributional outcomes, all of which can materially affect rankings among options (Markanday et al., 2019). This matters especially for heat, where benefits are often measured in avoided deaths or life-years while costs accrue to different actors

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

at different times; the monetisation of mortality risk, the treatment of short-term mortality displacement, and the choice between deaths avoided and life-years saved can each change benefit-cost ratios. Empirical applications confirm both the usefulness and the sensitivity of the approach: combinations of urban heat island measures in Austrian cities yielded positive benefit-cost ratios that depended on the measures and city context (Johnson et al., 2021), and the estimated benefits of heat-health warning systems vary widely depending on whether avoided mortality is valued as deaths avoided or years of life lost (Chiabai et al., 2018). These findings argue for reporting physical outcomes, such as avoided deaths and life-years, alongside discounted monetary indicators, and for preserving spatial and social information, since a citywide benefit-cost ratio can hide whether benefits reach the most heat-vulnerable populations, whether costs fall on public budgets or private households, and whether an intervention shifts risk between groups.

2.4 The public-private adaptation interface

The boundary between public and private adaptation is central to urban heat risk, because the actor who pays for adaptation is not always the actor who benefits from it. Tompkins and Eakin (Tompkins and Eakin, 2012) emphasise that adaptation goods differ in their providers and beneficiaries, and that public and private adaptation are often interdependent rather than separate categories. Air conditioning illustrates this interdependence: it is typically purchased and operated privately and can reduce individual heat exposure and mortality risk (Sera et al., 2020), yet widespread use raises electricity demand, increases emissions where power is carbon-intensive, and adds anthropogenic heat to the urban environment (Salamanca et al., 2014). Public policy may subsidise or regulate AC adoption for vulnerable households, but such policies can themselves raise infrastructure needs and externalities; conversely, public investment in trees, shading, cool public spaces, and warning systems can reduce the private need for mechanical cooling, though its benefits may be uneven if it is not spatially targeted.

This interface is also an equity issue. Households with lower income, poorer housing quality, insecure tenure, or limited access to energy services may be less able to benefit from private cooling even where heat risk is high (Mazzone et al., 2023). Public adaptation can correct some of these disparities, but only if it explicitly considers social vulnerability and access; without targeting, it may reinforce existing inequalities, for example where greening and cooling improvements concentrate in already advantaged neighbourhoods. Urban heat adaptation assessment should therefore distinguish public from private costs, direct from external effects, and average from distributional benefits, and it should evaluate policies against a realistic current-adaptation baseline rather than a hypothetical no-adaptation world.

These requirements motivate URBADAPT-HEAT, which couples a general probabilistic risk engine with urban heat-specific preprocessing, adaptation, and cost-benefit workflows and represents adaptation options as mechanisms rather than labels: vegetation primarily modifies the hazard; AC modifies the exposure-response function for cov-

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

ered populations while creating energy and waste-heat externalities; and warning systems modify behavioural and institutional response on selected high-risk days. Because these mechanisms can interact, their combined effect is not necessarily the sum of standalone benefits. CBA is implemented as a post-processing layer connected to gridded risk outputs, rather than as a separate spreadsheet exercise, so that mortality benefits, costs, co-benefits, externalities, and distributional outcomes are tracked consistently across scenarios and portfolios. Existing AC penetration is treated as part of the baseline, additional AC as an incremental private adaptation pathway, and vegetation and warning systems as public measures, with interactions such as AC waste heat, EWS-AC overlap, and vegetation-related reductions in AC electricity demand represented explicitly.

3 Modelling framework

3.1 URBADAPT

URBADAPT is a modular, city-agnostic geospatial framework for urban climate risk assessment and adaptation cost-benefit analysis. The framework integrates the CLIMADA probabilistic risk engine v6.1.0 (Bresch and Aznar-Siguan, 2021; Aznar-Siguan and Bresch, 2019) with purpose-built pre- and post-processing workflows to produce city-level estimates of climate-attributable impacts and the discounted costs and benefits of adaptation policies. The workflow is structured as a sequential pipeline of modular computational notebooks, each with a well-defined input-output interface to the preceding and following stages. Configuration is fully externalised to city-specific YAML files, so the same analytical code operates identically across cities by reading city-specific paths, parameters, and scenario choices from the configuration layer. Input data are organised under a common directory structure and can be retrieved automatically from shared storage via a manifest-driven synchronisation utility. All intermediate and final outputs are written to documented locations so individual pipeline stages can be re-run in isolation without reprocessing upstream steps.

The spatial domain for each application is the Functional Urban Area boundary from the Global Human Settlement Layer (GHS-FUA) database. All raster layers are harmonised to a common reference grid at approximately 100 m horizontal resolution, which is the native grid of the UrbClim urban climate product used for the heat hazard. The analysis covers four target years - 2020, 2030, 2040, and 2050 - representing a current baseline and a trajectory of future climate and socioeconomic change. CLIMADA objects (Hazard, Centroids, Exposures, ImpactFuncSet) are used throughout for mortality impact calculations; centroid coordinates are stored in WGS84 (EPSG:4326) as required by CLIMADA, while all raster operations are performed in EPSG:3035 to preserve area-conserving properties.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

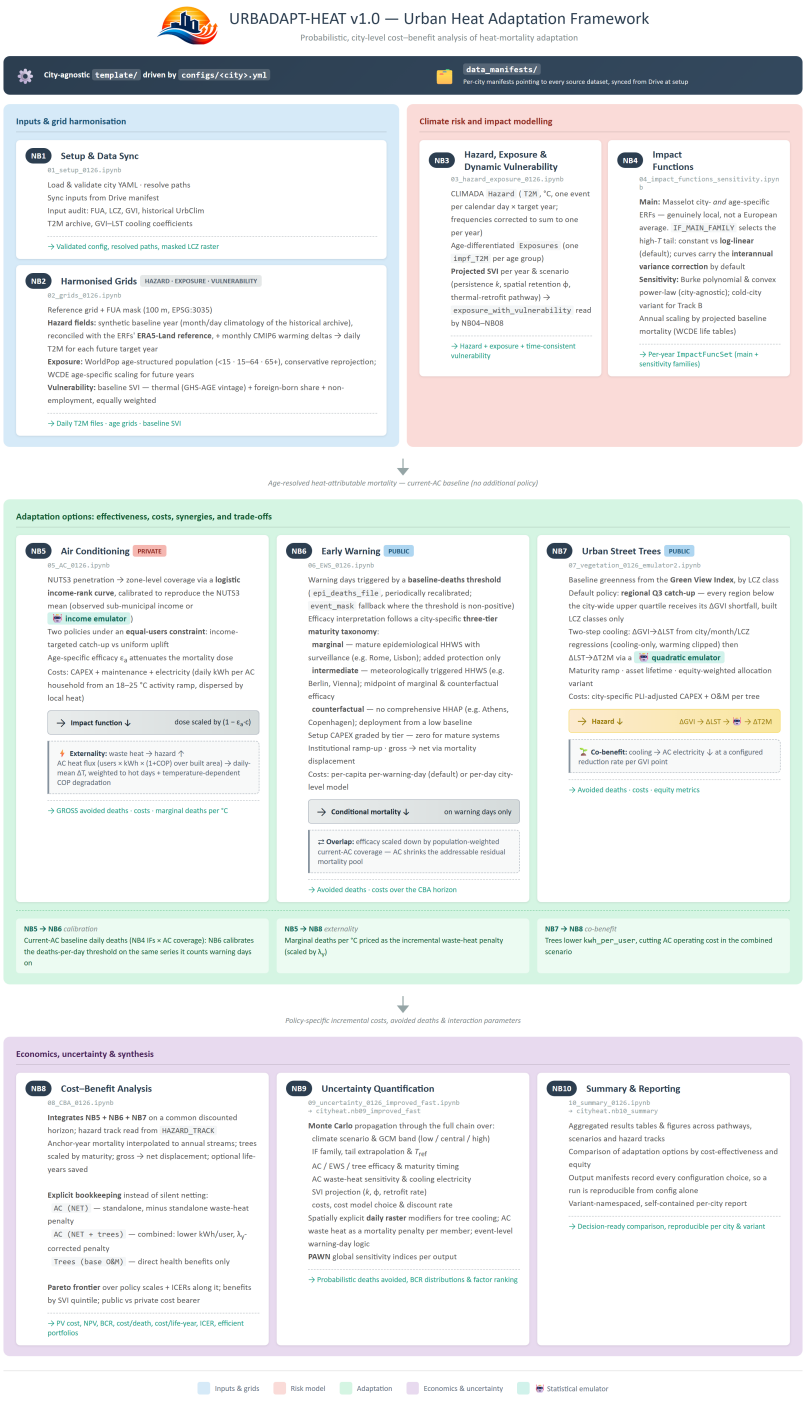


Figure 1. Overview of the URBADAPT-HEAT workflow. City-specific configuration files drive data harmonisation to a common 100 m grid, construction of CLIMADA hazard, exposure, and vulnerability objects, and estimation of baseline heat-attributable mortality. Three adaptation pathways are then evaluated in parallel: air conditioning, early warning systems, and urban trees. Their mortality benefits, costs, interactions, and equity implications are integrated in a 25-year cost-benefit analysis and portfolio optimisation module.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.2 URBADAPT-HEAT

URBADAPT-HEAT is the heat-specific implementation of the URBADAPT framework. It operationalises the general CLIMADA risk chain - hazard, exposure, vulnerability, and impact - through a set of components tailored to outdoor ambient heat, heat-attributable mortality, and three categories of public and private adaptation (Fig. 1). The modular structure is designed so that each component can be updated, replaced, or extended independently. The hazard module produces daily temperature fields for each target year on the common reference grid. The exposure module provides age-stratified population counts on the same grid. The vulnerability module constructs a composite social vulnerability index projected consistently with the chosen socioeconomic scenario. The adaptation module evaluates three policy pathways - air conditioning (AC), urban vegetation (street trees), and early warning systems (EWS) - and quantifies their mortality benefits and associated costs. A cost-benefit analysis (CBA) module integrates these outputs across a 25-year policy horizon with a consistent discounting framework. An uncertainty module propagates key structural and parametric uncertainties through the full chain.

3.3 Implementation, licence, and software requirements

URBADAPT-HEAT is implemented as a sequence of configurable Jupyter notebooks supported by the `cityheat` helper package, with each city described by a single YAML configuration file and an accompanying data manifest. The source code is available at https://github.com/URBADAPT/URBADAPT-HEAT (the exact version used in this paper, v1.0.1, is archived at https://doi.org/10.5281/zenodo.24165564, Aboudrar-Méda and Falchetta, 2026a) under the Creative Commons CC0 1.0 Universal public-domain dedication; the user manual is maintained as the repository wiki and the model website at https://urbadapt.github.io/ documents the notebook structure and contents [1]. Archived versions and input data are listed under Code and data availability.

3.4 Hazard component

The heat hazard is represented by daily mean near-surface air temperature (T_{2m}). Historical temperature fields are taken from the UrbClim urban-climate simulations distributed through the PROVIDE Climate Risk Dashboard and the VITO data service (De Ridder et al., 2015). The input archive covers 2008-2017 at a nominal spatial resolution of 100 m, and is provided in the European LAEA coordinate reference system (EPSG:3035). Input temperatures are converted from kelvin to degrees Celsius before the hazard is constructed. All temperature rasters are cropped and masked to the city-specific Global Human Settlement Layer Functional Urban Area (GHS-FUA) geometry specified

[1] The model runs under Python 3.10.20 with CLIMADA 6.1.0 as the underlying risk engine. Every direct dependency is pinned in `environment.yml`, including NumPy 1.26.4, pandas 2.1.4, SciPy 1.15.2, xarray 2025.6.1, GeoPandas 1.1.1, rasterio 1.4.3, Fiona 1.10.1, PyProj 3.7.1, scikit-learn 1.7.2 and SALib 1.5.1, and the complete solved environment of the runs reported here is archived as `environment.lock.yml`. The income emulator is a separate Python component with its own pinned requirements (Python 3.14, LightGBM 4.6.0 and scikit-learn 1.9.0). The climate emulator is not re-estimated at runtime: its coefficients are distributed as a fitted bundle, estimated in R 4.5.1 with `fixest` 0.13.2, whose fitting scripts are archived with the code.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

in the data manifest and city configuration.

Because continuous daily urban-climate simulations are not available for all projection years, URBADAPT-HEAT constructs a synthetic baseline daily temperature sequence from the 2008-2017 archive. The default baseline `climatology_mean`, assigns to each calendar day and grid cell the mean temperature for that month-day across the historical years. It therefore retains the observed spatial temperature pattern and calendar-day seasonal cycle, but does not represent interannual variability or a sub-daily temperature cycle. Three alternative baseline modes are implemented for sensitivity analysis. The `pixelwise_doy_max` mode uses the maximum temperature observed at each grid cell and calendar day across the historical years. The `domain_peak_day` mode uses the complete daily sequence from the historical year containing the hottest domain-mean day. The `warmest_summer` mode uses the complete sequence from the year with the highest June-August domain-mean temperature. These alternatives represent different assumptions about the persistence of historically hot conditions; they are not percentile transformations of individual projected years.

Reconciliation with the estimation reference. The exposure-response functions described in Sect. 3.7 are estimated against ERA5-Land city-mean daily temperature, whereas the hazard is constructed from UrbClim. UrbClim resolves the urban heat island at approximately 100 m and runs warmer than ERA5-Land through much of the distribution, by up to 2.3 °C at the 99th percentile in the warmest pilot cities and by approximately zero in the coolest. Supplying UrbClim temperatures to a curve indexed on ERA5-Land therefore displaces the whole temperature distribution relative to the city-age MMT and to the spline knots, which are themselves placed on ERA5-Land percentiles, and assigns a higher mortality risk to the same underlying conditions. URBADAPT-HEAT therefore reconciles the two references before the hazard is used, as part of the canonical configuration. Let $\bar{T}_{\mathrm{base}}(d)$ denote the domain mean of the baseline calendar on day d, let $\hat{F}_c$ denote the empirical cumulative distribution function of domain-mean daily temperature over the complete 2008-2017 UrbClim record for city c, and let Q_c^{ERA5} denote the quantile function of the ERA5-Land daily mean temperature distribution reported by Masselot et al. (2023) for the matching urban audit unit. The reference shift is

$$\delta_c(d) = Q_c^{\mathrm{ERA5}}\Big(\hat{F}_c\big(\bar{T}_{\mathrm{base}}(d)\big)\Big) - \bar{T}_{\mathrm{base}}(d), \tag{1}$$

and the corrected baseline is $T'_{\mathrm{base}}(d,x) = T_{\mathrm{base}}(d,x) + \delta_c(d)$. Three properties of this construction are deliberate. First, the shift is applied uniformly across grid cells, so each cell retains its anomaly about the city reference and the intra-urban temperature gradient on which the distributional results depend is preserved rather than compressed; quantile-mapping cell by cell would instead weaken that gradient. Second, the empirical distribution $\hat{F}_c$ is estimated on the complete observed record rather than on the smoothed baseline calendar, so the mapping reflects the actual distribution of daily temperatures rather than that of the climatology. Third, the correction is applied to the baseline calendar before the monthly deltas of Equation (2) are added, and it is estimated only on the 2008-2017 historical period; projected climate changes, and the temperature changes induced by the adaptation modules, therefore pass

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.



through unaltered.

Future temperature fields are generated with a monthly delta-change method. Monthly changes in mean near-surface air temperatures are taken from the PROVIDE scenario data for the four target years 2020; 2030, 2040, and 2050. The framework currently appraises the Current policies (`CurPol`), Gradual Strengthening of current policies (`GS`), and Shifting Pathways (`SP`) scenarios, together with a high-emissions SSP5-8.5 sensitivity pathway. For the band-aware PROVIDE inputs, a representative monthly change is first calculated for each global climate model as the mean of its estimates at the 45th and 55th percentiles. The low, central, and high warming bands are then defined as the 25th, 50th, and 75th percentiles, respectively, across the contributing climate models. These bands therefore describe inter-model spread in the monthly temperature change rather than a probabilistic confidence interval. Where only a legacy ensemble-average input is available, the workflow constructs the central trajectory only and does not infer low and high bands. The resulting annual-mean fields across the Functional Urban Area are mapped in Fig. 2, and a worked characterisation of the hazard fields for Rome under `CurPol` is given in Fig. A4.

For target year y, calendar month m, warming band b, and grid cell x, projected daily temperature is calculated as

$$T_{y,b}(d,x) = T_{\text{base}}(d,x) + \Delta T_{y,m(d),b}, \tag{2}$$

where $T_{\text{base}}(d,x)$ is the selected historical baseline sequence and $\Delta T_{y,m(d),b}$ is the corresponding monthly climate model temperature change. The same monthly increment is applied to every grid cell. Consequently, the method preserves the within-city spatial temperature contrasts of the baseline during each month, while allowing the magnitude and seasonality of warming to vary across scenarios, years, and uncertainty bands. It does not simulate changes in the spatial structure of the urban heat island, land cover, or urban morphology. The resulting daily rasters are converted to CLIMADA `Hazard` objects, with one event for each day and temperature intensity expressed in degrees Celsius. Event frequencies are set to the reciprocal of the number of days in the corresponding target year, so that the frequencies of the daily events within each target year subset sum to one. This convention allows the impacts of the daily events to be aggregated into an expected annual impact for each target year.

An optional threshold-based event representation is also implemented for diagnostic and sensitivity applications. In this mode, events are restricted using a threshold derived from the historical temperature distribution. This alternative representation is disabled in the canonical URBADAPT-HEAT v1 configuration: the main results use the complete daily temperature sequence and the continuous heat-mortality impact functions described in Sect. 3.7. Table 1 summarises the inputs and transformations used to construct the hazard component.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 1. Main inputs and transformations used to construct the URBADAPT-HEAT temperature hazard.

Component	Source or setting	Treatment in URBADAPT-HEAT
Historical temperature	UrbClim daily mean T_{2m}, 2008-2017, nominally 100 m	Converted from kelvin to degrees Celsius and aligned in EPSG:3035
Spatial domain	City-specific GHS-FUA geometry	Temperature rasters are cropped and masked to the configured functional urban area
Default baseline	`climatology_mean`	Grid-cell-specific month-day mean across 2008-2017
Alternative baselines	`pixelwise_doy_max`, `domain_peak_day`, and `warmest_summer`	Configurable sensitivity cases representing alternative historical heat conditions
Reference reconciliation	Masselot ERA5-Land daily mean distribution for the matching urban audit unit	Quantile mapping of the domain-mean baseline onto the ERA5-Land distribution, applied uniformly across grid cells so that cell-level anomalies are preserved
Projection method	PROVIDE monthly T_{2m} changes	Monthly delta added uniformly to the baseline temperature field
Climate uncertainty	Low, central, and high warming bands	25th, 50th, and 75th percentiles across per-model monthly estimates, where available
Target years	2020, 2030, 2040, and 2050	Complete daily calendar generated for each target year
CLIMADA representation	One hazard event per day	Event frequency equals the reciprocal of the number of days in the target year

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

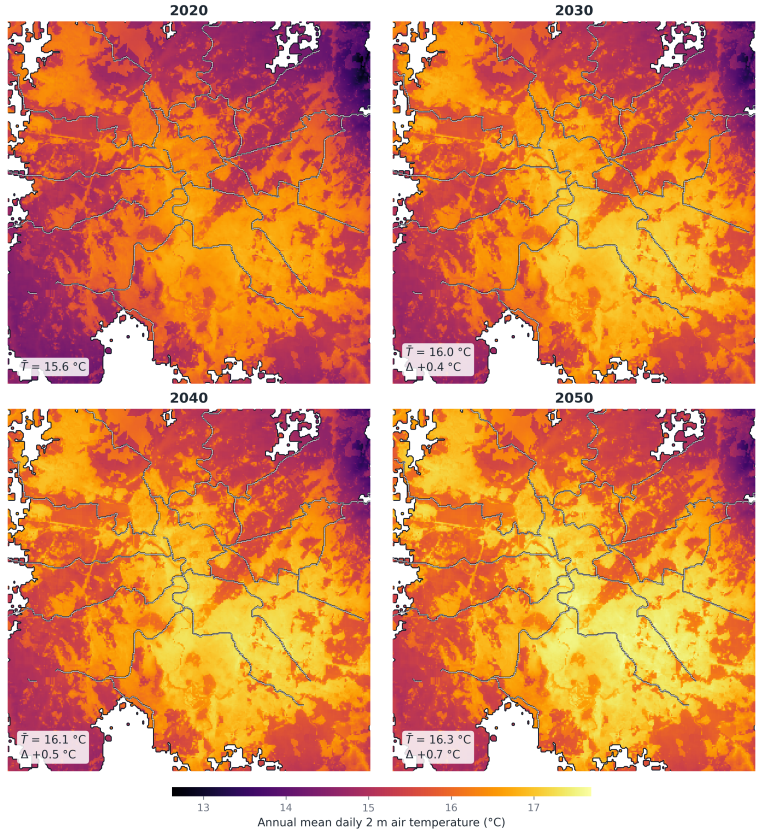


Figure 2. Spatial distribution of annual mean daily 2 m air temperature (T2M) across the Functional Urban Area for the 2020 synthetic baseline and for 2030, 2040, and 2050 under the PROVIDE Current Policies (`CurPol`) pathway, central uncertainty band. Each panel shows the full 100 m-resolution UrbClim grid masked to the FUA boundary, drawn on a shared temperature scale, with the sub-municipal administrative boundaries (*municipi*) of Rome overlaid as white lines (boundaries © OpenStreetMap contributors, ODbL 1.0); panel annotations report the FUA-mean temperature $\bar{T}$ and its change Δ relative to the 2020 baseline. Under the current monthly-delta perturbation approach, warming increments are applied domain-wide, so baseline intra-city temperature contrasts are largely preserved over time, with urban cores remaining systematically warmer than peri-urban and vegetated areas.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.5 Exposure component

Population exposure is represented on the same approximately 100 m grid as the temperature hazard and is divided into the three age groups used by the URBADAPT-HEAT impact functions: younger than 15 years, 15-64 years, and 65 years and older. The gridded population inputs are taken from the WorldPop R2025A age-structure product. The native five-year age classes are summed into the three model age groups, and the population-count rasters are reprojected to the hazard grid using sum-preserving resampling. This treatment preserves population totals when source and target grid cells differ in size or alignment. WorldPop provides the spatial age distribution used directly for 2020 and 2030. Population in 2040 and 2050 is constructed by retaining the 2030 WorldPop spatial distribution within each broad age group and scaling every grid cell by the corresponding national age-specific change from the Wittgenstein Centre Data Explorer (WCDE) demographic projections (KC et al., 2024). For age group a, city c, target year $y \in \{2040, 2050\}$, and grid cell x, population is calculated as

$$P_{a,c,y}(x) = P^{\mathrm{WP}}_{a,c,2030}(x) \frac{N^{\mathrm{WCDE}}_{a,n(c),y,s}}{N^{\mathrm{WCDE}}_{a,n(c),2030,s}}, \tag{3}$$

where $P^{\mathrm{WP}}_{a,c,2030}(x)$ is the 2030 WorldPop population, $n(c)$ denotes the country containing city c, N^{WCDE} is the national WCDE population of the corresponding age group, and s denotes the selected demographic scenario. The implemented scenarios comprise SSP1-SSP5 and the SSP2 zero-migration and double-migration variants. Population in 2020 and 2030 is taken directly from WorldPop and is therefore invariant across these demographic scenarios; scenario-dependent scaling begins in 2040.

This approach projects national demographic and ageing trends onto the 2030 within-city population pattern. It does not model subsequent intra-urban migration, spatially explicit urban expansion, neighbourhood redevelopment, or age-specific residential relocation within the functional urban area. The 2040 and 2050 grids should therefore be interpreted as scenario-conditioned population exposures on a fixed 2030 spatial template rather than as spatially explicit urban-growth projections.

The harmonised population grids are converted to CLIMADA `Exposures` objects, with one record for each populated grid cell and age group. Each record contains its population count, spatial coordinates, age-group impact-function identifier, and, where applicable, its social vulnerability attributes. The common computational city mask is defined by the intersection of the configured GHS-FUA, valid temperature coverage, and grid cells with positive total population in the 2020 WorldPop layer. This ensures that the hazard, exposure, and vulnerability components are evaluated over the same spatial support. Table 2 summarises the inputs and transformations used to construct the exposure component, and Fig. 3 maps the resulting age-disaggregated exposure for the 2020 baseline.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 2. Main inputs and transformations used to construct population exposure in URBADAPT-HEAT.

Component	Source or setting	Treatment in URBADAPT-HEAT
Population product	WorldPop R2025A age-structure rasters, approximately 100 m	Population-count grids are aligned to the hazard grid using sum-preserving resampling
Age groups	Native five-year WorldPop age classes	Aggregated to < 15, 15-64, and 65+ years
2020 exposure	WorldPop 2020	Used directly after spatial and age-group harmonisation
2030 exposure	WorldPop 2030	Used directly and retained as the spatial template for later years
2040 and 2050 exposure	WorldPop 2030 spatial pattern and WCDE national projections	Each age-group grid is multiplied by its national age-specific growth factor relative to 2030
Demographic scenarios	SSP1-SSP5, SSP2 zero migration, and SSP2 double migration	Affect projected population in 2040 and 2050; 2020 and 2030 remain scenario-invariant
Computational mask	GHS-FUA, valid temperature coverage, and positive 2020 population	Defines the common spatial support of the hazard, exposure, and vulnerability components
CLIMADA representation	One record per grid cell and age group	Population count is the exposure value and the age group selects the corresponding impact function

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

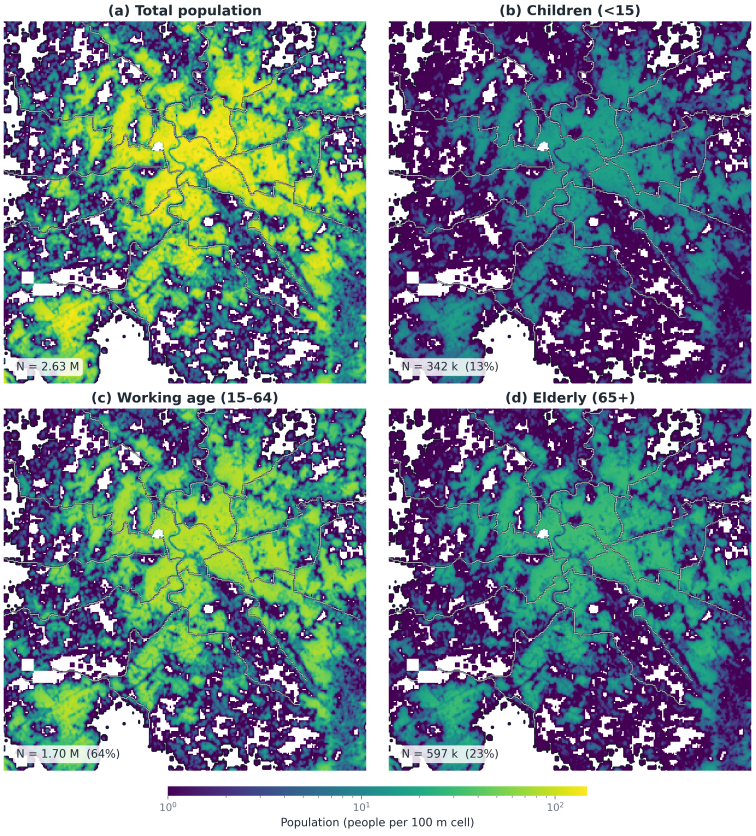


Figure 3. Exposed population across the Functional Urban Area for the 2020 baseline, on the 100 m reference grid: (a) total population (all age groups); (b) children under 15; (c) working-age adults (15-64); (d) elderly (65+). All panels share a logarithmic colour scale (people per 100 m cell, floored at 1); panel annotations report the population count N and, for the age groups, its share of the total. The dark contour marks the GHS-FUA boundary that defines the spatial analysis domain, and the white lines the sub-municipal *municipi* of Rome. The concentric gradient reflects the population density structure of a large European FUA; the same grid underlies all three age-group `Exposures` objects passed to CLIMADA.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.6 Vulnerability component

URBADAPT-HEAT constructs a spatially explicit Social Vulnerability Index (SVI) on the common analysis grid. The SVI is a model-constructed index used to describe the spatial distribution of three vulnerability proxies: the dominant age of the built stock, the share of residents in Eurostat's "place of birth elsewhere" category, and the non-employment share of the working-age population. Higher values indicate higher modelled vulnerability. The index is intended for distributional analysis and policy targeting: it is not an empirically estimated modifier of the heat-mortality exposure-response relationship.

Baseline components. The thermal component is derived from the 100 m GHS-AGE R2025A raster (Uhl et al., 2026). GHS-AGE does not report individual building construction years. For each grid cell, it estimates the epoch in which more than 50 % of the built-up surface observed in 2020 was first present. The 5-year product encodes the resulting dominant built-stock epoch using class codes 1-10, from the oldest class, present by 1975, to the most recent 2020 class. URBADAPT-HEAT reprojects these categorical data to the reference grid using nearest-neighbour resampling and defines baseline thermal vulnerability as:

$$z_{\mathrm{th},0}(x) = \frac{10 - A(x)}{9}, \qquad A(x) \in \{1, \ldots, 10\}, \tag{4}$$

where $A(x)$ is the GHS-AGE class at grid cell x. Thus, the oldest class receives a score of 1 and the most recent class a score of 0. This inversion is a modelling assumption that uses dominant built-stock age as a proxy for thermal performance; GHS-AGE itself does not observe insulation, ventilation, indoor temperature, or access to cooling.

The two social components are derived from the Eurostat 2021 Census population grid (Eurostat, 2025). The source grids have a spatial resolution of 1 km and contain extensive population counts. They are aligned by the approximately 100 m reference grid using sum-preserving reprojection before ratios are calculated. This alignment does not create additional social information within the original 1 km cells.

The first social proxy is:

$$r_{\mathrm{oth},0}(x) = \frac{C_{\mathrm{OTH}}(x)}{C_{\mathrm{T}}(x)}, \tag{5}$$

where C_{OTH} is the population in Eurostat's "place of birth elsewhere" category and $C_{\mathrm{T}}$is the total usual-resident population. For the four pilot countries, this category is distinct from Eurostat's separate category for persons born in another EU Member State. Consequently, $r_{\mathrm{oth},0}$ must not be interpreted as the total foreign-born share.

The second social proxy is the non-employment share

$$r_{\mathrm{ne},0}(x) = 1 - \frac{C_{\mathrm{EMP}}(x)}{C_{\mathrm{Y1564}}(x)}, \tag{6}$$

where C_{EMP} is the number of employed persons and C_{Y1564} is the population aged 15-64 years. This measure includes all working-age persons who are not recorded as employed and is therefore broader than the conventional unemployment rate.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

For the operational dynamic SVI, each social proxy $r_{j,0}$ is transformed to a baseline-anchored score using its within-city 5th and 95th percentiles:

$$z_{j,0}(x) = \mathrm{clip}\left[\frac{r_{j,0}(x) - Q_{0.05,j}}{Q_{0.95,j} - Q_{0.05,j}}, 0, 1\right], \qquad j \in \{\mathrm{oth}, \mathrm{ne}\}. \tag{7}$$

The quantile bounds are calculated once from the baseline year. Projected years are not re-ranked, which preserves comparability of component values over time. Percentile-rank versions of the two components are also retained by the software for legacy diagnostics, but they are not used in the operational dynamic SVI.

The three components are combined using equal weights. Let $\mathcal{J}_y(x)$ denote the set of components with finite values at location x in year y. The composite index is:

$$\mathrm{SVI}_y(x) = \frac{\sum_{j \in \mathcal{J}_y(x)} w_j z_{j,y}(x)}{\sum_{j \in \mathcal{J}_y(x)} w_j}, \qquad w_{\mathrm{th}} = w_{\mathrm{oth}} = w_{\mathrm{ne}} = \frac{1}{3}. \tag{8}$$

In the four pilot-city configurations, a grid cell must contain at least two valid components for an SVI value to be calculated. The weights are renormalised over the available components. When the SVI is subsequently attached to a CLIMADA exposure file, any remaining missing values within the computational city mask are replaced by the within-city median SVI so that all retained exposure records have a finite stratification value.

Dynamic projection of the social components. The source social rasters refer to the 2021 Census, while the GHS-AGE series ends in 2020. The vulnerability layer stored as the 2020 model output is therefore the baseline component layer attached to the 2020 population exposure; it should not be interpreted as a separately observed set of 2020 social indicators. The dynamic projection engine uses 2021 as its internal social-component baseline year and produces layers for 2020, 2030, 2040, and 2050. The resulting trajectory for Rome is shown in Fig. A3.

Near-term changes in the mean social-component scores are anchored to NUTS3-level series from the DRMKC Risk Data Hub (Antofie et al., 2022). Net migration is used as a trend proxy for the place-of-birth component, and the employment rate is used with a negative sign for the non-employment component. For component j, the modelled change between 2021 and 2030 is

$$\Delta_j^{\mathrm{D}} = a_j \lambda_j^{\mathrm{D}} \left(D_{j,2030} - D_{j,2021}\right), \tag{9}$$

where $a_{\mathrm{oth}} = 1$, $a_{\mathrm{ne}} = -1$, $\lambda_{\mathrm{oth}}^{\mathrm{D}} = 0.04$, and $\lambda_{\mathrm{ne}}^{\mathrm{D}} = 0.08$. These λ values are configured response coefficients: the DRMKC indicators are proxy trend anchors rather than direct projections of the two baseline Census ratios. For years between 2021 and 2030, the model linearly interpolates between the baseline component mean and the 2030 mean.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Long-run changes are linked to national projections of the Global Data Lab Global Vulnerability Index (GDL-GVI) under the Shared Socioeconomic Pathways (Huisman et al., 2025). This GVI is a national socioeconomic vulnerability index and is unrelated to the Green View Index used in the street-tree module. The source provides SSP1, SSP2, and SSP3 trajectories. URBADAPT-HEAT maps SSP2-ZM and SSP2-DM to SSP2, SSP4 and SSP3, and SSP5 to SSP2 because separate GDL-GVI trajectories for these scenarios are not available in the input database.

For $2030 < y \leq 2050$, the target city mean for social component j is

$$\mu_{j,y} = \mu_{j,0} + \Delta_j^{\mathrm{D}} + \lambda_j^{\mathrm{G}} \left(G_{s,y} - G_{s,2030}\right) + (k_j - 1)\Delta_j^{\mathrm{D}} \frac{y - 2030}{2050 - 2030}, \tag{10}$$

where $\mu_{j,0}$ is the baseline city mean, $G_{s,y}$ is the national GDL-GVI value for scenario s, and $\lambda_{\mathrm{oth}}^{\mathrm{G}} = 0.35$ and $\lambda_{\mathrm{ne}}^{\mathrm{G}} = 0.50$. The parameter k_j determines how much of the DRMKC derived 2021-2030 change remains by 2050: at 2030 the complete near-term change is present, whereas at 2050 its contribution is $k_j \Delta_j^{\mathrm{D}}$.

The projected city mean is downscaled to the reference grid by retaining a fraction $\phi_{j,y}$ of the baseline intra-city anomaly:

$$z_{j,y}(x) = \mathrm{clip}\left[\mu_{j,y} + \phi_{j,y}\left(z_{j,0}(x) - \mu_{j,0}\right), 0, 1\right]. \tag{11}$$

The value of $\phi_{j,y}$ is one at the baseline and is linearly interpolated to the component-specific 2030 and 2050 configuration values. A lower value therefore produces stronger convergence towards the projected city mean. Both k_j and $\phi_{j,y}$ are configurable model-response parameters rather than independently observed future quantities.

Dynamic projection of thermal vulnerability. Thermal vulnerability follows a separate projection. For the existing built stock, the baseline score decays exponentially:

$$\widetilde{z}_{\mathrm{th},y}(x) = z_{\mathrm{th},0}(x) \exp\left[-r_c f_s (y - 2021)\right], \tag{12}$$

where r_c is the city configuration's effective annual thermal-upgrading rate and f_s is a scenario multiplier. The implemented multipliers are 1.20 for SSP1, 1.00 for SSP2 and SSP2-DM, 0.95 for SSP2-ZM, 0.80 for SSP3, 0.85 for SSP4, and 1.10 for SSP5.

The model also calculates positive population growth relative to the 2020 population grid. Cell-level growth is normalised by its within-city 95th percentile and converted to a new-build share $q_y(x)$ using the configured growth-sensitivity coefficient and maximum share. Projected thermal vulnerability is then

$$z_{\mathrm{th},y}(x) = \left[1 - q_y(x)\right] \widetilde{z}_{\mathrm{th},y}(x) + q_y(x) z_{\mathrm{new}}, \tag{13}$$

with $z_{\mathrm{new}} = 0.15$ in all four pilot-city configurations. Cells with projected population but no valid baseline thermal component are also assigned z_{new}. The value 0.15 and the population-growth response coefficients are modelling

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

assumptions; they are not measurements of the actual future thermal performance of newly constructed buildings. Year- and scenario-specific SVI rasters are written for the target years and attached to the corresponding CLIMADA exposure objects. In URBADAPT-HEAT v1.0, the SVI does not multiply population exposure or alter the Masselot impact functions. It is used to stratify mortality and adaptation benefits, construct equity diagnostics, and support the optional SVI-weighted allocation of vegetation interventions. Accordingly, differences across SVI groups reflect the spatial distributions of population, heat, and policy coverage under a common city-level epidemiological response function; they do not represent an estimated causal effect of the SVI on individual heat-mortality risk. Table 3 summarises the inputs and transformations used to construct the vulnerability component, and Fig. 4 reports the population-weighted mean scores of its individual components.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Component	Input	Operational treatment
Thermal vulnerability	GHS-AGE R2025A, 100 m	Dominant built-stock epoch inverted to a 0-1 score; nearest-neighbour alignment
Place-of-birth proxy	Eurostat Census 2021 `OTH/T`, 1 km	"Place of birth elsewhere" share; robust baseline $Q_{0.05}$-$Q_{0.95}$ scaling
Non-employment proxy	$1-$`EMP/Y_1564`, Eurostat Census 2021, 1 km	Working-age non-employment share; robust baseline $Q_{0.05}$-$Q_{0.95}$ scaling
Composite SVI	Three component scores	Equal-weighted available-component mean; at least two valid components in the four-city application
Near-term dynamics	DRMKC NUTS3 net-migration and employment-rate series	Proxy changes mapped to component means between 2021 and 2030
Long-term dynamics	National GDL-GVI trajectories for SSP1-SSP3	Scenario-dependent mean shifts from 2030 to 2050, with explicit scenario remapping
Spatial downscaling	Baseline component anomalies and $\phi_{j,y}$	Retains a configurable fraction of the baseline intra-city pattern
Thermal projection	Effective upgrading rate, population growth, and new-build score	Exponential decay for existing stock combined with a population-growth-dependent new-build share
Operational role	Year- and scenario-specific exposure attributes	Equity stratification and policy targeting; no direct mortality multiplier

Table 3. Inputs and transformations used to construct the dynamic vulnerability component.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

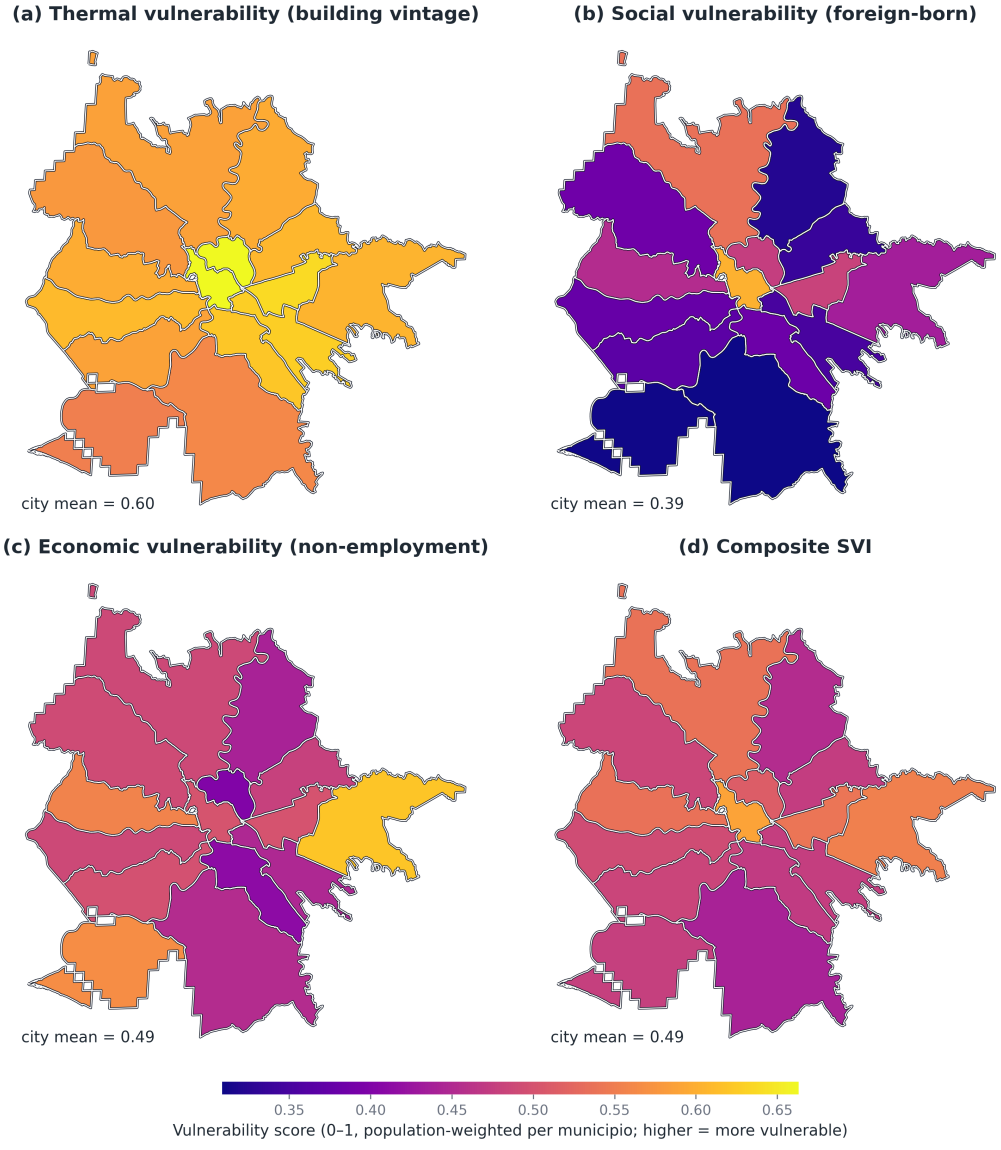


Figure 4. Population-weighted mean vulnerability scores by *Municipio* in Rome for 2050 under SSP2 (*Municipio* boundaries © OpenStreetMap contributors, ODbL 1.0): (a) thermal vulnerability; (b) the normalised Eurostat "place of birth elsewhere" proxy; (c) the normalised working-age non-employment proxy; and (d) the composite Social Vulnerability Index. The social-component panels show projected baseline-anchored scores, not raw Census percentages or observed future Census values. All panels use dimensionless scores on a 0-1 scale, with higher values indicating greater modelled vulnerability.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.7 Heat-mortality impact functions

Age-specific heat-mortality impact functions translate daily ambient temperature into a per-person probability of heat-attributable death for each exposed age group. In URBADAPT-HEAT v1.0, the canonical impact functions are derived from the city- and age-specific exposure-response curves published by Masselot et al. (2023). Masselot et al. estimated distributed-lag non-linear models (DLNMs) using daily all-cause mortality and ERA5-Land daily mean temperature for European cities through a three-stage design. First-stage DLNMs were estimated for 232 cities in the Multi-Country Multi-City (MCC) network. These estimates were subsequently combined through multivariate meta-regression using city-level characteristics and ordinary co-kriging of the residual spatial structure, producing final city-age coefficients for 854 European cities. URBADAPT-HEAT uses these final coefficients rather than the raw first-stage estimates.

The reconstruction is performed as a documented preprocessing step. The final coefficients distributed in Masselot's `coefs.csv`, the city-specific temperature distributions, the city-age minimum-mortality temperatures (MMTs), and the baseline mortality and population data are used to generate CLIMADA-compatible impact-function files. Notebook 04 then reads the resulting city-specific JSON files and promotes the configured Masselot family to the canonical downstream impact-function slot. Thus, the epidemiological coefficients are not re-estimated within URBADAPT-HEAT, and the raw mortality time series are not required to execute the model.

Masselot et al. report five age groups: 20-44, 45-64, 65-74, 75-84, and 85+. For each city c and original Masselot age group g, the cumulative log relative risk at daily mean 2 m temperature T is reconstructed from five final B-spline coefficients $\beta_{g,c,k}$:

$$\log \mathrm{RR}_{g,c}(T) = \sum_{k=1}^{5} \beta_{g,c,k}\, B_k(T;\boldsymbol{\kappa}_c) - \sum_{k=1}^{5} \beta_{g,c,k}\, B_k(\mathrm{MMT}_{g,c};\boldsymbol{\kappa}_c), \tag{14}$$

where B_k is the kth quadratic B-spline basis function and $\boldsymbol{\kappa}_c$ denotes the city-specific knot sequence. Following Masselot et al. (2023), the internal temperature knots are placed at the 10th, 75th, and 90th percentiles of the city's ERA5-Land daily mean temperature distribution, with boundary knots at the historical minimum and maximum temperatures. The curve is centred on the city-age MMT reported in Masselot's `cityage.csv`, so that $\mathrm{RR}_{g,c}(\mathrm{MMT}_{g,c}) = 1$.

The five coefficients in Equation (14) describe the cumulative exposure-response relationship over lags 0-21 days. They therefore represent the overall temperature-mortality association accumulated across the lag period, rather than separate effects for individual lag days.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

For the operational heat-only implementation, the heat-attributable fraction is defined as

$$\mathrm{AF}_{g,c}(T) = \begin{cases} \max\left[0, \dfrac{\mathrm{RR}_{g,c}(T)-1}{\mathrm{RR}_{g,c}(T)}\right], & T \geq \mathrm{MMT}_{g,c}, \\ 0, & T < \mathrm{MMT}_{g,c}. \end{cases} \tag{15}$$

The corresponding CLIMADA mean damage degree (MDD), expressed as the probability of a heat-attributable death per exposed person and event day, is

$$\mathrm{mdd}^{(0)}_{g,c}(T) = \frac{M^{\mathrm{base}}_{g,c}}{365.25}\,\mathrm{AF}_{g,c}(T), \tag{16}$$

where $M^{\mathrm{base}}_{g,c}$ is the annual per-person baseline all-cause mortality rate reported by Masselot et al. for city c and age group g. The superscript (0) denotes the reference-year curve before applying future changes in national baseline mortality.

The five Masselot age groups are aggregated to the three exposure groups used by URBADAPT-HEAT, $a \in \{<15,\ 15\text{-}64,\ 65+\}$. For the two adult groups, aggregation is performed on the per-person MDD curves using Masselot's city-age population counts:

$$\mathrm{mdd}^{(0)}_{a,c}(T) = \sum_{g \in \mathcal{G}_a} w_{g|a,c}\,\mathrm{mdd}^{(0)}_{g,c}(T), \qquad w_{g|a,c} = \frac{P_{g,c}}{\sum_{h \in \mathcal{G}_a} P_{h,c}}, \tag{17}$$

where $\mathcal{G}_{15\text{-}64} = \{20\text{-}44, 45\text{-}64\}$ and $\mathcal{G}_{65+} = \{65\text{-}74, 75\text{-}84, 85+\}$. Because Masselot's first-stage analysis excludes people younger than 20 due to insufficient mortality counts for stable estimation, the <15 impact function is set identically to zero. The URBADAPT 15-64 group is consequently represented by ages 20-64; ages 15-19 are absent from the epidemiological source data.

The aggregation in Equation (17) combines MDD values rather than fitting a new exposure-response curve. Consequently, the aggregated 15-64 and 65+ functions do not have a single newly estimated MMT: each underlying Masselot curve retains its own city-age MMT. The reconstruction was validated by applying the reconstructed relative-risk curves to the daily ERA5-Land temperature series used by Masselot et al. and comparing the resulting annual heat-attributable deaths with the published values. When the published attribution convention is reproduced, the annual totals agree to numerical precision, with percentage differences of order $10^{-12}\,\%$. The operational URBADAPT export additionally imposes the non-negative heat-only constraint in Equation (15). This changes the reconstructed annual totals only marginally: among the four pilot cities, the maximum absolute difference from the published annual total is approximately $0.064\,\%$.

Extrapolation beyond the historical temperature support. The Masselot B-spline basis is defined within each city's historical temperature support, $[T^{\min}_c, T^{\max}_c]$. Simulated temperatures can exceed $T^{\max}_c$ under future climate conditions, so URBADAPT-HEAT implements two explicit extrapolation rules before aggregating the original Masselot age groups.

The canonical *log-linear-tail* specification (`masselot_tail`) extends log relative risk above $T^{\max}_c$:

$$\log \mathrm{RR}_{g,c}(T) = \log \mathrm{RR}_{g,c}(T^{\max}_c) + s_{g,c}\,(T - T^{\max}_c), \qquad T > T^{\max}_c, \tag{18}$$

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

where $s_{g,c}$ is the non-negative finite-difference slope of the reconstructed log-relative-risk curve over approximately the final 1 °C of its historical support. If the estimated terminal slope is negative, it is set to zero. The extrapolated relative risk is subsequently converted to AF and MDD using Equations (15) and (16).

The alternative *constant-tail* specification (`masselot`) holds MDD at its value at $T_c^{\max}$ for all higher temperatures. It is retained as a within-Masselot structural sensitivity representing a conservative saturation assumption. The two variants make the extrapolation choice explicit; they should not be interpreted as guaranteed statistical bounds on the unknown response outside the observed temperature range. The relevance of the extrapolation assumption depends on how frequently and by how much a city's simulated temperature distribution exceeds its historical support.

After extrapolation, the curves are interpolated to a common 0-45 °C grid at 0.5 °C intervals and clipped to the $[0, 1]$ MDD range required by CLIMADA.

Interannual variance correction. The default `climatology_mean` baseline assigns to each calendar day the mean temperature across the historical years, so the model evaluates the response at the mean of the observed years. The quantity of interest is instead the burden of an average year, which is the mean of the responses to the individual years. Because the heat limb of the exposure-response relationship is convex, Jensen's inequality makes the first strictly smaller than the second, and the gap widens with interannual variance. The consequence is most severe in cool cities, where the climatological summer peak can fall below the city-age MMT even though individual observed years exceed it: the climatological burden then collapses towards zero while the true expectation does not.

Carrying every observed year through the hazard and averaging the resulting impacts would remove the discrepancy by construction, but it multiplies the size of the hazard object and the cost of the impact calculation by the number of observed years. URBADAPT-HEAT instead implements the correction on the impact function, where it is algebraically equivalent and free at run time. Writing f for a reconstructed per-person MDD curve and $\mathcal{E}_c$ for the pooled sample of daily temperature anomalies about the day-of-year mean in city c, the variance-corrected curve is

$$\widetilde{\mathrm{mdd}}^{(0)}_{a,c}(T) = \frac{1}{|\mathcal{E}_c|} \sum_{\epsilon \in \mathcal{E}_c} \mathrm{mdd}^{(0)}_{a,c}(T + \epsilon), \tag{19}$$

evaluated on the same temperature grid, so that $\widetilde{\mathrm{mdd}}^{(0)}$ is a drop-in replacement for $\mathrm{mdd}^{(0)}$ and the hazard, the event count, the event frequencies and all downstream components are unchanged. Values beyond the grid are held flat, which reproduces the constant-tail convention for the `masselot` family; for the `masselot_tail` family the log-linear extrapolation is already embedded in the curve being convolved and is therefore preserved.

Two aspects of $\mathcal{E}_c$ are constrained by the implementation. The anomalies are pooled across days and years into a single sample, and the resulting correction is a single curve per city and age group, because CLIMADA associates impact functions with exposure points rather than with hazard events, so a month-varying curve cannot be expressed. The sample is restricted to June-August. This is not a concession to that constraint but an improvement on it: daily anomalies are strongly heteroscedastic, and pooling the full calendar applies summer-sized excursions to mild days, which in a cool city manufactures exceedance of an MMT that the season never approaches. Restricting the sample to

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

the warm season loses nothing, because the modelled response outside it is approximately zero. Both this correction and the reference reconciliation of Equation (1) are enabled in the canonical configuration and are jointly necessary (Sect. 3.12); each can be disabled independently, the reference correction in the city configuration and the variance correction by selecting the uncorrected impact-function family, which is how the attribution reported there was obtained. Evaluated against the exact per-year calculation, the warm-season single curve attains a median error of 1.6 %, against 2.0 % for twelve month-specific curves and 6.1 % for an annually pooled sample (Sect. 3.12).

Future baseline-mortality scaling. The shape and city-age centring of the Masselot curves are retained over time, while their MDD magnitude is scaled to represent changes in age-specific national baseline all-cause mortality. For city c, age group $a \in \{<15, 15\text{-}64\}$, and target year y:

$$\mathrm{mdd}_{a,c,y}(T) = s^{\mathrm{mort}}_{a,c,y}\,\mathrm{mdd}^{(0)}_{a,c}(T), \qquad s^{\mathrm{mort}}_{a,c,y} = \frac{M^{\mathrm{WCDE}}_{a,n(c),y}}{M^{\mathrm{WCDE}}_{a,n(c),2020}}, \tag{20}$$

where $M^{\mathrm{WCDE}}_{a,n(c),y}$ is the population-weighted national mortality rate for the country $n(c)$ containing city c, derived from WCDE v3 age-specific survival ratios and population projections. The 65+ group is treated differently, because its crude death rate already contains the projected shift of the group towards older ages, whereas the heat response steepens with age. For 2030, 2040, and 2050 its curve is recomposed from the Masselot curves of the 65-74, 75-84, and 85+ sub-bands g:

$$\mathrm{mdd}_{65+,c,y}(T) = \mathrm{clip}\left[\sum_g \frac{w^0_{g,c}\,\pi_{g,n(c),y}}{\sum_{g'} w^0_{g',c}\,\pi_{g',n(c),y}}\, s^{\mathrm{mort}}_{g,n(c),y}\,\mathrm{mdd}^{(0)}_{g,c}(T), 0, 1\right], \tag{21}$$

where $w^0_{g,c}$ is the population weight of sub-band g in the city's Masselot source period, which defines the 2020 65+ curve, $\pi_{g,n(c),y}$ is the ratio of the sub-band's projected share of the national 65+ population in year y to its 2020 share, and $s^{\mathrm{mort}}_{g,n(c),y}$ is the sub-band's national death-rate ratio, defined as in Equation (20). Each sub-band curve is thus scaled by its own death-rate trend and weighted by its projected share of the 65+ population, so ageing within the group enters once, through the weights, and the crude 65+ death-rate trend, which already contains it, is not applied in addition. If the three sub-bands shared one heat-attributable fraction and the city's age structure and death rates matched the national ones, Equation (21) would reproduce the single 65+ scalar of Equation (20); it departs from it because the response of the 85+ group, whose share grows, is steeper. In 2050 the recomposed 65+ curve lies 6 % above the single-scalar curve in Rome and Lisbon, 11 % in Athens, and 13 % in Copenhagen, where the projected 85+ share of the national 65+ population roughly doubles. The 2020 curves and the curves of the younger groups are unchanged. The scaling factors are calculated for 2020, 2030, 2040, and 2050 under the selected demographic SSP and stored in `baseline_heat_scaling_{city}.csv` and, for the 65+ sub-bands, `baseline_heat_subband_{city}.csv`. They are subsequently embedded in the year-specific Masselot JSON files consumed by Notebook 04. In the canonical four-city application, the demographic scaling follows SSP2. Climate-policy scenarios alter the temperature hazard and do not independently modify the impact-function shape.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Air-conditioning attenuation. Air-conditioning coverage is applied downstream as an attenuation of the year-specific Masselot MDD. For age group a, year y, and grid cell x, the effective MDD under local AC coverage fraction $c_y(x)$ is:

$$\mathrm{mdd}^{\mathrm{eff}}_{a,c,y}(T,x) = \mathrm{mdd}_{a,c,y}(T)\,\mathrm{clip}\left(1-\varepsilon_a c_y(x), 0, 1\right), \tag{22}$$

where ε_a is the age-specific efficacy of AC in reducing heat-mortality risk among covered residents. The central configuration uses $\varepsilon_{<15} = 0.20$, $\varepsilon_{15\text{-}64} = 0.30$, and $\varepsilon_{65+} = 0.40$, following the structure of the epidemiological evidence synthesised by Sera et al. (2020). Alternative efficacy configurations are retained for sensitivity analysis. Because the canonical Masselot MDD for < 15 is zero, the corresponding AC efficacy has no effect on canonical heat-attributable deaths in this age group.

Alternative impact-function families. The polynomial and power-law functions based on Burke et al. (2025) are retained as alternative structural sensitivity families. They do not enter the deterministic canonical run, which uses the Masselot-log-linear-tail specification. The Burke functions differ structurally from Masselot because they impose a parametric temperature-response form and, for the standard-city implementation, use a fixed reference temperature of 20 °C rather than city-age-specific MMTs. They also retain a non-zero response for the < 15 group. The four families are compared for Rome, by age group, in Fig. A1. Fig. 5 shows the canonical functions together with the projected annual heat-attributable deaths for Rome.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

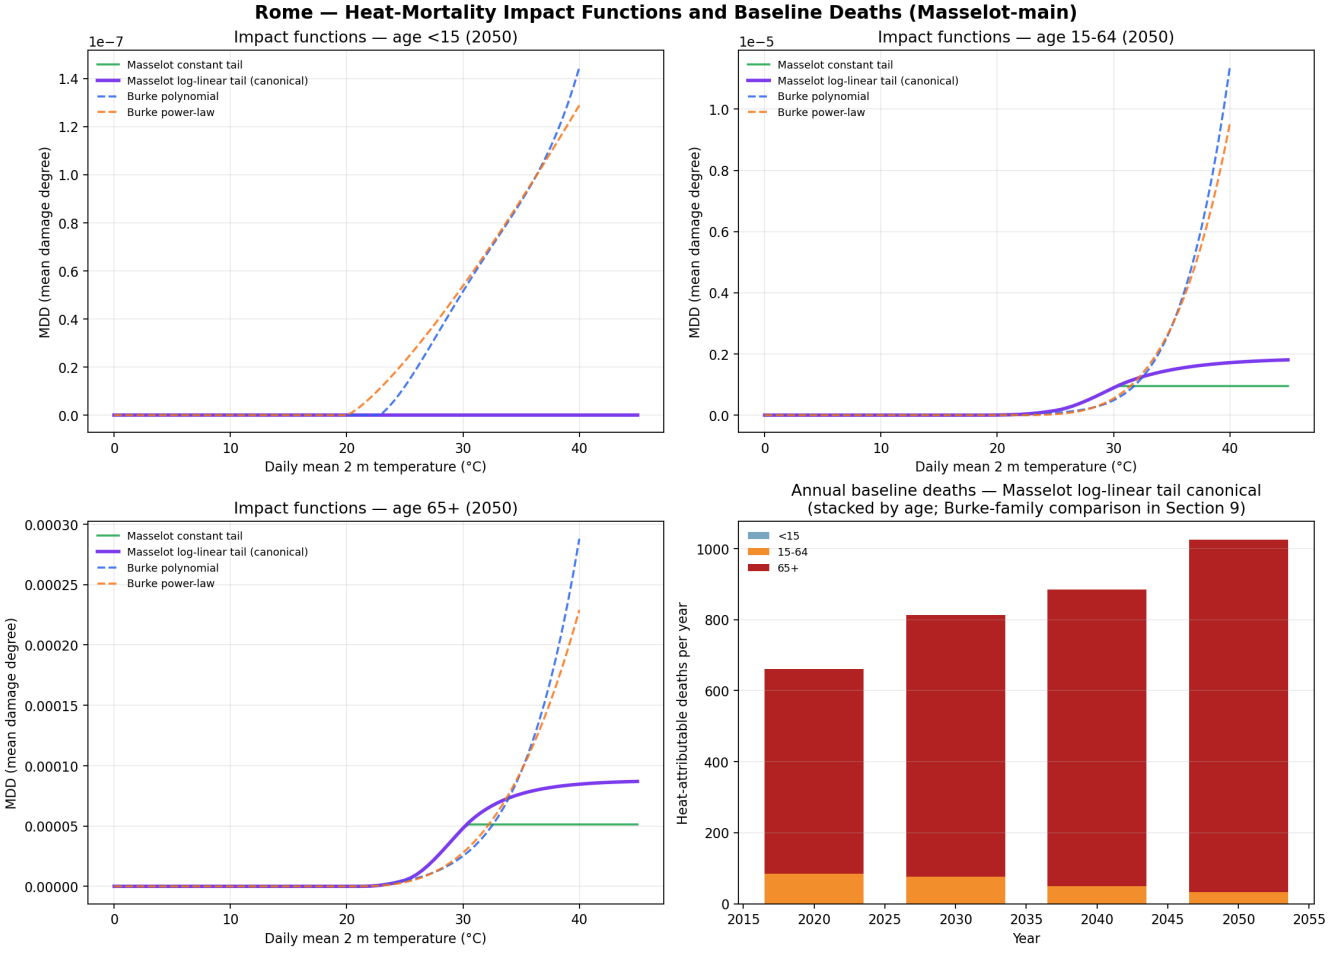


Figure 5. Heat-mortality impact functions and projected annual heat-attributable deaths for Rome. Panels (a)-(c) show the age-specific mean damage degree (MDD; probability of heat-attributable death per exposed person and event day) as a function of daily mean 2 m temperature in 2050. The canonical Masselot log-linear-tail curve is shown together with the Masselot constant-tail extrapolation sensitivity and the Burke polynomial and power-law structural sensitivity families. The Masselot < 15 curve is identically zero because people younger than 20 are excluded from the underlying epidemiological analysis. Panel (d) shows annual heat-attributable deaths at the four target years under the canonical Masselot log-linear-tail functions, before attenuation by the AC and EWS adaptation modules. The difference between the two Masselot tail specifications arises only above Rome's historical maximum daily mean temperature.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.8 Emulation components

3.8.1 Climate emulator for rapid variable translation

URBADAPT-HEAT uses a pooled statistical emulator to represent the relationship between daily maximum land-surface temperature (LST) and daily mean 2 m air temperature (T2M) as a function of the current temperature state, Green View Index (GVI), urban morphology, and Local Climate Zone (LCZ). The coefficient bundle deployed in the present workflow (`emulator_bundle_pooled_all_quadratic_holdout_safe.json`) is estimated from 33,391,358 point-day observations across 130 cities over 2008-2017. Observation locations correspond to the available GVI point sample and are restricted to built LCZ classes 1-10.

The emulator is estimated in R using the `fixest` package. Let $T^c_{it} = T2M_{it} - \overline{T}_{\text{train}}$ denote daily mean T2M centred on the pooled training mean, which equals 17.0117 °C in the deployed bundle. The fitted model is

$$\begin{aligned}\text{LST}_{it} \sim\ & \beta_1 T^c_{it} + \beta_2 \left(T^c_{it}\right)^2 + \beta_3 g_i + \beta_4 \log(1 + h_{b,i}) + \beta_5 \log(1 + v_{b,i}) \\ & + \beta_6 \log(1 + \rho_{\text{pop},i}) + \beta_7 w_i + \beta_8 z_i \\ & + \gamma_1 \left(T^c_{it} \times g_i\right) + \gamma_2 \left(T^c_{it} \times w_i\right) + \gamma_{3,\text{LCZ}(i)} T^c_{it} \;\Big|\; \alpha_{\text{month}(t)} + \alpha_{\text{LCZ}(i)},\end{aligned} \tag{23}$$

where g_i is GVI rescaled from its original 0-100 scale to 0-1; $h_{b,i}$ and $v_{b,i}$ are building height and built volume from the Global Human Settlement Layer; $\rho_{\text{pop},i}$ is population density; w_i is the water-body covariate; z_i is terrain elevation; and $\text{LCZ}(i)$ is the Local Climate Zone class. Month and LCZ fixed effects account for common seasonal and urban-form differences, while standard errors are clustered by city-specific observation point. A single pooled coefficient bundle is consequently applied across cities without city-specific refitting.

For the vegetation module, the emulator is used to translate a prescribed LST change into a T2M-equivalent change rather than to predict the absolute LST level. Holding the time-invariant covariates constant, the local slope of the fitted LST-T2M relationship is

$$b_{it} = \beta_1 + 2\beta_2 T^c_{it} + \gamma_1 g_i + \gamma_2 w_i + \gamma_{3,\text{LCZ}(i)}. \tag{24}$$

For a prescribed ΔLST_{it}, notebook 07 solves

$$\Delta\text{LST}_{it} = \beta_2 \left(\Delta\text{T2M}_{it}\right)^2 + b_{it} \Delta\text{T2M}_{it} \tag{25}$$

and selects the quadratic root closest to the corresponding linear approximation , $\Delta\text{LST}_{it}/b_{it}$. The deployed vegetation calculation supplies the baseline rasterised GVI and LCZ class at each grid cell. No aligned water raster is currently passed to this translation step, so the water-interaction term is evaluated at zero. Invalid roots, non-positive local slopes, sign-inconsistent translations, or absolute changes exceeding 10 °C are replaced using the median positive local slope.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Transfer performance is evaluated by refitting the model while holding out one complete city at a time. Across the 130 held-out-city evaluations, the unweighted mean predictive R^2 is 0.718, the mean RMSE is 4.68 °C, and the mean absolute error is 3.62 °C. The mean RMSE for the hottest decile of observed LST days is 6.66 °C[2]

3.8.2 Income distribution emulator

Sub-city income observations are not available consistently across the full 40-city application. The current configurations use observed income in 23 cities and emulator-derived income ordering in the remaining 17. The emulator is included in URBADAPT-HEAT to provide the within-city socioeconomic ordering required by the AC-downscaling and equity modules when measured income is unavailable.

The operational target is not absolute income, which is difficult to compare across countries, reference years, currencies, and statistical definitions. Instead, the current configuration models the precomputed within-city income percentile, $p_{\mathrm{inc}}^{\mathrm{obs}} \in [0, 1]$, directly. Lower values identify poorer zones and higher values identify richer zones. The code also supports an alternative log-transformed relative-income target, but that alternative is not used in the deployed configuration described here.

The model uses 16 zone-level predictors: population density, built volume per capita, mean building height, the shares of compact, open, industrial and vegetated LCZ classes, the population shares aged under 15 and 65 or older, employment rate, foreign-born population share, residential churn, health- and education-facility densities, distance to the city centre, and zone area. Each of these predictors is transformed to a within-city z-score. Raw city population is retained as an additional continuous context variable, and country is represented by one-hot indicators. Missing numerical predictor values are median-imputed.

The deployed estimator is a LightGBM gradient-boosted regression model with 800 trees, a learning rate of 0.03, 31 leaves, a minimum of 20 observations per leaf, and row- and feature-subsampling fractions of 0.8. Training weights give every country the same total weight while preserving population weighting among zones within each country. This prevents countries with many sub-city units from dominating the pooled loss.

The deployed model is fitted to 5,135 labelled sub-city zones in 145 cities and 14 countries. Great Britain is excluded from this fit because its zones lack the Eurostat 2021 census-grid covariates required by the configured completeness filter. Validation follows the intended transfer setting and holds out one entire country at a time. Within each held-out city, performance is assessed using the population-weighted Spearman correlation between the observed and predicted zone ordering. The zone-count-weighted mean of these city-specific correlations is 0.592. When the same out-of-fold predictions are pooled within each held-out country before scoring, the corresponding zone-count-weighted mean correlation is 0.597. For the poorest and richest observed fifths of zones, the respective out-of-fold identification rates are 0.537 and 0.453, compared with 0.20 under random assignment. Performance

[2]Among the four illustrative cities demonstrated in this paper (Sect. 4), Rome, Lisbon, and Copenhagen are represented in the training-city pool and have corresponding city-holdout evaluations. Athens is absent from the 130-city fitted pool; application of the pooled bundle to Athens is therefore a fully out-of-sample spatial transfer and is treated accordingly in the limitations.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

nevertheless varies substantially among countries, and emulator-derived ranks should therefore be interpreted as proxy socioeconomic distributions rather than substitutes for locally observed income.

For deployment, the fitted values provide a zone-level ordering score. Notebook 05 converts that ordering to the same midpoint, population-weighted within-city percentile used for observed income:

$$p_{\mathrm{inc},i} = \frac{\sum_{j<i} P_j + \frac{1}{2} P_i}{\sum_j P_j}, \tag{26}$$

where zones are ordered from poorer to richer and P_i denotes the modelled population of zone i. The percentile is subsequently used by the logistic AC-downscaling function and by income-targeted policy allocation. The income-source switch is configuration-driven (`observed` or `emulator`); emulator predictions are joined using the same normalised zone key as the policy geography, and the workflow stops if the default 95 % minimum matching coverage is not achieved[3].

3.9 Adaptation modules

The adaptation modules represent three distinct policy pathways - air conditioning, urban vegetation, and early warning systems - and a set of combined-policy scenarios. The pathway-specific reference states preserve the relevant existing physical infrastructure. AC benefits are measured relative to mortality under observed current AC penetration, while the vegetation intervention starts from the observed urban greenness and estimates the additional benefit of the configured planting policy. The EWS module also uses the residual mortality burden under current AC as its mortality baseline. Existing EWS maturity is represented through the configured marginal, intermediate, or counterfactual interpretation, which determines the efficacy and cost attributed to the modelled intervention; protection already delivered by an operational EWS is not separately subtracted from the baseline daily mortality series. Benefits are therefore measured relative to pathway-appropriate current-state references rather than to a single common mortality baseline that explicitly embeds the protective effects of all three adaptations. Table 4 lists the input data and main parameters of the adaptation component.

[3]All four cities analysed in this paper's illustrative application (Sect. 4) - Rome, Athens, Lisbon, and Copenhagen - are currently configured to use observed income. The income emulator therefore supports the scalability of the wider 40-city framework but does not affect the four-city results presented in this paper.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Input / Parameter	Source / Description	Notes
Air conditioning		
City-level AC penetration	NUTS3 regional statistics from Falchetta et al. (2024)	Baseline by target year
Income downscaling	Observed sub-city income where available; emulator-derived within-city percentiles otherwise	Logistic income-rank curve
AC efficacy (ε_a)	Sera et al. (2020), age-differentiated	0.20 / 0.30 / 0.40 central scenario
Electricity demand	2.5-15 kWh per AC household per cooling day, from city-mean T2M	Spread across zones by temperature
Waste heat	AC heat flux over the built-up area; 0.02 K per W m^{-2} (Li et al., 2026)	Daily penalty; COP loss (Kou et al., 2026)
Urban vegetation		
Green View Index	Falchetta and Hammad (2025), 356-city dataset	Baseline street-level greenness
LCZ raster	`lcz_filter_v3.tif`	Scope restricted to built classes
Cooling coefficients	GVI-LST regression, city/month/LCZ	Cross-European calibration
LST-to-T2M emulator	UrbClim and Falchetta and Hammad (2025) pooled bundle	Two-step cooling bridge
Tree costs	€10 million per summed regional ΔGVI point Falchetta et al. (2026a); city-specific per-tree references for the O&M conversion	25-year linear rollout; 12-year maturity; 5-year starting age
Early warning systems		
System maturity	Three-class taxonomy: epidemiological HHWS, meteorological HHWS, or traffic-light warnings only	Determines marginal, intermediate, or counterfactual interpretation
Warning trigger	Harmonised citywide daily heat-mortality threshold	Calibrated under current AC in the 2020 reference year
Warning season	City-specific seasonal window	05-01 to 09-30 (Rome, Athens); 05-15 to 09-30 (Lisbon, Copenhagen)
Efficacy	Age-specific marginal, midpoint intermediate, or flat counterfactual efficacy	Low, central, and high scenarios
Threshold projection	Reference-year threshold population-scaled at each modelled anchor year	Track-B event mask retained as a non-positive-threshold fallback
Setup CAPEX	€0 / €75,000 / €150,000 for the three maturity classes	Plus city-specific fixed and warning-day OPEX
Cost model	Pavanello and Valenti (2025) per-capita activation cost; Chiabai et al. (2018) alternative	25-year discounted accounting

Table 4. Adaptation component: input data and main parameters

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.9.1 Air conditioning

Air conditioning is represented as a private adaptation option that reduces heat-related mortality among the covered population. It attenuates the heat-mortality dose for the covered population rather than modifying the outdoor temperature hazard or population exposure. For age group a, year y, and grid cell i, the mortality dose under AC coverage $c_i(y)$ is obtained by multiplying the unprotected dose by

$$1 - \varepsilon_a c_i(y), \tag{27}$$

as defined in Equation 22. In the central, moderate-efficacy configuration, ε_a equals 0.20, 0.30, and 0.40 for the <15, 15-64, and 65+ age groups, respectively. The baseline is therefore the current-AC state, with non-zero existing penetration, rather than a counterfactual without AC.

Regional projections of AC penetration are read at the 2020, 2030, 2040, and 2050 anchor years from Falchetta et al. (2024). The relevant regional units are specified separately in each city configuration. Where more than one regional identifier is configured, the corresponding values are averaged within each scenario and year. These regional penetration targets are downscaled to the policy-zone geography using the within-city income percentile $p_{\mathrm{inc},i}$ described above. All four cities examined in this paper use observed sub-city income data.

The initial income-gradient shape is

$$q_i = c_{\min} + \frac{c_{\max} - c_{\min}}{1 + \exp\left[-s\left(p_{\mathrm{inc},i} - 0.5\right)\right]}, \tag{28}$$

where the current default parameters are $c_{\min} = 0.05$, $c_{\max} = 0.98$, and $s = 2.2$. For each anchor year, this shape is calibrated as

$$c_i(y) = \mathrm{clip}\left[c_{\min} + k_y\left(q_i - c_{\min}\right), c_{\min}, c_{\max}\right], \tag{29}$$

where the scalar k_y is found numerically so that the population-weighted mean of $c_i(y)$ reproduces the regional penetration target for that year. The resulting zone values are rasterised to a common model grid and used to attenuate the age-specific mortality doses. Populated grid cells of the modelled area that lie outside the policy zones are assigned the zones' population-weighted mean coverage, which equals the regional target, and keep it under both allocation policies described below, so that policy-induced users are added only within the zones.

Two stylised allocation policies are compared under an equal-users constraint:

1. *Income-targeted catch-up*: zones are ordered from poorer to richer using $p_{\mathrm{inc},i}$. Whole zones are selected until their cumulative 2020 population is as close as possible to 50 % of the modelled city population. In the selected zones, coverage is raised to at least 0.70, without reducing zones already above that level and subject to the overall coverage bounds of 0.05 and 0.98.

2. *Uniform expansion*: the same percentage-point increment is applied to all valid zones, subject to the 0.98 upper bound and without reducing existing coverage. For every anchor year, the increment is calibrated so that, using

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

the fixed 2020 zone populations, the number of additional users equals that generated by the income-targeted policy.

Thus, the targeted share can differ slightly from exactly 50 % because policies are assigned to whole administrative zones. The comparison nevertheless holds the additional number of users constant and isolates the consequences of their spatial and socioeconomic allocation.

Annual electricity use per AC user, μ_y, follows the daily heat of each anchor year, through one rule applied identically in every city. On a day t with city-mean daily 2 m air temperature T_t (°C), cooling activity is

$$a_t = \mathrm{clip}\left(\frac{T_t - 18}{25 - 18}, 0, 1\right), \tag{30}$$

and an AC household uses

$$e_t = \begin{cases} e_{\mathrm{on}} + (e_{\mathrm{full}} - e_{\mathrm{on}})\, a_t, & a_t > 0, \\ 0, & a_t = 0, \end{cases} \tag{31}$$

kWh of electricity, with $e_{\mathrm{on}} = 2.5$ kWh on the first cooling day and $e_{\mathrm{full}} = 15$ kWh on a day at full activity. The range spans a single unit running for a few hours at the onset to several hours of full-load operation of a split unit of about 1 kW electrical input on a very hot day. Annual electricity per AC user is $\mu_y = \sum_{t \in y} e_t / h$, with $h = 2.3$ persons per household, the EU-27 mean. In 2020 this gives about 1,470, 1,900, 970, and 90 kWh per AC household in Rome, Athens, Lisbon, and Copenhagen. In the cooler cities such totals are close to national statistics, whereas in the hottest they exceed them, so there the rule is an upper envelope for both the electricity cost and the waste heat it drives. The rule replaces regional (NUTS) projections of electricity use per AC user, which did not follow the cities' climates.

Within the city, per-user consumption is dispersed across zones using annual-mean daily T2M. Let $H_{i,y}$ denote the population-weighted mean temperature in zone i and H_y its population-weighted city mean. The unscaled zone-level electricity use is

$$e^{\mathrm{raw}}_{i,y} = \mu_y\, r^{\gamma}_{i,y}, \qquad r_{i,y} = \mathrm{clip}\left(\frac{H_{i,y}}{H_y}, 0.25, 4\right), \tag{32}$$

with unit elasticity, $\gamma = 1$. A year-specific multiplicative factor then rescales $e^{\mathrm{raw}}_{i,y}$ so that its mean, weighted by the population currently using AC, exactly reproduces μ_y. Annual electricity use under coverage scenario s is subsequently calculated as

$$E^s_y = \sum_i P_i(y) c^s_i(y) e_{i,y}. \tag{33}$$

The city mean is therefore set by the daily heat of the anchor year, and hotter zones are assigned proportionally higher consumption per user.

The model also accounts for the outdoor warming externality associated with AC waste heat. AC units reject their electricity, times $1 + \mathrm{COP}$, outdoors as sensible heat, with $\mathrm{COP} = 3.5$ for the installed split-unit stock. The AC users

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

of the current-AC and policy-AC branches, summed over all populated cells of the modelled area, are interpolated from the anchor years across the 25-year assessment horizon. On each day the heat rejected by these users, at e_t per AC household (Equation (31)), gives a heat flux averaged over the city's built-up area, the Local Climate Zones 1-10 of the functional urban area (about 1,490, 1,300, 880, and 800 km^2 in Rome, Athens, Lisbon, and Copenhagen). The flux raises city-wide daily-mean air temperature by $S = 0.02\,\mathrm{K}$ per $\mathrm{W\,m^{-2}}$, increased slightly by a first-order correction for the efficiency loss of units running in warmer air (Kou et al., 2026). The sensitivity S is taken at the upper end of the city-scale evidence: modelling studies give a median of 0.017 K per $\mathrm{W\,m^{-2}}$ of anthropogenic heat flux (interquartile range 0.008-0.029; Li et al., 2026), observations in Tokyo and Osaka 0.007-0.009 (Kikegawa et al., 2014), simulations of Paris about 0.005 per $\mathrm{W\,m^{-2}}$ of peak flux (Tréméac et al., 2012; de Munck et al., 2013), Berlin about 0.01 (Jin et al., 2020), ten central-European cities 0.006-0.007 (Karlický et al., 2026), and the continental United States about 0.01 (Wang et al., 2023). The upper value is chosen because European cities are compact; the uncertainty analysis samples 0.01-0.035, the upper bound corresponding to Phoenix under extreme heat (Salamanca et al., 2014). This flux formulation replaces the earlier lookup from AC penetration to nighttime warming based on Salamanca et al. (2014), whose penetration levels are the air-conditioned share of floor area in Phoenix at extreme-heat loads (about $8\,\mathrm{W\,m^{-2}}$ of daily-mean AC heat at 65 %); read as European household penetration, the lookup implied roughly ten times the warming per $\mathrm{W\,m^{-2}}$ that city-scale studies support.

Waste-heat mortality follows the days on which AC operates and how much it runs. Each day's warming is applied to that day's marginal heat deaths per degree, obtained from a uniform +1 řC perturbation of the corresponding branch, so the waste-heat deaths of a branch equal its warming at full activity times the sum of the daily marginal deaths weighted by e_t/e_{full}. The daily weighting is computed at each anchor year and interpolated across the horizon. It places the penalty on the hot days on which AC runs and heat deaths concentrate, rather than spreading the warming evenly over the year. Waste-heat deaths are evaluated separately under the current-AC and policy-AC branches, and the penalty assigned to policy expansion is their difference. The net mortality benefit of AC is consequently the direct protective benefit minus this incremental waste-heat penalty. The governing equations are given in Sect. 3.9.4 (Equations (50)–(52)). The only study that carries AC waste heat through to mortality finds it taking back about 7-8 % of the protective benefit of AC at 96-98 % household AC ownership in Japanese cities (Chua et al., 2023).

AC costs are specified per modelled additional user. The current configurations use CAPEX values of €1,498 in Rome, €1,349 in Athens, €1,336 in Lisbon, and €2,156 in Copenhagen. Annual maintenance equals 4 % of CAPEX, and equipment lifetime is 12 years in all four cities. Electricity tariffs are €0.17 $\mathrm{kWh^{-1}}$ in Rome, €0.14 $\mathrm{kWh^{-1}}$ in Athens, €0.17 $\mathrm{kWh^{-1}}$ in Lisbon, and €0.30 $\mathrm{kWh^{-1}}$ in Copenhagen.

The main CBA evaluates the income-targeted policy relative to the evolving current-AC trajectory over 2020-2044, using a 25-year horizon and a 3 % annual discount rate. Anchor-year coverage and electricity inputs are interpolated annually. CAPEX is assigned to non-negative cohorts of additional users, with replacement expenditure every 12 years when the replacement falls within the assessment horizon. Maintenance is charged on the incremental user

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

stock, and electricity expenditure is based on incremental consumption. Consequently, the CBA includes only the costs and benefits attributable to policy-induced AC expansion beyond the current-AC reference. Fig. 6 maps the resulting zone-level penetration for Rome in 2050.

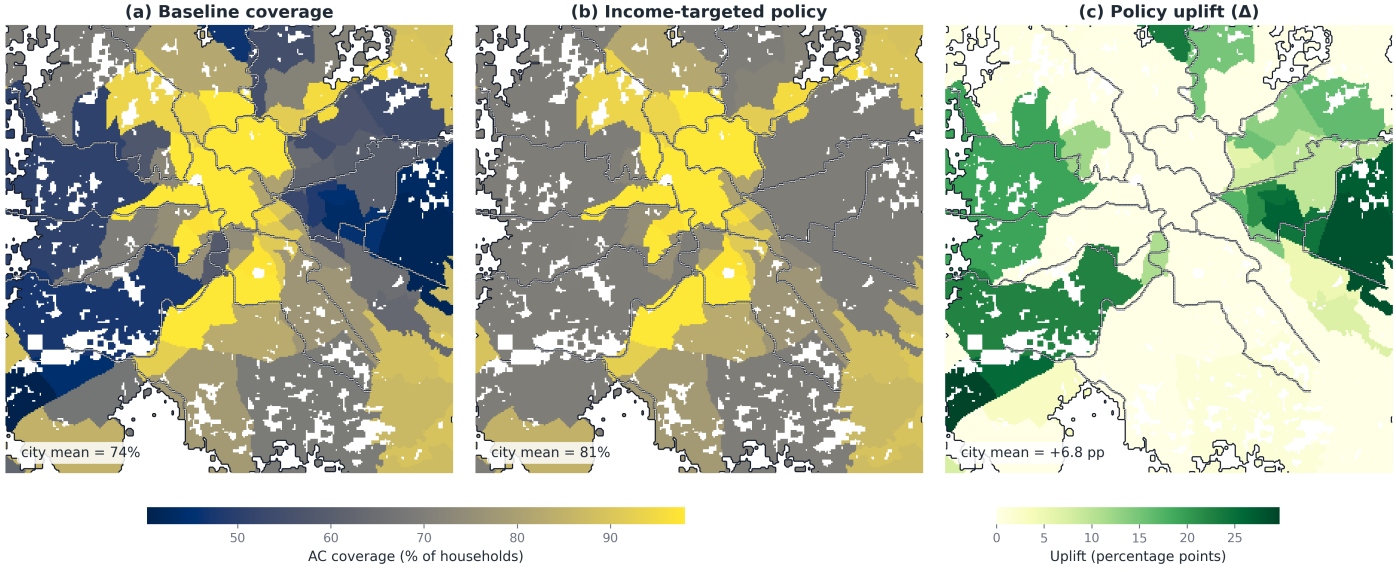


Figure 6. Zone-level (CAP postal code) air conditioning penetration for Rome in 2050, mapped on the 100 m reference grid with *municipi* boundaries overlaid for reference (boundaries © OpenStreetMap contributors, ODbL 1.0). (a) Baseline AC coverage under the exogenous adoption trajectory (logistic income-rank downscaling of the NUTS3 target; see Fig. A2 in the Supplement). (b) AC coverage after the income-targeted catch-up policy, which prioritises lower-income zones. (c) Policy uplift (Δ): the largest gains concentrate in lower-income peripheral zones where baseline penetration is lowest. Panels (a) and (b) share a common coverage scale (percent of households); panel (c) is in percentage points. Panel annotations report the population-weighted city-mean value.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.9.2 Street trees

Urban vegetation is represented as a public adaptation pathway that modifies the outdoor thermal environment. In contrast to air conditioning, which reduces indoor exposure for the households that adopt it, the street-tree intervention acts on the temperature hazard experienced by the population within the treated areas. The intervention is operationalised as an increase in the Green View Index (GVI), rather than through an explicit simulation of individual trees.

Baseline GVI is obtained from the georeferenced street-level observations Falchetta and Hammad (2025). Values expressed on a 0-1 scale are converted to a 0-100 scale, and observations falling within each cell of the reference grid are averaged. Grid cells without an observation are assigned the value of the nearest observed cell within the city mask. Local Climate Zone (LCZ) data are independently reprojected to the reference grid, and the intervention is restricted to built LCZ classes 1-10.

The policy is defined over the administrative geography specified for each city. For the four illustrative cities presented below: Municipi in Rome, Communities in Athens, Freguesias in Lisbon, and Bydeler in Copenhagen. These sub-municipal boundaries are retrieved programmatically from OpenStreetMap via the Overpass API, selecting the administrative level that resolves the sub-municipal units of each city (`admin_level` 10 for Rome). OpenStreetMap data are © OpenStreetMap contributors and are made available under the Open Database License (ODbL) v1.0; the boundary geometries reproduced in this paper and the maps derived from them are produced works under that licence. For each policy region r, baseline greenness G_r^0 is calculated as the unweighted spatial mean of grid-cell GVI over the built-LCZ pixels belonging to that region. The policy target is the 75th percentile of the regional baseline-GVI distribution, rounded upward to the nearest 0.1 GVI point:

$$G^* = \frac{\lceil 10\, Q_{0.75}\big(\{G_r^0\}\big) \rceil}{10}. \tag{34}$$

The prescribed uplift in each region is therefore

$$\Delta G_r = \max\big(0, G^* - G_r^0\big). \tag{35}$$

Regions already at or above the target receive no intervention. Within each treated region, the same regional uplift is assigned to all eligible built-LCZ grid cells. Thus, the central policy is a regional catch-up scenario: it raises below-target regions to a common benchmark but does not optimise the spatial allocation according to mortality benefits or socioeconomic vulnerability.

The cooling effect of the intervention is estimated using city-, calendar-month-, and LCZ-specific GVI-land-surface-temperature coefficients. For a grid cell i located in region $r(i)$, the land-surface-temperature perturbation is

$$\Delta \mathrm{LST}_{i,m} = \min\big(\theta_{c,m,\ell(i)}, 0\big)\, \Delta G_{r(i)}, \tag{36}$$

where $\theta_{c,m,\ell(i)}$ is the coefficient for city c, month m, and LCZ class $\ell(i)$. The coefficient input table retains only estimates meeting the statistical-significance criterion of the underlying city-month-LCZ regressions. In the operational

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

cooling-only implementation, retained positive coefficients are set to zero so that the intervention cannot generate warming; month-LCZ combinations for which no statistically significant coefficient is available are likewise assigned zero cooling. Consequently, a zero operational coefficient can represent either the absence of a retained negative coefficient or a retained positive coefficient that has been clipped under the cooling-only assumption.

These coefficients must be interpreted at the spatial scale and under the physical assumptions of the source simulations. The UrbClim runs used to estimate them do not prescribe supplementary irrigation of urban vegetation. Under hot and dry conditions, particularly during summer in water-limited Mediterranean settings, limited modelled soil-moisture availability can suppress evapotranspiration and weaken, eliminate, or potentially reverse the grid-scale association between GVI and land-surface temperature. Moreover, the nominal 100 m grid does not explicitly resolve the pedestrian-scale shading experienced directly beneath individual tree canopies. Weak or zero operational cooling coefficients in water-limited settings should therefore be interpreted as an absence of statistically supported grid-scale LST cooling under the simulated non-irrigated conditions and at the available spatial resolution, rather than as evidence that irrigated and well-maintained street trees, or their local shade, provide no thermal benefit.

The resulting monthly ΔLST fields are translated into daily near-surface-air-temperature changes using the pooled quadratic LST-T2M emulator described above. This translation is evaluated using the baseline GVI, LCZ, and daily baseline temperature of each grid cell, and is subject to the emulator guardrails described in Sect. 3.8.

Mortality is subsequently recalculated under four internally consistent combinations: baseline air-conditioning coverage without trees, baseline air-conditioning coverage with trees, expanded air-conditioning coverage without trees, expanded air-conditioning coverage with trees. The standalone tree benefit is defined by comparing the first two combinations, whereas the conditional tree benefit in the presence of the air-conditioning policy is obtained from the latter two. All other demographic, socioeconomic, and impact-function inputs are held constant in these comparisons.

An equity-oriented allocation is retained as a sensitivity analysis. It redistributes the same total modelled greening budget used by the central policy among regions whose baseline GVI is below G^*. Allocation weights are proportional to the population-weighted mean Social Vulnerability Index (SVI) of the eligible regions:

$$\Delta G_r^{\mathrm{eq}} = \mathbb{I}\big(G_r^0 < G^*\big)\, \frac{\overline{\mathrm{SVI}}_r}{\sum_{s:G_s^0<G^*} \overline{\mathrm{SVI}}_s} \sum_s \Delta G_s, \tag{37}$$

where $\overline{\mathrm{SVI}}_r$ is the population-weighted regional SVI. This sensitivity therefore prioritises more socially vulnerable eligible regions. It does not impose a region-specific upper bound at G^*, and consequently should be interpreted as a reallocation experiment rather than an alternative catch-up rule.

For costing, the scale of the intervention is summarised by

$$\Delta I = \sum_r \Delta G_r, \tag{38}$$

that is, the unweighted sum of regional GVI uplifts. Capital expenditure is configured at €10 million per unit of ΔI. The model does not derive an explicit number of planted trees or use an assumed planting density. Instead,

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

city-specific per-tree reference costs are used to convert the capital-cost scale into annual operation-and-maintenance expenditure:

$$C_{\mathrm{O\&M},\,\Delta I} = \frac{C_{\mathrm{O\&M,\,tree}}}{C_{\mathrm{CAPEX,\,tree}}} C_{\mathrm{CAPEX},\,\Delta I}. \tag{39}$$

The configured capital and annual maintenance reference values per tree are respectively €680 and €63 for Rome, €612 and €57 for Athens, €606 and €56 for Lisbon, and €978 and €91 for Copenhagen. These values imply city-specific annual maintenance costs per unit of ΔI, while preserving the common €10-million capital-cost scale.

Capital expenditure is distributed uniformly over a 25-year implementation period. Tree cohorts follow an underlying 12-year linear maturity curve, with the central specification representing planting stock with a starting age of five years; a zero-starting-age specification is retained as a sensitivity. Benefits and maintenance expenditures are scaled by the same cohort rollout and maturity profile, and no benefit or maintenance expenditure is assigned in a cohort's planting year. Trees have a configured lifetime of 25 years. All monetary flows are evaluated over the 2020-2044 horizon and discounted at 3 % annually. Fig. 7 compares the central regional-GVI catch-up allocation with the SVI-weighted equity sensitivity.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.



Figure 7. Comparison of the central regional-GVI catch-up policy and the SVI-weighted allocation sensitivity. The upper-left panel reports region-level GVI uplifts under the two allocation rules, while the upper-right panel shows the redistribution produced by the equity sensitivity relative to the central policy. The lower-left panel relates the allocated uplifts to baseline regional GVI, and the lower-right panel compares the corresponding avoided heat-attributable deaths at the modelled anchor years. The two policies preserve the same total summed regional GVI uplift, but differ in its distribution across eligible regions.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.9.3 Early warning systems

Early warning systems are represented as a public, city-level mortality-reduction policy layered on top of the current-AC baseline. They do not modify the temperature hazard or the population exposure grids. Instead, they reduce modelled heat-attributable mortality on days for which a warning is activated.

EWS maturity taxonomy. European heat-warning arrangements differ in their epidemiological basis, institutional integration, and health-system response (Casanueva et al., 2019). URBADAPT-HEAT therefore assigns each city to one of the three baseline-maturity classes set out in Table 5:

Baseline class	Policy represented	Efficacy interpretation	Setup CAPEX
Epidemiologically based mature HHWS	Marginal enhancement of an existing heat-health warning and action system with an organised health-sector response and epidemiological surveillance	Age-specific marginal	€0
Meteorologically or historically based HHWS	Upgrade of an operational heat-warning or heat-action system that lacks real-time epidemiological surveillance	Age-specific intermediate	€75,000
Traffic-light warnings only	Deployment of a comprehensive HHWS or heat-health action plan from a low institutional baseline	Flat counterfactual	€150,000

Table 5. Three-class EWS maturity taxonomy and its implementation in URBADAPT-HEAT. HHWS denotes a heat-health warning system.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

The first class includes countries with mature epidemiologically based systems, represented in the configurations by archetypes such as `mature_hhwws` and `mature_hybrid_icaro`. The second class is represented by `meteo_hhwws`; twelve city configurations currently use this intermediate class. The third class contains cities with meteorological or traffic-light alerts but without a comprehensive HHWS, generally represented by `weak_threshold`. Copenhagen uses the specialised `cold_city_extreme_focus` archetype but belongs to this third policy class.

The taxonomy characterises baseline institutional maturity and determines efficacy interpretation and setup expenditure. Table B2 gives the full class definitions, the corresponding modelling assumptions and limitations, and the country assignment with its official sources. It does not reproduce the heterogeneous operational triggers used by actual national or subnational systems. For comparability, all central model runs use the same harmonised mortality-based warning trigger described below. Local surveillance protocols, communication quality, public compliance, subnational differentiation, and health-system response capacity are therefore not simulated explicitly.

In the four-city application, Rome and Lisbon belong to the first class and are evaluated as marginal enhancements of mature systems. Athens and Copenhagen belong to the third class and are evaluated as counterfactual deployments of comprehensive HHWSs. None of the four demonstration cities belongs to the intermediate class, although that class is implemented in the wider city configuration.

Warning activation. The harmonised warning trigger is calibrated epidemiologically rather than defined by a fixed temperature threshold. Notebook 05 first reconstructs daily citywide heat-attributable deaths under current AC coverage using the age-specific impact functions and AC efficacies described above. For city c, a reference threshold $\theta_{c,2020}$ is selected from the modelled daily mortality values so that the number of warning-season days at or above the threshold is as close as possible to the configured activation-day target:

$$\theta_{c,2020} = \arg\min_{\theta} \left| \sum_{d \in \mathcal{S}_c} \mathbf{1}\left(D^{\mathrm{AC}}_{c,d,2020} \geq \theta\right) - N^{*}_{c} \right|, \tag{40}$$

where $D^{\mathrm{AC}}_{c,d,2020}$ denotes modelled daily heat-attributable deaths under current AC, $\mathcal{S}_c$ is the city-specific warning season, and N^{*}_{c} is the target activation count. The target is 23 days in all four pilot cities, as in every other city configuration. In the four pilot configurations, candidate warning days are restricted to a predefined warm-season window: 1 May–30 September for Rome and Athens and 15 May–30 September for Lisbon and Copenhagen. The Lisbon window follows the documented operating period of the Portuguese ICARO surveillance system, while the remaining windows are configuration assumptions delimiting seasonal threshold calibration and activation.

In v1.0, thresholds updating is implemented as population scaling rather than repeated epidemiological recalibration. At each modelled anchor year y, the reference threshold is multiplied by the ratio of projected city population to reference-year population:

$$\theta_{c,y} = \theta_{c,2020} \frac{P_{c,y}}{P_{c,2020}}. \tag{41}$$

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

A warning is activated when a day lies within the configured season and modelled heat-attributable deaths are at or above the applicable threshold:

$$W_{c,d,y} = \mathbf{1}(d \in \mathcal{S}_c)\,\mathbf{1}\big(D^{\mathrm{AC}}_{c,d,y} \geq \theta_{c,y}\big)\,. \tag{42}$$

The workflow additionally supports a non-positive-threshold contingency. When enabled, warning days are obtained by intersecting the configured warning season with the Track-B extreme-event mask. Among the four pilot configurations, this contingency is enabled only for Copenhagen and is invoked only if its epidemiologically calibrated threshold is non-positive; in the reported runs that threshold is positive and the contingency is not invoked.

Three efficacy interpretations. For age group a, city c, and low, central, or high scenario s, the effective EWS efficacy is selected according to the city configuration:

$$\varepsilon_{a,c,s} = \begin{cases} \varepsilon^{\mathrm{marg}}_{a,c,s}, & I_c = \text{marginal}, \\ \dfrac{\varepsilon^{\mathrm{marg}}_{a,c,s} + \varepsilon^{\mathrm{cf}}_{c,s}}{2}, & I_c = \text{intermediate}, \\ \varepsilon^{\mathrm{cf}}_{c,s}, & I_c = \text{counterfactual}, \end{cases} \tag{43}$$

where $\varepsilon^{\mathrm{marg}}_{a,c,s}$ is the age-specific marginal efficacy and $\varepsilon^{\mathrm{cf}}_{c,s}$ is the flat counterfactual efficacy. The intermediate interpretation is therefore age-differentiated and is calculated dynamically as the arithmetic midpoint between the two corresponding configured efficacy values. This midpoint is a modelling assumption used to represent the intermediate scope for improvement of meteorologically based HHWSs.

For the four pilot cities, the central marginal efficacies for the < 15, 15-64, and 65+ age groups are $(0.05, 0.07, 0.15)$ in Rome and $(0.03, 0.05, 0.10)$ in Lisbon. The central flat counterfactual efficacies are 0.20 in Athens and 0.15 in Copenhagen. Although efficacy parameters are retained for the < 15 group, that group contributes no EWS benefit in the current Masselot implementation because its heat-mortality impact function is zero.

For each age group, gross avoided deaths are calculated as

$$A^{\mathrm{gross}}_{a,c,y,s} = D^{W}_{a,c,y}\,\varepsilon_{a,c,s}\,r_y\,\big(1 - \overline{C}_{c,y}\,\omega_{c,s}\big)\,, \tag{44}$$

where $D^{W}_{a,c,y}$ is baseline heat-attributable mortality on warning days, r_y is the implementation factor, $\overline{C}_{c,y}$ is mean gridded current-AC coverage, and $\omega_{c,s}$ is the configured AC-EWS overlap factor. The overlap adjustment represents the smaller residual mortality pool addressable by an EWS as AC coverage increases. Under the central scenario, $\omega_{c,s}$ equals 0.50 in Rome, 0.30 in Athens, and 0.20 in Lisbon and Copenhagen.

All three EWS classes currently use the same implementation profile. Efficacy is set to 10 % of its full value during the first three policy years and to its full value thereafter:

$$r_y = \begin{cases} 0.10, & y = 2020, 2021, 2022, \\ 1, & y \geq 2023. \end{cases} \tag{45}$$

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Thus, the current implementation uses a three-year low-efficacy period followed by a step to full efficacy. It does not yet apply class-specific ramp durations, even though institutional implementation times could plausibly differ among the three maturity classes.

Gross avoided deaths are adjusted for short-term mortality displacement. If $h_{a,c,s}$ denotes the configured age- and scenario-specific displacement fraction, net avoided deaths and life-years saved are

$$A^{\text{net}}_{a,c,y,s} = A^{\text{gross}}_{a,c,y,s}\,(1 - h_{a,c,s}), \tag{46}$$

$$L_{a,c,y,s} = A^{\text{net}}_{a,c,y,s}\ell_{a,c}, \tag{47}$$

where $\ell_{a,c}$ is the configured residual life expectancy for age group a. Gross avoided deaths, displacement-adjusted net avoided deaths, and life-years saved are retained as distinct outputs.

EWS costs. The default cost calculation uses the per-capita, per-activation-day formulation of Pavanello and Valenti (2025). Variable annual operating expenditure is based on US$0.014 per resident per warning day, converted using the configured factor of 0.92 euros per US dollar. A city-specific annual standing operating cost is added independently of the number of warning days. Annual undiscounted expenditure is

$$C^{\text{EWS}}_{c,y} = C^{\text{setup}}_{c}\mathbf{1}(y = 2020) + r_y\left[C^{\text{fixed}}_{c} + 0.014 \times 0.92 \times P_{c,y}N^{W}_{c,y}\right], \tag{48}$$

where $N^{W}_{c,y}$ is the number of warning days. Setup CAPEX follows the maturity taxonomy:

$$C^{\text{setup}}_{c} = \begin{cases} e0, & I_c = \text{marginal}, \\ e75{,}000, & I_c = \text{intermediate}, \\ e150{,}000, & I_c = \text{counterfactual}. \end{cases} \tag{49}$$

The configured fixed annual operating costs for the four pilot cities are €212,009 for Rome, €190,830 for Athens, €189,083 for Lisbon, and €305,022 for Copenhagen. Accordingly, setup expenditure is zero for Rome and Lisbon and €150,000 for Athens and Copenhagen. The same implementation factor r_y is applied to fixed and warning-day operating expenditure, whereas setup expenditure is incurred in full in the first year.

The alternative Chiabai et al. (2018) specification replaces the Pavanello variable-cost term with a city-specific incremental cost per warning day. The configured values for the four pilot cities are €6,572 for Rome, €5,916 for Athens, €5,862 for Lisbon, and €9,456 for Copenhagen. In the 25-year calculation, these per-day values are population-scaled relative to 2020. Fixed annual operating expenditure and class-specific setup CAPEX remain part of the accounting under this alternative.

Warning-day counts, age-specific warning-day deaths, population, and mean AC coverage are linearly interpolated between the 2020, 2030, 2040, and 2050 model outputs to construct annual trajectories for the 2020-2044 CBA horizon. Costs are discounted at 3 % annually. EWS cost-effectiveness is reported using net avoided deaths and life-years saved, thereby retaining the mortality-displacement adjustment in both denominators. Fig. 8 reports the resulting early-warning-system diagnostics.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

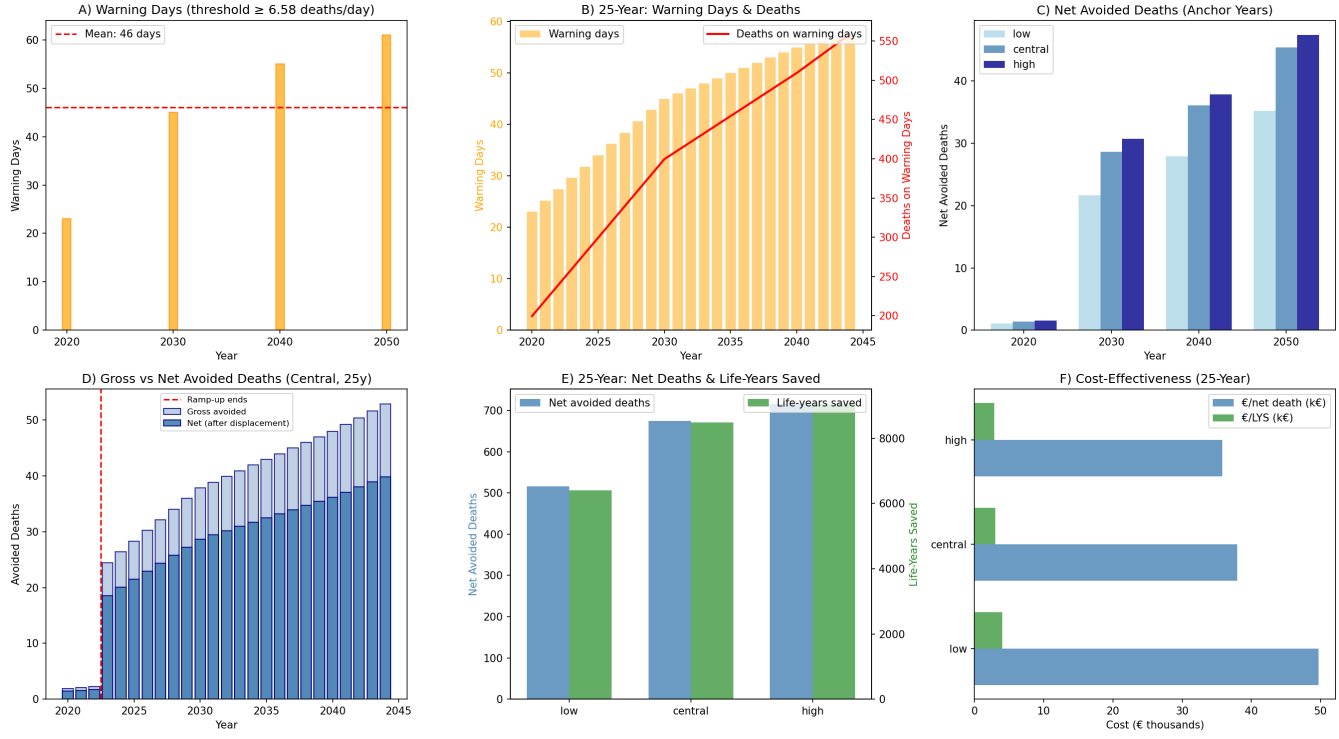


Figure 8. Early-warning-system diagnostics. Panel (A) reports warning-day counts at the modelled anchor years, with the mean count shown by the dashed line. Panel (B) shows the interpolated 25-year trajectories of warning days and baseline heat-attributable deaths occurring on warning days. Panel (C) compares net avoided deaths at the anchor years under the low, central, and high efficacy scenarios. Panel (D) reports annual gross and displacement-adjusted net avoided deaths under the central scenario; the dashed vertical line marks the end of the three-year low-efficacy period. Panel (E) compares cumulative 25-year net avoided deaths and life-years saved across efficacy scenarios. Panel (F) reports the corresponding cost per net death avoided and cost per life-year saved. The same diagnostic structure is generated separately for each pilot city.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.9.4 Feedbacks among adaptation options

URBADAPT-HEAT represents three groups of interactions among the adaptation pathways: the outdoor waste-heat externality generated by AC, overlap between current AC and EWS protection, and two vegetation-AC interactions affecting electricity demand and the AC waste-heat penalty. These interactions are implemented as first-order adjustments: the model does not solve an iterative equilibrium among outdoor temperature, cooling demand, AC efficiency, and adaptation uptake.

AC waste-heat externality. The waste-heat calculation distinguishes the evolving current-AC trajectory from the income-targeted policy-AC trajectory. For each branch $b \in \{\mathrm{cur}, \mathrm{pol}\}$, the number of AC users $U_{b,y}$ is summed over the populated cells of the modelled area from the gridded population and AC-coverage fields at the anchor years and linearly interpolated to annual values. AC units reject the electricity they use, multiplied by $1 + \mathrm{COP}$, outdoors as sensible heat, with $\mathrm{COP} = 3.5$ representing the installed split-unit stock. On a day at full cooling activity, the resulting heat flux averaged over the city's built-up area A, the area of Local Climate Zones 1–10 within the functional urban area, is

$$Q_{b,y}^{\mathrm{full}} = \frac{U_{b,y}}{h} \frac{10^3 \, e_{\mathrm{full}} \, (1 + \mathrm{COP})}{24 \, A} \quad \left[\mathrm{W\,m^{-2}}\right], \tag{50}$$

where $h = 2.3$ persons per household, $e_{\mathrm{full}} = 15\,\mathrm{kWh}$ is the household AC electricity on a day at full cooling activity (Equation (31)), and A is in square metres. The flux raises the city-wide daily-mean air temperature by $S\,Q_{b,y}^{\mathrm{full}}$, with $S = 0.02\,\mathrm{K}\,(\mathrm{W\,m^{-2}})^{-1}$. A first-order correction for the efficiency loss of units operating in warmer outdoor air amplifies this warming:

$$\Delta T_{b,y}^{\mathrm{full}} = S\,Q_{b,y}^{\mathrm{full}} \left[1 + \kappa\, S\, Q_{b,y}^{\mathrm{full}} \frac{1 + 1/\mathrm{COP}_{\mathrm{ref}}}{\mathrm{COP}_{\mathrm{ref}}}\right], \tag{51}$$

where the central configuration uses $\kappa = 0.065\,^{\circ}\mathrm{C}^{-1}$ and $\mathrm{COP}_{\mathrm{ref}} = 3$, following the configured interpretation of Kou et al. (2026); at the warming levels of the pilot cities the correction is of the order of 0.1 %. The sensitivity S is taken at the upper end of the city-scale evidence. Modelling studies give a median of $0.017\,\mathrm{K}$ per $\mathrm{W\,m^{-2}}$ of anthropogenic heat flux (interquartile range 0.008–0.029; Li et al., 2026); the observed weekday-weekend contrast in Tokyo and Osaka gives 0.007–0.009 (Kikegawa et al., 2014), Paris 0.005 per $\mathrm{W\,m^{-2}}$ of peak flux (Tréméac et al., 2012; de Munck et al., 2013), Berlin about 0.01 (Jin et al., 2020), ten central-European cities 0.006–0.007 (Karlický et al., 2026), and the continental United States about 0.01 (Wang et al., 2023). The upper value is chosen because European cities are compact, and the uncertainty analysis samples S over 0.01–0.035, the upper bound corresponding to Phoenix under extreme heat (Salamanca et al., 2014).

On day d, the warming scales with that day's AC electricity, $\Delta T_{b,d} = \Delta T_{b,y}^{\mathrm{full}}\, e_d / e_{\mathrm{full}}$, where e_d is the household electricity of Equation (31) at the city-mean temperature of day d. For each AC branch, Notebook 05 estimates the marginal mortality response $M_{b,d,y}$ on each day by uniformly adding 1 °C to the daily hazard and recalculating

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

mortality under that branch's spatial AC-coverage field. The annual waste-heat mortality penalty is then

$$H_{b,y}^{\mathrm{WH}} = \sum_{d=1}^{n_y} \Delta T_{b,d}\, M_{b,d,y} = \Delta T_{b,y}^{\mathrm{full}} \sum_{d=1}^{n_y} \frac{e_d}{e_{\mathrm{full}}}\, M_{b,d,y}, \tag{52}$$

where n_y is the number of days in year y. Weighting each day by its AC electricity places the penalty on the hot days on which AC operates and on which heat deaths concentrate. The electricity-weighted share of the marginal response, $\sum_d (e_d/e_{\mathrm{full}}) M_{b,d,y} / \sum_d M_{b,d,y}$, and the annual marginal response $\sum_d M_{b,d,y}$ are evaluated at the anchor years and linearly interpolated over the assessment horizon. Only the difference between the policy and current branches is attributed to the AC expansion policy:

$$\Delta H_y^{\mathrm{WH}} = H_{\mathrm{pol},y}^{\mathrm{WH}} - H_{\mathrm{cur},y}^{\mathrm{WH}}. \tag{53}$$

Waste heat is therefore represented as an incremental mortality penalty subtracted from the direct mortality benefit of policy-induced AC expansion. It is not treated as a monetary cost. Waste heat associated with existing AC penetration remains part of the current-AC reference branch. In the four pilot cities, the existing AC stock produces a daily-mean warming at full cooling activity of 0.009 °C in Copenhagen to 0.044 °C in Athens in 2020, and the income-targeted expansion adds at most 0.005 °C.

This flux formulation replaces the look-up table from AC penetration to nighttime warming used in earlier versions of the framework (Salamanca et al., 2014), which remains available as a configuration option for sensitivity analysis. The penetration levels of that table are the air-conditioned share of floor area in Phoenix at extreme-heat loads, about 8 $\mathrm{W\,m^{-2}}$ of daily-mean AC heat at 65 %; read as European household penetration, it implied roughly ten times the warming per $\mathrm{W\,m^{-2}}$ that city-scale studies support. The only study that carries AC waste heat through to mortality finds that it takes back about 7–8 % of the protective benefit of AC at 96–98 % household AC ownership in Kansai cities (Chua et al., 2023). All parameters of the electricity and waste-heat rules are set in a single module of the `cityheat` package, can be overridden for an individual city in its configuration file, and are documented with their sources in the repository.

EWS-AC overlap. EWS benefits are evaluated on the residual heat-attributable mortality burden that remains after current AC protection. The model applies a further overlap adjustment to represent the reduced scope for EWS action at higher AC penetration. For city c, year y, and efficacy scenario s, the adjustment factor is

$$\eta_{c,y,s}^{\mathrm{EWS-AC}} = 1 - \overline{C}_{c,y}^{\mathrm{cur}}\, \omega_{c,s}, \tag{54}$$

where $\overline{C}_{c,y}^{\mathrm{cur}}$ is the unweighted mean of the gridded current-AC coverage field and $\omega_{c,s}$ is the configured overlap parameter. This is the factor appearing in Equation (44). The implemented adjustment uses annual current-AC coverage: it is neither population-weighted nor calculated specifically over warning days. EWS efficacy is not recalculated under the income-targeted policy-AC expansion in the current implementation.

Vegetation-AC interactions. Vegetation interacts with AC through two separate channels. First, Notebook 08 applies a reduced-form vegetation adjustment to annual electricity use per AC user. For policy region r and year y,

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

the fractional reduction is

$$R_{r,y}^{\mathrm{elec}} = \beta_{\mathrm{GVI}} \Delta G_r \frac{M_{\mathrm{summer}}}{12} m_y^{\mathrm{elec}}, \tag{55}$$

where ΔG_r is the regional GVI uplift, M_{summer} is the configured number of summer months, and m_y^{elec} is the maturity factor used specifically for this interaction. In all four pilot configurations, $M_{\mathrm{summer}} = 3$. The coefficient β_{GVI} is resolved at runtime for each city from the city-masked mean JJA daily-maximum T2M over the configured baseline UrbClim years. The framework applies piecewise-linear interpolation between the temperature-electricity-response anchors 0.000, 0.001, 0.008, and 0.020 per GVI point at JJA mean daily-maximum temperatures of 20, 25, 30, and 35 °C, respectively, with values clamped to this temperature range, following the reduced-form interpretation of Falchetta et al. (2026a). The completed pilot-city runs give $\beta_{\mathrm{GVI}} = 0.00927$ for Rome, 0.01314 for Athens, 0.00331 for Lisbon, and 0.00000 for Copenhagen, corresponding to JJA mean daily-maximum temperatures of 30.53, 32.14, 26.65, and 20.00 °C, respectively. If the required UrbClim daily-maximum fields are unavailable, or if a manual configuration override is requested, the configured value 0.008 is used as a fallback.

The interaction-specific maturity factor is

$$m_y^{\mathrm{elec}} = \mathrm{clip}\left(\frac{y - 2020 + 5}{12}, 0, 1\right). \tag{56}$$

This is a simple calendar-year maturity profile and does not include the cohort-based 25-year planting rollout used for the main tree benefit and cost streams. Adjusted annual electricity use per user is

$$e_{r,y}^{\mathrm{veg}} = e_{r,y}\left(1 - R_{r,y}^{\mathrm{elec}}\right). \tag{57}$$

Standalone AC costs are calculated without this adjustment. In the explicit AC-plus-trees comparison, the vegetation adjustment is applied to both current and policy AC users before their electricity consumption is differenced. The resulting reduction in incremental policy electricity expenditure is reported as an AC operating-cost interaction. Electricity savings for current AC users and for all AC users in the combined scenario, together with the associated reduction in electricity-related carbon dioxide emissions, are also reported as vegetation co-benefits; they are not subtracted from the standalone tree-policy cost.

Second, vegetation cooling lowers the AC electricity, and hence the waste heat, of each day. Notebook 08 loads the approximately monthly T2M-equivalent cooling maps exported by Notebook 07 and calculates their unweighted spatial mean over the physical vegetation-policy domain. At each modelled anchor year, the monthly cooling is scaled using the maturity profile in Equation (56) and added to the daily city-mean temperature series. The mean daily electricity weights under the baseline and tree-cooled temperatures are

$$q_y^{\mathrm{base}} = \frac{1}{n_y} \sum_{d=1}^{n_y} \frac{e\left(\overline{T}_{d,y}\right)}{e_{\mathrm{full}}}, \tag{58}$$

$$q_y^{\mathrm{tree}} = \frac{1}{n_y} \sum_{d=1}^{n_y} \frac{e\left(\overline{T}_{d,y} + \overline{\Delta T}_{m(d),y}^{\mathrm{tree}}\right)}{e_{\mathrm{full}}}, \tag{59}$$

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

where $e(T)$ is the household AC electricity of Equation (31) on a day with city-mean temperature T and $\overline{T}_{d,y}$ is the unweighted spatial mean of the daily hazard. Their ratio defines

$$\lambda_y = \mathrm{clip}\left(\frac{q_y^{\mathrm{tree}}}{\max(q_y^{\mathrm{base}}, 10^{-9})}, 0, 1\right). \tag{60}$$

Values of λ_y are calculated at the four anchor years and linearly interpolated over the 25-year assessment horizon. When the required vegetation and hazard outputs are available, the incremental AC waste-heat penalty in the combined scenario is

$$\Delta H_y^{\mathrm{WH,tree}} = \lambda_y \Delta H_y^{\mathrm{WH}}. \tag{61}$$

This is a first-order interaction: vegetation modifies the electricity-weighted waste-heat penalty, but the framework does not iteratively recompute AC penetration, electricity demand, COP, vegetation cooling, and mortality until convergence.

3.10 Cost-benefit analysis

Notebook 08 integrates the outputs of the AC, EWS, and vegetation modules over a common 25-year assessment horizon, from 2020 to 2044. The central specification uses an annual discount rate of $r = 0.03$. Because avoided mortality is retained as a physical outcome rather than assigned a monetary value, the principal economic comparison implemented in v1.0 is a discounted cost-effectiveness analysis within the broader CBA framework. The model reports present-value costs, equivalent annual costs, cumulative avoided deaths, and cost per avoided death; it does not calculate a monetary benefit-cost ratio or net monetary benefit using a value of statistical life.

For a stream of monetary costs $C_{j,t}$ associated with policy j, present value is calculated as

$$\mathrm{PV}(C_j) = \sum_{t=0}^{T-1} \frac{C_{j,t}}{(1+r)^{\tau_{j,t}}}, \qquad T = 25, \tag{62}$$

where $\tau_{j,t}$ reflects the timing convention of the corresponding policy module. Tree and AC expenditures calculated in Notebook 08 are treated as end-of-year flows, whereas the EWS series imported from Notebook 06 treats the 2020 setup and operating expenditure as a time-zero flow. The equivalent annual cost is

$$\mathrm{EAC}_j = \frac{\mathrm{PV}(C_j)}{\dfrac{1-(1+r)^{-T}}{r}}. \tag{63}$$

Policy cost streams. AC costs are evaluated incrementally relative to the evolving current-AC trajectory. Current and policy coverage, population, and electricity use per user are interpolated annually from the 2020, 2030, 2040, and 2050 anchor values. The incremental user stock in year t is

$$U_t^{\mathrm{add}} = U_t^{\mathrm{pol}} - U_t^{\mathrm{cur}}. \tag{64}$$

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

New installation cohorts are calculated from non-negative annual increases in this stock:

$$U_t^{\text{new}} = \max\left(0, U_t^{\text{add}} - U_{t-1}^{\text{add}}\right), \tag{65}$$

with the complete incremental stock in 2020 treated as the first cohort. CAPEX is charged to each cohort when installed and again at the configured equipment-replacement interval if the replacement occurs within the 25-year horizon. Annual maintenance is charged on the contemporaneous incremental user stock as the configured fraction of unit CAPEX. Electricity expenditure is based on the non-negative difference between policy and current electricity consumption, multiplied by the city-specific tariff. The standalone AC present-value cost is consequently

$$\text{PV}(C_{\text{AC}}) = \text{PV}(C_{\text{CAPEX}}) + \text{PV}(C_{\text{maint}}) + \text{PV}(C_{\text{elec}}). \tag{66}$$

An interaction-adjusted AC cost is also calculated after applying the vegetation-related electricity reduction described in Equation (55). This adjustment changes the electricity expenditure attributable to incremental AC users; it does not alter AC CAPEX or maintenance expenditure.

Tree costs are calculated from the total summed regional GVI uplift, ΔI, defined in Equation (38). The full intervention is divided into 25 equal annual planting cohorts. Each cohort's CAPEX is charged in its planting year. O&M begins in the following year and is scaled by the cohort's maturity, using the configured 12-year maturity period, five-year starting age, and 25-year lifetime. Notebook 08 reads the CAPEX per unit of GVI uplift directly from `trees.capex_per_index_pt_eur`. The annual O&M cost per unit of GVI uplift is derived by multiplying this CAPEX scale by the ratio `om_per_tree_per_year_eur/capex_per_tree_eur`, following Equation (39). Thus, the calibrated per-tree reference values determine the O&M-to-CAPEX ratio; they are not used to derive CAPEX per unit of GVI uplift. The central tree cost is the sum of the resulting CAPEX and O&M present values. A high-O&M sensitivity is additionally constructed by rescaling the O&M stream so that its present value equals five times the tree CAPEX present value.

EWS costs are imported from the 25-year calculation in Notebook 06. They include the maturity-class-specific setup expenditure, fixed annual operating expenditure, and warning-day-dependent operating expenditure described in Sect. 3.9.3. The default analysis uses the Pavanello formulation, while the Chiabai specification is retained as an alternative. The imported EWS present value preserves the configured implementation ramp, annual population trajectory, warning-day trajectory, and EWS maturity class.

Mortality-benefit streams. All health benefits are measured relative to the relevant current-infrastructure reference trajectory. The common cross-policy endpoint is cumulative avoided heat-attributable deaths over 2020-2044.

For trees, anchor-year mortality benefits are linearly interpolated to annual values and multiplied by the cohort rollout and maturity factor described in Sect. 3.9. The standalone tree benefit compares current AC without the policy uplift in GVI against current AC with the tree policy. Notebook 07 also exports the conditional tree benefit calculated on top of policy AC, but the standalone tree row in the cross-policy cost-effectiveness table uses the current-AC comparison.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

For AC, the gross benefit is the annual mortality difference between policy AC and current AC before accounting for outdoor waste heat. The standalone net AC benefit is

$$A_{\mathrm{AC},t}^{\mathrm{net}} = A_{\mathrm{AC},t}^{\mathrm{gross}} - \Delta H_t^{\mathrm{WH}}, \tag{67}$$

where ΔH_t^{WH} is the incremental waste-heat mortality penalty in Equation (53). In the vegetation-conditioned AC diagnostic, this penalty is replaced by $\lambda_t \Delta H_t^{\mathrm{WH}}$, and incremental AC electricity expenditure is recalculated after vegetation cooling. This row represents AC performance conditional on the presence of trees; it does not add the tree policy's own health benefit or tree-policy cost and should therefore not be interpreted as the complete cost and benefit of an AC-plus-trees portfolio.

For EWS, Notebook 08 uses the central-scenario net avoided-death trajectory imported from Notebook 06. These benefits already include the implementation ramp, AC-overlap adjustment, and age-specific mortality displacement. Mortality displacement and residual-life-year conversion are specific to the EWS module in the current implementation; they are not applied to AC or tree benefits. life-years saved remain an additional EWS output, while avoided deaths provide the common denominator for comparison among all three policies.

The primary cost-effectiveness ratio for policy j is

$$\mathrm{CER}_j = \frac{\mathrm{PV}(C_j)}{\sum_{t=0}^{T-1} A_{j,t}}, \tag{68}$$

where $A_{j,t}$ is the applicable annual avoided-death stream. Monetary costs are discounted, whereas avoided deaths in this primary ratio are summed without discounting. Ratios are reported only when cumulative avoided deaths are positive. The standalone comparison includes trees under the central and high-O&M specifications, gross and net AC, and central EWS. The interaction-adjusted AC diagnostic is reported separately when the vegetation-electricity or vegetation-waste-heat interaction is available.

Budget-constrained portfolio search. The portfolio routine represents each intervention by a continuous policy-scale variable $s_j \in [0, 1]$, where zero denotes no additional implementation and one denotes the complete modelled policy. These quantities are fractions of full policy implementation, not shares of the available budget. The implemented search assumes that standalone costs and benefits scale linearly:

$$C(\mathbf{s}) = s_{\mathrm{tree}}\mathrm{PV}(C_{\mathrm{tree}}) + s_{\mathrm{AC}}\mathrm{PV}(C_{\mathrm{AC}}) + s_{\mathrm{EWS}}\mathrm{PV}(C_{\mathrm{EWS}}), \tag{69}$$

$$A(\mathbf{s}) = s_{\mathrm{tree}}A_{\mathrm{tree}} + s_{\mathrm{AC}}A_{\mathrm{AC}}^{\mathrm{net}} + s_{\mathrm{EWS}}A_{\mathrm{EWS}}^{\mathrm{net}}. \tag{70}$$

The portfolio search therefore uses the standalone net AC, standalone tree, and central EWS results. The vegatation-related electricity saving and λ_t waste-heat adjustment are reported in the explicit interaction diagnostics but are not incorporated into every candidate portfolio in the current linear optimisation.

For the reference calculation, each policy scale is evaluated from 0 to 1 in increments of 0.02 under a present-value budget constraint of €1.5 billion. Among the feasible combinations, the maximum-benefit solution is the combination

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

with the largest cumulative avoided-death total. A second solution minimises average cost per avoided death among combinations delivering at least 50 % of the maximum feasible benefit. The search is repeated for eight budgets spanning €0.5-2.0 billion using a coarser scale increment of 0.04.

A portfolio is classified as Pareto-efficient if no other feasible combination has both a lower present-value cost and a greater avoided-death benefit. For consecutive points ordered along the resulting frontier, the incremental cost-effectiveness ratio is

$$\mathrm{ICER}_k = \frac{\mathrm{PV}(C_k) - \mathrm{PV}(C_{k-1})}{A_k - A_{k-1}}. \tag{71}$$

Additional diagnostics impose stylised implementation constraints, including a minimum 20 % tree-policy scale and a maximum 80 % AC-policy scale. These are scenario diagnostics rather than optimised political or behavioural constraints inferred from observed city decisions.

Distributional accounting and sensitivity. Distributional results are retained alongside the aggregate cost-effectiveness analysis. Notebook 05 reports AC benefits using both equal-area and population-weighted SVI quintiles, while Notebook 07 reports the distribution of vegetation benefits and the SVI-weighted tree-allocation sensitivity. These distributional metrics are not currently included as an objective or constraint in the portfolio search.

Notebook 08 also applies an explicit public-private cost-allocation assumption. Tree and EWS expenditures are assigned entirely to the public sector. For AC, 20 % of CAPEX is assigned to a public subsidy and 80 % to households, while maintenance and electricity expenditure are assigned entirely to the private sector. These shares are accounting assumptions in the present notebook rather than parameters inferred from city-specific subsidy programmes (Malloy and Ashcraft, 2020; Walker et al., 2024; Brousseau et al., 2024).

Discount-rate sensitivity is reported at 1 %, 3 %, 5 %, and 7 %. Tree CAPEX and O&M streams are rediscounted using their year-specific profiles. The AC and EWS diagnostic values are approximated by rescaling their central present values using the corresponding annuity-factor ratio; this sensitivity should therefore be interpreted as a ranking diagnostic rather than as a complete recalculation of every underlying policy cash flow.

Notebook 08 also reports three alternative accountings beside the central one. First, avoided deaths are discounted at the cost rate, with the same end-of-year factors as costs, whereas the primary ratio of Equation (68) divides discounted costs by undiscounted deaths. Second, the straight-line residual value of capital still in service at the end of the horizon, namely AC equipment from each cohort's last replacement within the horizon and trees planted late in the roll-out, is discounted from the final year and netted off the present-value cost. Third, the uniform AC expansion used as the distributional comparator is costed with the same CAPEX, replacement, maintenance, and electricity rules as the income-targeted policy, so that the cost side of the equal-users comparison is calculated rather than assumed. Fig. 9 reports the budget-constrained portfolio diagnostics.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

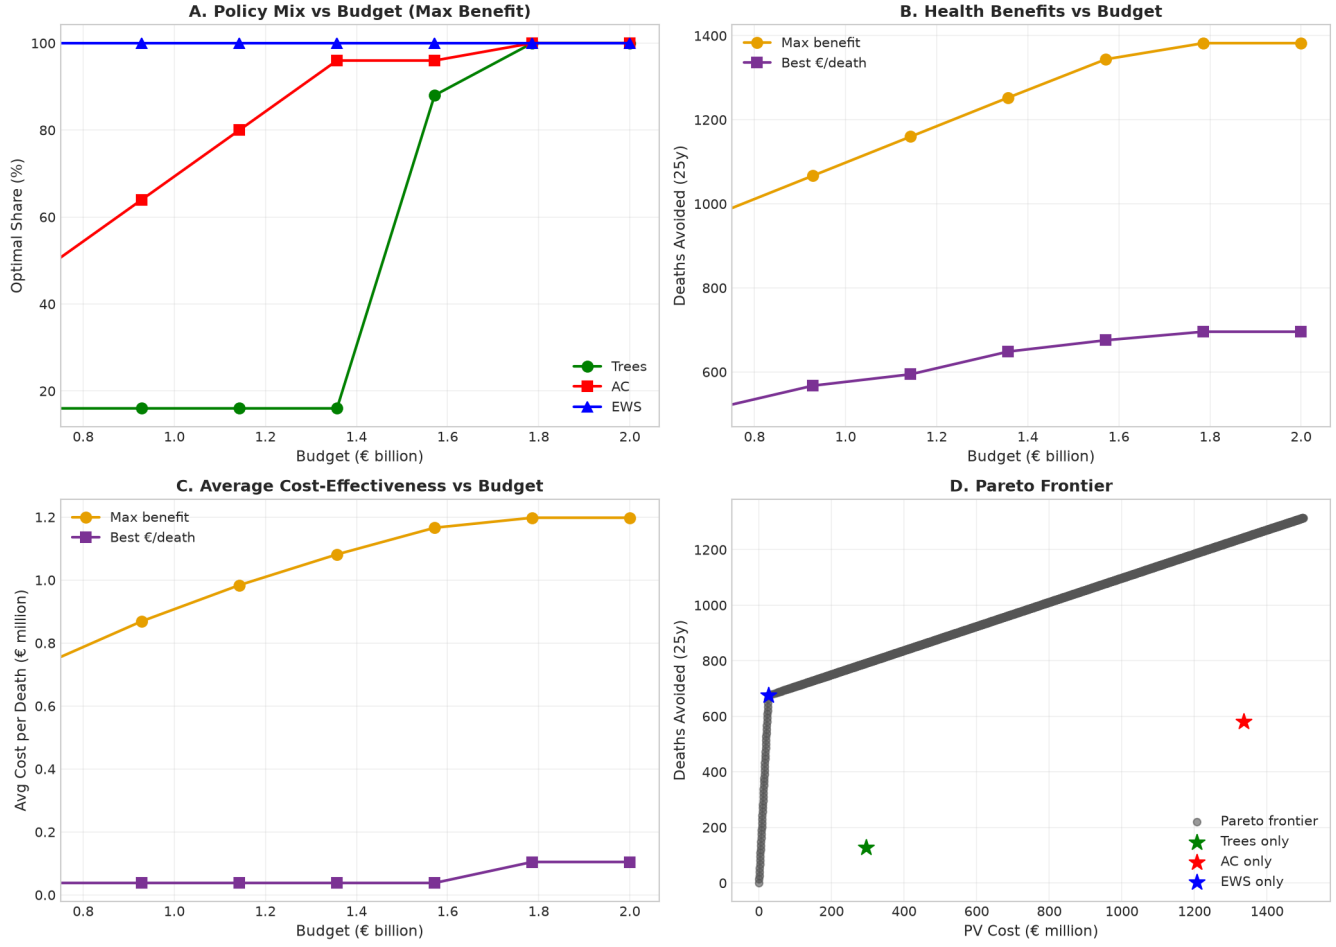


Figure 9. Budget-constrained portfolio diagnostics. Panel (A) reports the optimal scale of each full policy under the maximum-benefit objective across the tested budgets; the plotted percentages denote fractions of full implementation, not shares of expenditure. Panel (B) compares the cumulative 25-year avoided deaths delivered by the maximum-benefit solution and by the minimum-cost-per-death solution subject to the 50 % minimum-benefit condition. Panel (C) reports their average present-value cost per avoided death. Panels (A)-(C) are drawn from a budget of €0.75 billion. Panel (D) shows the Pareto-efficient combinations in present-value-cost-avoided-death space, together with the three full-scale standalone policies. The portfolio calculations use linearly scaled standalone costs and benefits and do not include the vegetation-AC interaction terms within individual candidate mixes.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.11 Uncertainty and sensitivity analysis

URBADAPT-HEAT propagates structural and parametric uncertainty jointly using a Latin Hypercube sample (LHS), implemented in Notebook 09 with the SciPy quasi-Monte Carlo sampler. The production uncertainty analysis reported here uses $N = 128$ Latin-hypercube draws. Under the canonical Masselot-main workflow, the standard daily-mean hazard track declares 57 input coordinates: 55 coordinates shared across the modelling chain and two additional EWS warning-trigger coordinates governing the target number of activation days and the threshold-recalibration interval. Some coordinates are conditional on other sampled choices and are therefore inactive for particular branches or draws; the count of 57 consequently describes the declared sampling space rather than 57 independently active processes in every model evaluation. Table B4 enumerates the sampled coordinates with their baseline values and sampled ranges.

Global sensitivity is evaluated using the PAWN distribution-based method with ten conditioning intervals. For each parameter and outcome, the implementation exports the minimum, mean, median, maximum and coefficient of variation of the Kolmogorov-Smirnov statistic. The manuscript summary figures use the median PAWN statistics. The same LHS draws are propagated to annual heat-attributable mortality, the daily heat-death distribution, the 25-year cost-effectiveness outcomes of AC, trees, and EWS, and the projected vulnerability indicators (Fig. 10). The ensemble medians of the cost-effectiveness outcomes sample the adaptation settings (waste-heat switch, efficacy, penetration, EWS interpretation) and therefore differ from the single central specification of Table 7. Some coordinates known not to enter a given outcome nevertheless appear among the largest PAWN scores (Fig. 10e–g); their non-zero values reflect finite-sample screening variability rather than a modelled causal pathway, and these rankings are exploratory and have not been convergence-tested.

Impact function scope

The headline LHS samples only the two Masselot specifications used in the main analysis: the constant extrapolation above the city-specific historical maximum temperature (`masselot`) and the non-negative log-linear extrapolation (`masselot_tail`). The impact-function-family coordinate therefore represents the structural uncertainty associated with extrapolating the Masselot curves beyond their observed temperature ranges.

The Burke polynomial and power-law functions are excluded from the headline distribution. For the $N = 128$ analysis reported here, Notebook 09 also runs two conditional Latin-hypercube experiments, on the same sample size, with the impact-function family fixed separately to `burke_polynomial` and `burke_powerlaw`. All other applicable uncertainty coordinates remain sampled in these experiments. The resulting Burke ranges are therefore conditional uncertainty bands rather than deterministic point estimates (Fig. 10c,h). These experiments reuse the 128 headline draws with only the impact-function family changed. For Rome, the median [P5–P95] annual heat-attributable mortality is 488 [303–836] deaths for Masselot, 65 [0.019–542] for the Burke polynomial and 50 [0.18–538] for the

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Burke power-law family; the wider Burke ranges mainly reflect the reference temperature, which is sampled over 18–26 °C in the Burke branches but fixed at 20 °C for Masselot.

Within each impact-function branch, the model samples independent multiplicative MDD factors over $[0.80, 1.20]$ for the < 15, 15-64, and 65+ exposure groups, a mortality-displacement fraction over $[0.00, 0.30]$, and a PAA scale over $[0.80, 1.00]$. Here PAA denotes the CLIMADA percentage, or fraction, of affected exposure; it is not a population-attributable fraction. Because the Masselot reconstruction assigns zero heat mortality risk to the < 15 group, the sampled < 15 MDD multiplier has no effect in the Masselot branches. A reference-temperature coordinate is retained for compatibility with the Burke workflow, but it is decoded as 20 °C and does not shift the Masselot curves, which already incorporate city- and age-specific minimum-mortality temperatures.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

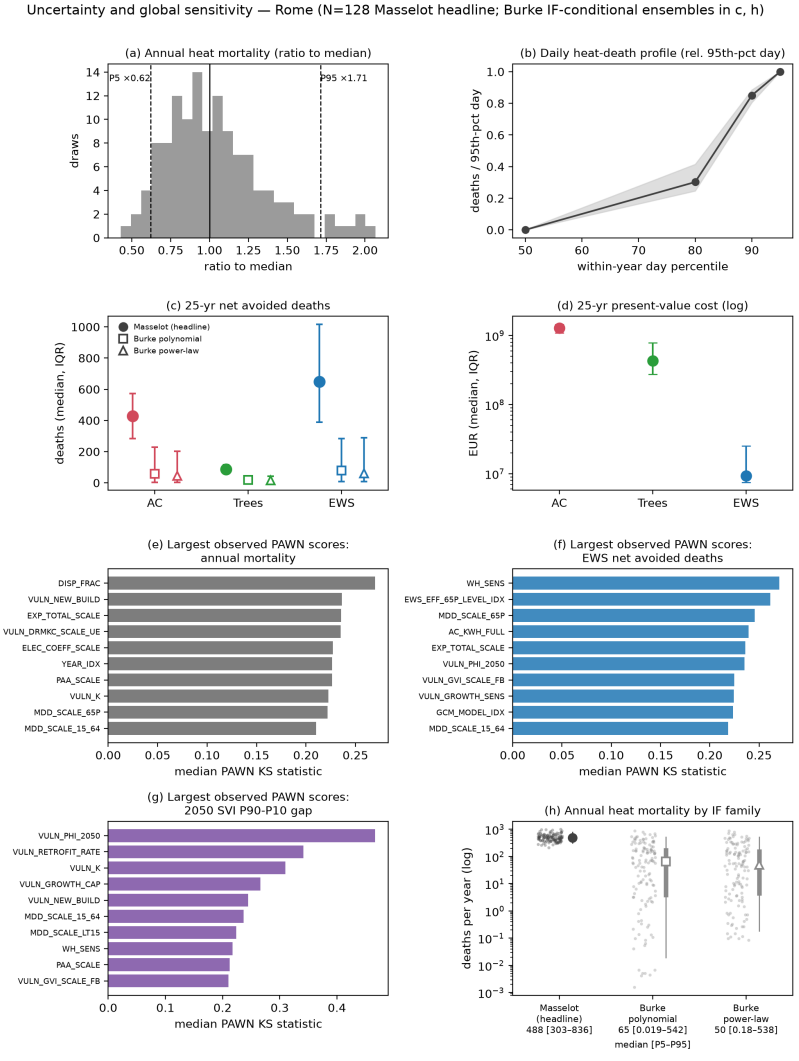


Figure 10. Uncertainty and global sensitivity for Rome under the Masselot headline Latin Hypercube design ($N = 128$ draws; the Burke families are evaluated in separate IF-conditional experiments, shown in (c) and (h), Sect. 3.11). (a) Aggregate annual heat-attributable mortality as a ratio to its median (median = 1; 5th/95th percentiles ×0.62/×1.71). (b) Within-year daily heat-death profile relative to a typical 95th-percentile day, with its 5th-95th-percentile band. (c,d) Propagated 25-year cumulative net avoided deaths and present-value costs (median and interquartile range; cost on a logarithmic scale) for the three pathways. In (c), filled circles show the Masselot headline ensemble and open squares and triangles the Burke polynomial and power-law IF-conditional ensembles; costs do not depend on the impact function and are shown once in (d). (e–g) The ten largest observed median PAWN Kolmogorov-Smirnov statistics for annual mortality (e), 25-year net EWS-avoided deaths (f), and the 2050 SVI P90–P10 gap (g). (h) Annual heat-attributable mortality under each impact-function family on a logarithmic scale: individual draws, interquartile range (thick bar), 5th–95th-percentile range (thin bar) and median (marker). Parameter names follow the extended URBADAPT-HEAT sampling dictionary described in Table B4.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Hazard and exposure uncertainty. The sampled hazard and exposure coordinates comprise the analysis year, exposure SSP, total-exposure scaling over $[0.90, 1.10]$, baseline-hazard construction mode, climate scenario, low-central-high warming band, climate-information source, and GCM selection where per-GCM inputs are available. The climate scenarios are CurPol, GS, SP, and SSP5-8.5, as defined in Sect. 3.4. Climate-source and GCM coordinates are conditional: GCM identity affects the calculation only when the per-GCM source is selected, while unavailable scenario-band combinations are resolved using the central band recorded by the workflow.

Adaptation and economic uncertainty. AC uncertainty includes the SSP-dependent penetration trajectory, age-specific efficacy scenario, waste-heat enable switch, the waste-heat sensitivity S over $[0.01, 0.035]\,\mathrm{K}\,(\mathrm{W\,m^{-2}})^{-1}$, the full-activity household electricity e_{full} over $[10, 20]\,\mathrm{kWh}$ per day, COP-degradation enable switch and case, electricity tariff, and equipment lifetime. The sampled e_{full} also sets each member's AC electricity use, and hence its electricity cost, recalculated from the member's own daily temperatures. Unit CAPEX is sampled as 0.8, 1.0, or 1.2 times the calibrated city-specific value in the active configuration. Annual maintenance remains the configured city-specific maintenance rate multiplied by the sampled CAPEX, and equipment lifetime is sampled over 9, 12, and 16 years.

Tree uncertainty includes the cooling-coefficient scale over $[0.50, 1.50]$, a tree-policy intensity coordinate over $[0.06, 0.20]$, maturity-ramp duration, starting tree age, a CAPEX multiplier of 0.8, 1.0, or 1.2, and an O&M multiplier of 1 or 5. The policy-intensity coordinate rescales the Notebook 07 cooling maps relative to the configured central value `trees.cap_uplift_0_1` $= 0.12$; it is not the maximum regional GVI increase imposed by the Notebook 07 Q3 catch-up policy. The cost multipliers are applied to the calibrated per-GVI-point cost quantities used by the active city configuration.

EWS uncertainty includes the interpretation class, efficacy levels, AC-overlap level, age-specific mortality-displacement levels, implementation-ramp duration, variable-cost formulation, target activation days, and threshold-recalibration interval. The interpretation coordinate samples a local maturity-class bracket around the class assigned in the city configuration: a marginal system samples marginal and intermediate interpretations; an intermediate system samples marginal, intermediate, and counterfactual interpretations; and a counterfactual system samples intermediate and counterfactual interpretations. This coordinate changes the epidemiological effectiveness attributed to the observed system. It does not replace the city's configured infrastructure class or its fixed setup and operating costs. The Pavanello and Chiabai formulations provide alternative warning-day-dependent cost specifications.

The discount rate is sampled over 2 %, 3 %, and 5 %. The vegetation-related reduction in AC electricity demand is represented by an enable-disable coordinate and a coefficient multiplier over $[0.50, 1.50]$ around the runtime city-specific coefficient described above.

Vulnerability uncertainty. The vulnerability projection contains ten sampled coordinates: the persistence parameter k over $[0.55, 0.95]$; the 2050 spatial-smoothing parameter ϕ_{2050} over $[0.50, 0.90]$; DRMKC scaling for foreign-born population over $[0.02, 0.08]$ and unemployment over $[0.04, 0.16]$; GVI scaling for the same two components over $[0.15, 0.55]$ and $[0.25, 0.75]$, respectively; the building-retrofit rate over $[0.005, 0.020]$; vulnerability-growth sensitivity over $[0.50, 1.00]$; the cumulative vulnerability-growth cap over $[0.25, 0.45]$; and new-build vulnerability over

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

[0.10, 0.25]. The sampled exposure SSP also selects the corresponding socioeconomic trajectory used by the future vulnerability projection.

In URBADAPT-HEAT v1.0, these vulnerability perturbations are output-only uncertainty dimensions: they alter the projected SVI distribution and equity indicators but do not modify the mortality calculations. This separation should be retained when interpreting the mortality and vulnerability PAWN results.

Spatial propagation. Vegetation cooling is propagated through spatially explicit daily hazard adjustments: the monthly cell-level cooling maps of Notebook 07 are scaled by the sampled policy intensity and maturity and translated to air temperature day by day at each cell's daily temperature, which preserves the within-city heterogeneity of the cooling. AC waste heat is not added to the hazard field. As in Notebook 05, each ensemble member evaluates Equations (50)–(53) with its own sampled S, e_{full} and COP-degradation case, its own AC users and daily city-mean temperatures, and marginal mortality responses obtained from $+1\,^{\circ}\text{C}$ perturbations of its own current-AC and policy-AC branches. Waste heat therefore enters as a city-wide mortality penalty rather than as a spatially resolved temperature change.

To assess the sample-size sensitivity of the annual-mortality distribution, we compare the reported $N = 128$ design with a separately generated $N = 64$ Latin-hypercube design using the same model engine, input definitions, and random seed (Table B5). Across the four pilot cities, the median annual-mortality estimates differ by at most 3.5 %, and P5 estimates by at most 15.7 %; the $N = 64$ P95 estimates are 3.2–35.9 % higher. We use the larger $N = 128$ design for the reported distributions. This two-design comparison describes sample-size sensitivity but does not establish convergence of mortality quantiles, PAWN rankings, or adaptation cost-benefit outcomes.

3.12 Verification and evaluation

The principal verification test is the reconstruction of the published Masselot exposure-response curves described in Sect. 3.7. Applying the reconstructed relative-risk curves to the daily ERA5-Land temperature series used by Masselot et al. (2023), and reproducing the published attribution convention, returns the published annual heat-attributable totals to numerical precision, with percentage differences of order 10^{-12} %. The operational URBADAPT-HEAT export additionally imposes the non-negative heat-only constraint of Equation (15); across the four pilot cities this shifts the reconstructed annual totals by at most approximately 0.064 %. The residual discrepancy is therefore attributable to a deliberate modelling constraint rather than to numerical error in the implementation.

Evaluation of the hazard-curve reconciliation. Because the impact-function reconstruction is exact, any disagreement between modelled baseline mortality and the published city-level estimates of Masselot et al. (2023) must arise in the interface between the hazard and the curve. The two corrections of Equations (1) and (19) address that interface, and they were evaluated by running the complete workflow for 40 European cities four times: without correction, with each correction alone, and with both. Table 6 compares the resulting modelled 2020 baseline mortality rates with the published estimates. The two corrections act on different features of the disagreement and are only jointly sufficient. The variance correction governs the dispersion, reducing the ratio of the largest to the smallest

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 6. Modelled 2020 baseline heat mortality per 100 000 inhabitants relative to the published estimates of Masselot et al. (2023), for 40 European cities, under each combination of the two hazard-curve corrections. The spread is the ratio of the largest to the smallest city-level ratio.

Configuration	Median ratio	Spread	Within a factor of two
No correction	0.58	87.3	19/40
Variance only	1.64	5.0	30/40
Reference only	0.25	66.5	9/40
Both corrections	0.98	3.2	39/40

city-level ratio from 87 to 5, while leaving a systematic overestimate. The reference correction governs the level but barely reduces the dispersion, and applied alone it is worse than no correction at all, because in the uncorrected configuration the two biases partly cancel. Applied together they return a median ratio of 0.98 with 39 of 40 cities within a factor of two.

The variance correction was additionally verified against the exact calculation it approximates. At the level of the response curve, comparing Equation (19) with the mean of the responses to the individual observed years gives a median error of 1.6 % and a worst case of 9.5 % across nine cities spanning the three climate classes of the sample. Because that comparison evaluates the curve at city-mean temperatures, it attributes the entire convexity gap to interannual variance, whereas the model integrates over grid cells and part of the same effect is already supplied by the within-city spatial temperature distribution; it is therefore an upper bound on the correction rather than a prediction of it. The correction was consequently also verified inside the complete workflow, by driving each observed year separately through the full chain for the city of the wider sample in which the two measures differ most and averaging the resulting baseline mortality; Equation (19) reproduces that exact multi-year mean to 0.17 %.

Three invariance checks confirm that the corrections act only on the historical reference. The reference shift of Equation (1) is identical across target years to within $1.7\times10^{-3}\,^{\circ}$C; the projected warming increments are reproduced exactly in the corrected configuration (+0.383, +0.581 and +0.813 °C in 2030, 2040 and 2050 respectively); and the population-weighted cooling generated by the tree module is identical to within $10^{-9}\,^{\circ}$C across all 40 cities and scenarios.

Two components are evaluated out of sample by cross-validation, as described in Sect. 3.8. For the climate emulator, refitting while holding out one complete city at a time gives, across the 130 held-out-city evaluations, an unweighted mean predictive R^2 of 0.718, a mean RMSE of 4.68 °C and a mean absolute error of 3.62 °C. For the hottest decile of observed land-surface-temperature days the mean RMSE rises to 6.66 °C, so predictive accuracy is weakest on precisely the days that dominate heat-attributable mortality. For the income emulator, refitting while holding out one complete country at a time, the zone-count-weighted mean of the city-specific population-weighted Spearman correlations between observed and predicted zone ordering is 0.592, and 0.597 when out-of-fold predic-

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

tions are pooled within each held-out country before scoring. The out-of-fold identification rates for the poorest and richest observed fifths of zones are 0.537 and 0.453 respectively, against 0.20 under random assignment. Performance varies substantially among countries, so emulator-derived ranks are interpreted throughout as proxy socioeconomic distributions rather than as substitutes for locally observed income data.

4 Model implementation and illustrative results in four European cities

4.1 Study cities and configuration

To demonstrate its by-design scalability, URBADAPT-HEAT is applied here to four European cities spanning contrasting climatic, demographic, socio-economic, and urban morphological conditions: Rome (Italy), Athens (Greece), Lisbon (Portugal), and Copenhagen (Denmark). The sample combines a Mediterranean heat-exposed cluster (Rome, Athens, Lisbon) with a northern-latitude cool-climate city (Copenhagen), where heat-attributable mortality is lower in absolute terms but concentrated in summer extremes.

For each city the deterministic pipeline was executed from data harmonisation through CBA (Notebooks 01-08) with city-specific YAML configuration files and data manifests but identical analytical code. Input data (UrbClim T2M, WorldPop, Census 2021, GHS-AGE, GVI) were obtained from the sources described in Sect. 3, with city-specific NUTS3 identifiers anchoring AC penetration and EWS cost parameters. All four cities run on the standard daily-mean hazard track with a deaths-threshold EWS activation logic calibrated city by city against the reference-year heat-deaths series, so trigger-based indicators (warning-day counts, threshold-exceedance metrics) are directly comparable across the pilots. An optional event-exceedance hazard representation (Track-B), with event-mask warning activation and a cold-city impact-function variant, remains configurable in the city YAML files for extreme-event or cold-climate applications, but is disabled at runtime under the Masselot-main canonical configuration, because the per-city minimum-mortality temperatures embedded in the Masselot impact functions already give a well-defined heat-attributable signal on the standard track in every pilot, including Copenhagen. The city-specific vulnerability and calibration parameters used for the pilots are reported in Table B3.

4.2 Impact functions for the pilot sample

The canonical impact functions are city-specific Masselot curves anchored on per-city minimum-mortality temperatures (MMT; see Sect. 3.7). Across the three older-age curves, the MMT ranges are 20.32-21.50 °C in Rome, 23.43-24.36 °C in Athens, 19.23-20.48 °C in Lisbon, and 18.17-18.42 °C in Copenhagen, so the sample brackets the broad south-to-north acclimatisation gradient, and the age-specific variation within each city, without recourse to city-specific functional-form switches.

The cities differ, however, in the empirical basis of their Masselot coefficients. Rome, Athens, and Lisbon are MCC cities, so their final coefficients remain informed by locally observed daily mortality data after meta-regression and

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

spatial smoothing. Copenhagen is not an MCC city: its coefficients are predictions from the meta-regression and spatial co-kriging stages, conditional on the city's characteristics and location but without a Danish mortality time series entering the first-stage estimation. This introduces additional structural uncertainty for Copenhagen, revisited in Sect. 5.1. Comparability with the alternative impact-function families is likewise uneven: for Rome, Athens, and Lisbon the two Burke families can be compared with Masselot on the standard daily-mean absolute-temperature scale, whereas Copenhagen's earlier Burke implementation was calibrated in a cold-city event-exceedance coordinate system, so its Burke-Masselot comparison is not a directly equivalent fixed-temperature comparison and requires separate interpretation.

4.3 Illustrative results and framework behaviour across contrasting city contexts

URBADAPT-HEAT is designed to support three classes of urban adaptation decision that, taken together, are difficult to inform from any existing single-purpose tool: (i) within-city prioritisation across heterogeneous adaptation pathways with explicit cost-benefit accounting; (ii) cross-city comparison of the drivers of heat-attributable burden and of the effectiveness of comparable policies; and (iii) equity and distributional analysis of who bears costs, who receives benefits, and how the public-private split affects the political feasibility of each pathway. The four-city pilot illustrates each in turn.

The city-specific Masselot impact functions and age-stratified exposure trajectories produce substantial differences in the modelled burden (Table 7). In the climate-only branch before current-AC attenuation, the 2050 heat-attributable burden among people aged 65+ is approximately 1,196 deaths in Athens, 1,044 in Rome, 326 in Lisbon, and 59 in Copenhagen; applying the evolving current-AC reference reduces the corresponding 2050 all-age totals to 915, 756, 280, and 55, respectively. These magnitudes arise jointly from the city-specific exposure-response curves and minimum-mortality temperatures, the daily temperature distributions, population size, and age structure, and should not be interpreted as isolating the causal contribution of any one factor.

Propagating the structural and parametric dimensions of Sect. 3.11 through the headline Masselot Latin Hypercube design ($N = 128$) yields distributions of aggregate annual heat-attributable mortality – pooled across the sampled target years, climate scenarios, and model parameters, rather than uncertainty interval around a single fixed-year deterministic estimate. The median [5th-95th percentile] annual heat-death distributions are 613 [399-1,004] in Athens, 488 [303-836] in Rome, 169 [110-285] in Lisbon, and 30.2 [15.7-70.3] in Copenhagen. Across the three Mediterranean pilots, P5 and P95 are approximately 0.62–0.65 and 1.64–1.71 times the respective median; Copenhagen has the widest relative range, approximately 0.52–2.33 times its median. These distributions propagate the sampled hazard, exposure, impact-function-extrapolation, and adaptation dimensions jointly. They do not include the Masselot coefficient-simulation uncertainty, whose propagation is identified as a priority for a future version (Sect. 5.1).

Two classes of interaction are reported explicitly rather than being silently absorbed into standalone pathway accounts. First, the CBA module quantifies the incremental AC waste-heat externality: over 2020-2044 the penalty

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

is approximately 13 deaths relative to 593 gross AC-avoided deaths in Rome (2.1 %), 4.3 relative to 255 in Athens (1.7 %), 4.3 relative to 144 in Lisbon (3.0 %), and 0.6 relative to 46 in Copenhagen (1.2 %). Second, the vegetation-AC electricity interaction is city-specific: for incremental policy AC users, the runtime temperature-calibrated vegetation feedback reduces present-value electricity expenditure by approximately €1.76 million in Rome, €0.73 million in Athens, and €0.33 million in Lisbon, and is zero in Copenhagen under the central run because its calibrated coefficient is zero at the lower 20 °C temperature anchor.

The income-targeted and uniform AC policies add the same number of users at each anchor year by construction, but income targeting is not assumed to maximise aggregate mortality reduction. Relative to the uniform allocation, the income-targeted policy changes cumulative gross AC benefits over 2020-2044 by approximately −0.8 % in Rome, +6.6 % in Athens, −15.3 % in Lisbon, and −7.0 % in Copenhagen. Because the Masselot exposure-response function is spatially uniform within each city, these differences arise from the spatial alignment of coverage, population, and hazard rather than from a steeper epidemiological response assigned to lower-income residents; the alternative allocation is therefore an explicit distributional policy experiment, not a guaranteed efficiency improvement. SVI-stratified benefit diagnostics are reported for the spatial AC and tree analyses, and the public-private expenditure split as a separate accounting layer, whereas the current EWS implementation is age-specific and citywide rather than spatially targeted by SVI.

Figure 11 summarises the annual avoided-death trajectories of the three adaptation pathways over the 25-year horizon. EWS benefits are initially attenuated by the configured three-year implementation period, during which efficacy is set to 10 % of its mature value, and increase from 2023 when full efficacy is applied; their subsequent evolution reflects changes in warning-day mortality, warning-day frequency, population, AC overlap, and the underlying heat burden. AC benefits follow the exogenous penetration trajectory and the spatial allocation of additional users, while tree benefits accrue gradually through the 25-year planting rollout and the configured 12-year maturity profile. Table 7 reports the corresponding headline outputs for the four pilot cities.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

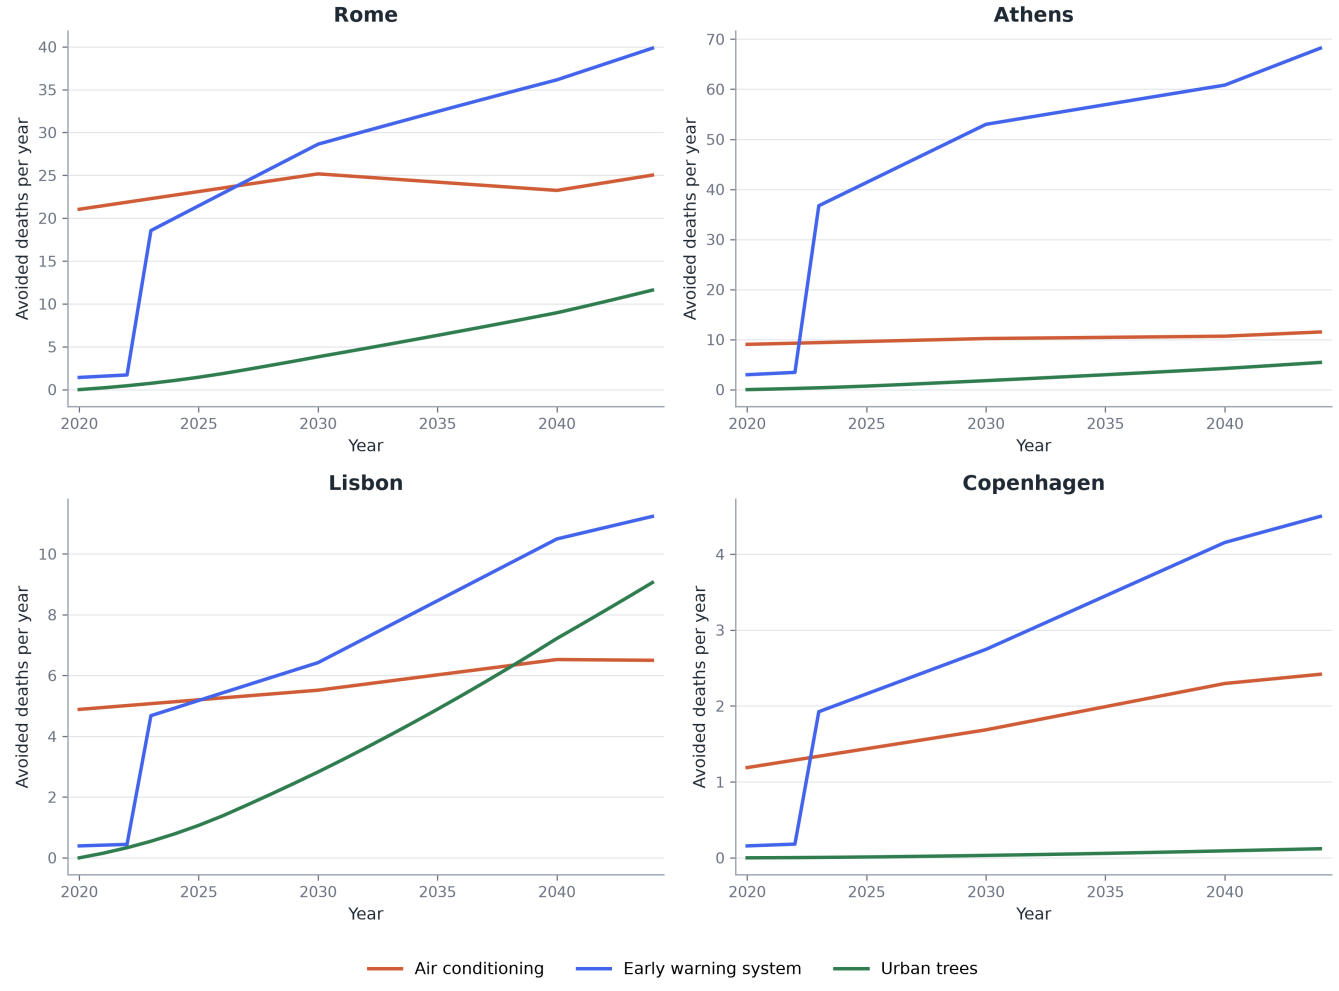


Figure 11. Twenty-five-year adaptation-benefit trajectories for the four pilot cities under the central deterministic specification. Each panel reports gross AC avoided deaths relative to current AC, net EWS avoided deaths after the implementation ramp, AC-overlap adjustment, and age-specific mortality displacement, and standalone tree avoided deaths under the planting-rollout and maturity profiles. The EWS interpretation is marginal in Rome and Lisbon and counterfactual in Athens and Copenhagen. Vertical scales differ among cities.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Metric	Rome	Athens	Lisbon	Copenhagen
Applied annual-mean ΔT2M, 2050−2020 (°C)	0.70	0.73	0.64	0.78
Heat deaths before current-AC attenuation, 2020 (yr^{-1})	661	848	223	27
Heat deaths before current-AC attenuation, 2050 (yr^{-1})	1075	1259	348	62
AC avoided deaths, 25 y (gross / net)	593 / 580	255 / 251	144 / 140	46 / 45
EWS net avoided deaths, 25 y	675^{m}	1203^{c}	176^{m}	72^{c}
Trees avoided deaths, 25 y	127	60	98	1.2
Cost per avoided death (€ million)				
Aggregate standalone ratio	1.20	0.32	4.24	10.7
Air conditioning (net)	2.30	1.35	3.89	22.1
Early warning system	0.038	0.023	0.105	0.185
Urban trees	2.33	2.05	12.16	210

Table 7. Illustrative outputs for the four pilot cities under the canonical Masselot log-linear-tail impact functions and central deterministic specification. Copenhagen's Masselot coefficients are predictions from the meta-regression and spatial co-kriging stages rather than estimates informed by a Danish mortality series, which carries additional structural uncertainty for its entries in this table; see Sect. 4.2. The temperature change is the day-weighted annual mean of the monthly CurPol central-band increments applied in Notebook 02. The two annual-mortality rows report the climate-only branch before attenuation by existing AC; by contrast, all adaptation benefits are evaluated against the evolving current-AC reference branch. Adaptation benefits are undiscounted cumulative avoided deaths over 2020-2044. EWS uses the marginal interpretation (^{m}) in Rome and Lisbon and the counterfactual interpretation (^{c}) in Athens and Copenhagen; its cross-city magnitudes should therefore not be interpreted as estimates of one homogeneous intervention. Cost-effectiveness divides 25-year present-value expenditure by undiscounted cumulative avoided deaths. AC is net of its incremental waste-heat penalty, trees use base O&M, and EWS uses the central efficacy scenario. The aggregate standalone ratio is the sum of the three standalone present-value costs divided by the sum of their standalone benefits; it is an accounting diagnostic, not the result of the interaction-aware portfolio search.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Under the configured central specifications, EWS is the most cost-effective standalone pathway in each of the four pilot cities and also yields the greatest modelled cumulative avoided-death total in each of them (Fig. 12). Its present-value cost per net death avoided ranges from approximately €23,000 in Athens to €185,000 in Copenhagen, compared with approximately €1.35-22.1 million for net AC and €2.05-210 million for trees under base O&M, so that EWS dominates the full-scale AC and tree endpoints in the standalone cost-benefit plane of every pilot. This ranking is conditional on each city's configured EWS interpretation – marginal enhancement of a mature system in Rome and Lisbon, counterfactual comprehensive deployment in Athens and Copenhagen – and is not a comparison of one homogeneous EWS treatment across four cities. Among the two physical cooling measures, AC produces the greater gross mortality benefit in every city, and its waste heat takes back 1.2-3.0 % of that benefit; AC has the highest standalone present-value cost in Rome, Athens, and Copenhagen, whereas the base-O&M tree pathway has the highest cost in Lisbon. Both physical measures require substantially larger capital and operating expenditures than EWS and provide co-benefits that the mortality-only denominator does not represent, and both can nevertheless enter the portfolio search when additional benefits beyond full EWS implementation are sought or when stylised implementation constraints are imposed.

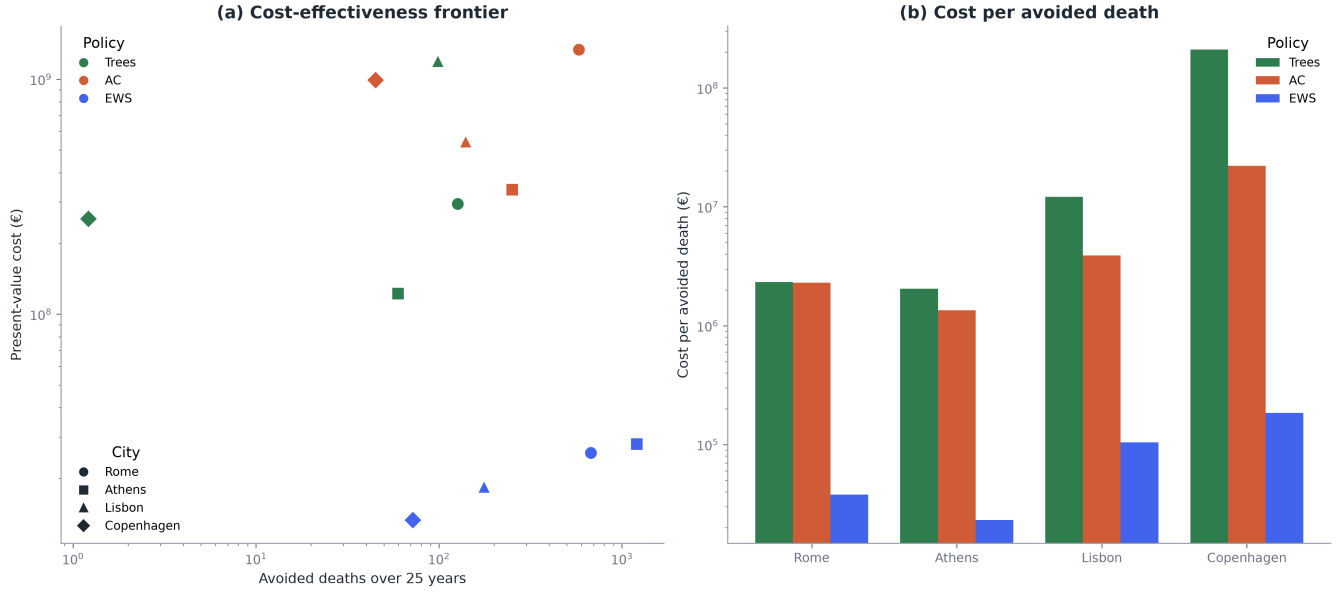


Figure 12. Cross-city cost-effectiveness results over the 25-year horizon. Panel (a) compares the standalone present-value cost and cumulative avoided-death endpoint of each policy and city. Under the city-specific central specifications, EWS has a lower cost than the standalone AC and tree endpoints in every pilot city, and a greater modelled benefit in each of them. Panel (b) reports present-value cost per avoided death on a logarithmic scale. AC is evaluated net of its incremental waste-heat penalty, trees use base O&M, and EWS uses the configured marginal or counterfactual interpretation.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Taken together, the pilots illustrate both the decision support the deterministic CBA delivers for city-level planners – standalone cost-effectiveness comparisons, budget-constrained policy scales, a cost-avoided-death Pareto frontier, incremental cost-effectiveness ratios, and stylised implementation-constraint diagnostics – and the boundaries of the current implementation. Distributional indicators are reported alongside the aggregate search but are not included as an optimisation objective or constraint and are not recalculated for every candidate portfolio mix; likewise, Notebook 09 is configured to propagate uncertainty to mortality, policy cost-effectiveness, and vulnerability endpoints without re-optimising the complete portfolio for every LHS draw. These separations preserve a transparent distinction between aggregate efficiency, uncertainty, and equity diagnostics.

5 Discussion

5.1 Limitations and priorities for future development

URBADAPT-HEAT v1.0 is designed as a transparent and extensible decision-support framework that advances urban-scale heat-risk assessment and adaptation appraisal. To place its contributions and residual limitations in perspective, Table B1 critically evaluates URBADAPT-HEAT v1.0 against the methodological frontier identified by the recent review and gap analysis of frameworks and methods for projecting urban temperature-related mortality under climate scenarios by Liu et al. (2026). This comparison identifies the following limitations, which should be made explicit when interpreting the results and guiding future model development.

(i) First, parts of the modelling chain are intentionally synthetic or anchored. The 2020 hazard year is a synthetic baseline derived from the historical UrbClim climatology rather than an observed 2020 weather year, and future warming is applied as monthly city-level deltas added uniformly to the daily fields within each month. This preserves baseline intra-city spatial contrasts while shifting absolute temperature levels, but it does not represent pixel-specific differential future warming patterns. On the epidemiological side, the canonical impact functions are faithful reconstructions of the city-specific B-spline DLNM curves published by Masselot et al. (2023), rather than re-estimations from raw mortality microdata, which lie outside the scope of an open-source decision-support framework. Notebook 09 is configured to propagate selected structural and parametric dimensions affecting the mortality calculation, including the Masselot extrapolation-tail formulation, age-specific MDD scaling, and mortality displacement, while Burke-conditional LHS runs are specified separately as structural sensitivity analyses. These dimensions do not replace propagation of the Masselot coefficient uncertainty discussed below. For exposure, the 2040 and 2050 population grids are obtained by scaling the 2030 spatial template with age-specific national growth factors from WCDE. Endogenous within-city relocation and redevelopment dynamics are therefore not modelled.

(ii) Second, the representation of adaptation options carries a set of stylisation choices that should be understood as structurally uncertain rather than locally exact. Several key calibration coefficients are transferred across contexts, including the temperature sensitivity to AC waste heat, the per-household cooling-electricity rule, the COP response structure, the EWS cost formulations, and the vegetation cooling bridge, and should therefore be inter-

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

preted as best-evidence parameters rather than city-specific local truths. AC penetration trajectories are exogenous NUTS-based inputs spatially downscaled within cities; the model does not endogenously solve household cooling adoption behaviour as a function of changing hazard, prices, or policy incentives. Flat or weakly changing trajectories in the source data, such as the Athens SSP2 trajectory around 2020-2030, are consequently inherited by design. Direct mortality protection is parametrised by the continuous local AC-coverage share through $1-\varepsilon_a c_i(y)$, whereas electricity use, which follows the number and intensity of cooling days, enters the energy-demand, cost, and waste-heat calculations separately. In v1.0, ε_a is therefore not a function of cooling-use intensity. Recent evidence from the contiguous United States estimates the temperature–mortality association as a continuous nonlinear function of household AC electricity use and suggests that the associated protection plateaus beyond median usage (Chu et al., 2026). Allowing AC efficacy, or more generally the temperature–mortality relationship, to vary with realised cooling use is a possible future extension, subject to the availability and transferability or corresponding usage–health evidence for European cities. The vegetation pathway uses an empirical ΔGVI$\rightarrow\Delta$LST$\rightarrow\Delta$T2M emulator bridge, which is practical and scalable but remains a statistical approximation to microclimatic and physical processes. In particular, the source UrbClim simulations do not prescribe supplementary irrigation of urban vegetation. The city-month-LCZ regressions can therefore produce weak or positive GVI-LST associations under hot and dry conditions, particularly during summer in water-limited Mediterranean settings, when modelled soil moisture becomes limiting and evaporative cooling declines. Because the operational implementation clips positive coefficients to zero, it neither attributes warming to additional trees nor represents the additional cooling that managed irrigation might sustain. At the nominal 100 m resolution, it also does not explicitly resolve pedestrian-scale canopy shading. Accordingly, zero modelled cooling denotes the absence of retained grid-scale cooling under these assumptions, not a general conclusion that street trees are ineffective. Irrigation requirements and their associated water, energy, and monetary costs are not represented in v1.0. Transfer performance also differs among the four pilots: Rome, Lisbon, and Copenhagen are represented in the 130-city fitted pool and have corresponding leave-one-city-out evaluations, whereas Athens is absent from that pool. Application of the pooled emulator bundle to Athens is therefore a fully out-of-sample spatial transfer. Additional adverse effects of adaptation policies, such as a possible increase in winter cold mortality in temperate regions resulting from increased street-tree cover Wang et al. (2025), are not represented because the present framework focuses on heat-related mortality. Finally, adaptation interactions are represented in first-order form: AC waste-heat effects are computed from a city-level heat flux - the electricity drawn by AC households on each cooling day, the heat their units reject, the Kou COP degradation, and a sensitivity of air temperature to that flux transferred from city-scale studies - and applied as a city-wide warming, but the framework does not currently solve fully iterative coupled atmosphere-behaviour feedbacks at each time step. The sensitivity is set at the upper end of the city-scale evidence and sampled up to the value implied for Phoenix, and the electricity rule gives annual totals above national statistics in the hottest pilots (about 1,500 kWh per AC household in Rome, against 230–490 kWh reported for Italy), so the modelled waste-heat penalty and AC electricity costs are more likely overstated than understated there.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

(iii) Third, the canonical Masselot impact functions carry structural limitations that propagate to all downstream results. The four pilot cities differ in the empirical basis underlying their Masselot exposure-response functions. Rome, Athens, and Lisbon are MCC cities whose final curves remain informed by locally observed daily all-cause mortality data after meta-regression and spatial smoothing. Copenhagen is not an MCC city; its coefficients are predictions from the meta-regression and ordinary spatial co-kriging stages, conditional on Copenhagen's characteristics and location, without a Copenhagen mortality time series entering the first-stage estimation. The Masselot tail-mode bracket sampled in the headline LHS addresses only the extrapolation of the exposure-response relationship beyond each city's historical maximum temperature. It does not quantify the additional uncertainty arising from Copenhagen's model-predicted coefficients.

The Masselot impact functions are also spatially uniform within each city: for a given age group, the same city-specific exposure-response curve applies to every grid cell inside the FUA. The framework can therefore stratify mortality and adaptation outcomes according to spatial differences in population, temperature, SVI, and policy coverage, but it does not estimate a cell-specific modification of the epidemiological dose-response relationship by social vulnerability.

Finally, the statistical uncertainty distributed with the Masselot model is not currently propagated through the headline LHS. The Zenodo archive provides 1000 coefficient-simulation draws in `coef_simu.csv` and per-curve variance-covariance matrices in `vcov.csv`, whereas URBADAPT-HEAT currently uses the published point estimates in `coefs.csv`. The Masselot tail-mode bracket, generic age-specific MDD scaling, and Burke-conditional sensitivity runs do not substitute for propagation of this coefficient and covariance uncertainty. Incorporating these statistical uncertainty inputs is therefore a clear priority for a future framework version.

(iv) Fourth, in the cost-benefit analysis module, annual costs and benefits over the 25-year horizon are interpolated from anchor-year model outputs, which smooths potential non-linear transitions between anchors. The core welfare metric is heat-mortality based; several co-benefits and disbenefits, including wider ecosystem, comfort, and urban-amenity effects, are only partially monetised or are reported outside the primary mortality-based objective. Relatedly, the framework deliberately reports mortality outcomes in physical terms, as avoided deaths, and as cost per avoided death, rather than collapsing benefits and costs into a single monetised net present value or benefit-cost ratio. A natural extension would monetise avoided mortality through a value of a statistical life (VSL), representing willingness to pay for a marginal reduction in mortality risk. Standardised European reference values exist for this purpose; for example, the EUROCONTROL Standard Inputs for Economic Analyses recommend an average VSL of €4.7 million for the EU27+UK in 2024 prices, drawing on the European Commission external-cost handbook European Commission, DG MOVE (2019); EUROCONTROL (2025). Applying a VSL would allow benefit-cost ratios and net present values to be reported directly, but it would also introduce first-order normative and empirical choices that materially affect results: whether to adopt a VSL or an age-adjusted value of a life year (VOLY), given that heat mortality is concentrated among older individuals with shorter residual life expectancy; how to transfer values across countries and income levels; and how to discount avoided deaths accruing decades ahead. Because

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

these choices belong to applied welfare appraisal rather than model description, monetised valuation is deliberately left beyond the scope of this paper and is instead identified as a configurable post-processing option for future applications.

(v) Fifth, uncertainty analysis in v1.0 is broad but not exhaustive. The workflow uses Latin Hypercube sampling with PAWN sensitivity analysis, a setup suited to mixed discrete-continuous input spaces. With $N = 128$ evaluations across 57 declared coordinates and ten PAWN conditioning intervals, the PAWN indices are interpreted as screening-oriented sensitivity diagnostics. The $N = 64$-$N = 128$ comparison evaluates the stability of annual-mortality distributional summaries and does not constitute a convergence assessment of the PAWN rankings. Moreover, PAWN outputs are marginal distribution-based sensitivities and do not directly provide explicit second-order interaction indices in Sobol form, while spatially explicit sensitivity fields are not yet a standard output. Within the same uncertainty chain, the ten sampled vulnerability coordinates affect the projected vulnerability endpoints but do not enter the heat-mortality calculation. Annual heat-death sensitivity can instead arise from the sampled hazard, exposure, impact-function, and adaptation dimensions. The vulnerability coordinates should therefore be interpreted as an uncertainty analysis of the vulnerability projections themselves, rather than as modifiers of individual heat-mortality risk.

Future development priorities for URBADAPT-HEAT include: (a) expanded local calibration data for adaptation parameters and costs; (b) richer dynamic population and neighbourhood-transition modelling; (c) a more comprehensive representation of adaptation options not currently covered by the framework, including building retrofitting, cool roofs, and behavioural change; (d) iterative thermo-behavioural feedback coupling for AC waste-heat interactions; (e) broader welfare accounting of non-mortality co-benefits and externalities; (f) enhanced sensitivity diagnostics, including interaction-focused and spatially resolved analyses, alongside the existing global uncertainty outputs; and (g) propagation of the Masselot coefficient-simulation and variance-covariance uncertainty through the headline LHS, extending the current point-estimate reconstruction. Beyond these model-development priorities, a companion paper applying URBADAPT-HEAT to the complete set of 40 European cities and reporting the substantive urban heat-adaptation appraisal results is currently in preparation. Extensions such as the monetised, VSL-based valuation described above are natural candidates for that applied setting.

5.2 Scaling to different climate hazards

By design, URBADAPT-HEAT is one instance of a deliberately hazard-agnostic architecture. The URBADAPT pipeline combines parts of the analysis that are common to any urban climate risk (the hazard-exposure-vulnerability characterisation) from the parts that are hazard-specific: the definition of the exposed elements at risk (people), the dose- or damage-response functions, and the set of adaptation options represented. Extending the framework to additional hazards therefore requires substituting these hazard-specific modules while reusing the surrounding infrastructure and configuration layer unchanged. First, each new hazard requires a high-resolution intra-urban hazard layer analogous to the UrbClim temperature fields used here. Candidate hazards and data sources include, for ex-

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

ample, pluvial and fluvial flooding (from urban hydrodynamic or rainfall-runoff models and event catalogues) (Fan et al., 2025). Where a hazard is naturally event-based rather than represented as a continuous field, the event/frequency structure - used here to assign each daily heat event a within-year frequency - generalises directly to return-period event sets. In addition, the heat implementation represents exposure as age-stratified population because the endpoint is mortality. Hazards acting primarily on the built environment, such as flooding, instead require exposure layers of physical and economic assets - building stock, critical infrastructure, and their replacement or content values - and a correspondingly different impact metric, with the mortality dose-response functions replaced by depth-damage or concentration-response functions. Third, the public-private, mechanism-based taxonomy that structures the heat adaptation module also needs to be transferred to other hazards. Just as heat measures act by modifying the hazard (street trees), attenuating the impact function (air conditioning), or reducing day-specific exposure (early warning systems), flood measures span hazard reduction (sustainable urban drainage, retention basins, permeable surfaces), asset-level impact reduction (property-level flood resilience), and warning-based response. Retaining the same hazard-modifying, exposure-reducing, and impact-attenuating classification keeps the cost-benefit, interaction, and equity accounting consistent across hazards and makes cross-hazard prioritisation tractable within a single framework. Finally, different hazards impose different resolution and temporal-aggregation trade-offs. Heat is well represented by daily mean fields aggregated to annual mortality on a 100 m grid, whereas pluvial flooding may demand event-based return period maps. Harmonising these onto the common reference grid and the four-target-year trajectory will in some cases require explicit up- or downscaling and careful treatment of event timing, and the appropriate compromise between spatial detail, temporal resolution, and computational cost will be hazard-specific. A longer-term objective of the URBADAPT programme is to represent compound and cascading urban climate risks (Wetterhall et al., 2025) and to appraise adaptation portfolios whose benefits and costs span more than one hazard simultaneously.

6 Conclusions

This paper introduced URBADAPT-HEAT v1.0, an open-source, city-agnostic geospatial framework for urban heat-risk assessment and public-private adaptation cost-benefit analysis. Built on the CLIMADA risk engine and configured entirely through city-specific YAML files, the framework couples high-resolution UrbClim temperature fields and scenario-consistent PROVIDE monthly temperature-change fields derived from CMIP6 models, age-stratified population exposure projected with Wittgenstein Centre demographics, a dynamic composite social vulnerability index, and city-specific age-resolved heat-mortality impact functions reconstructed from Masselot et al. (2023). Onto this risk chain it adds a mechanism-based adaptation module covering three pathways - air conditioning, urban street trees, and early warning systems - evaluated against a realistic current-adaptation baseline, together with a 25-year discounted cost-benefit, portfolio-optimisation, and equity-accounting layer that surfaces the interactions

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

and externalities (air-conditioning waste heat, vegetation-AC electricity co-benefits, and EWS-AC overlap) which single-pathway appraisals typically omit.

The four-city piloting - Rome, Athens, Lisbon, and Copenhagen - demonstrates the framework across heterogeneous climatic, demographic, and morphological settings. Under the configured central specifications, EWS is the most cost-effective standalone pathway in each pilot city and also the pathway with the greatest modelled cumulative avoided-death benefit. Rome and Lisbon use the marginal EWS interpretation, whereas Athens and Copenhagen use the counterfactual interpretation. The framework additionally supports an intermediate interpretation, although none of the four pilot-city configurations uses it. AC produces the greatest gross benefit among the two physical cooling measures but entails substantial private expenditure and a waste-heat penalty that takes back 1–3 % of that benefit; the relative mortality performance and cost-effectiveness of urban vegetation vary strongly among cities, while several vegetation co-benefits remain outside the mortality-only objective. Projected heat-attributable burdens differ by more than an order of magnitude across the pilots through the joint influence of city-specific exposure-response relationships, temperature distributions, population size, and age structure, with existing AC represented separately in the current-adaptation reference branch. The income-targeted AC and equity-weighted vegetation variants allow alternative spatial allocations and their distributional consequences to be compared, without assuming that equity-oriented targeting necessarily maximises aggregate avoided mortality.

URBADAPT-HEAT is released as a reproducible and extensible decision-support tool rather than a finished model. Its principal limitations - synthetic and anchored elements of the hazard and exposure chain, adaptation calibration coefficients transferred across contexts, spatially uniform within-city impact functions, and uncertainty propagation that does not yet incorporate the coefficient-simulation and variance-covariance information distributed with the Masselot model - are documented explicitly to guide interpretation and prioritise future development. As the heat-specific implementation of the broader, hazard-agnostic URBADAPT architecture, the framework is designed to be extended to additional urban climate hazards and, ultimately, to the appraisal of adaptation portfolios spanning multiple compounding hazards. By integrating hazard, exposure, vulnerability, impact, and public-private adaptation in a single transparent workflow, URBADAPT-HEAT aims to move urban heat adaptation appraisal from fragmented, single-measure studies toward consistent, comparable, and equity-aware decision support.

Code and data availability. The URBADAPT-HEAT source code is available at https://github.com/URBADAPT/URBADAPT-HEAT under the Creative Commons CC0 1.0 Universal public-domain dedication. The exact version used to produce the results in this paper, v1.0.1, is archived at https://doi.org/10.5281/zenodo.24165564 (Aboudrar-Méda and Falchetta, 2026a). The archive contains the model code, the city configurations, the cluster job scripts with the settings of each run, the pinned software environment, the scripts that generate every figure and table, and the code behind two fitted model inputs: the LST-T2M climate emulator, as run, and the street-greenery cooling coefficients, estimated with the code of Falchetta et al. (2026b). Its file `urban-heat/REPRODUCIBILITY.md` gives the code commits and settings of each run and documents a rerun of the deterministic chain from the raw inputs that reproduced the archived outputs exactly. The model output for the four

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

pilot cities, including the uncertainty ensembles, is archived at https://doi.org/10.5281/zenodo.24163360 (Aboudrar-Méda and Falchetta, 2026c), and their input data at https://doi.org/10.5281/zenodo.24163362 (Aboudrar-Méda and Falchetta, 2026b). Table B6 lists every input dataset with its provider, version, identifier and licence. Datasets that have their own persistent identifiers (the global Local Climate Zone map of Demuzere et al. (2022), GHS-AGE and WorldPop), the UrbClim daily fields and the PROVIDE warming deltas, whose distribution terms state no licence, and the OpenStreetMap boundaries (© OpenStreetMap contributors, ODbL 1.0), which are retrieved at run time, are not redistributed; the city manifests in the code archive record how each is obtained. Documentation is available at https://github.com/URBADAPT/URBADAPT-HEAT/wiki, and the model website gives additional details on the notebook structure and contents: https://urbadapt.github.io/.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Appendix A: Supplementary Figures

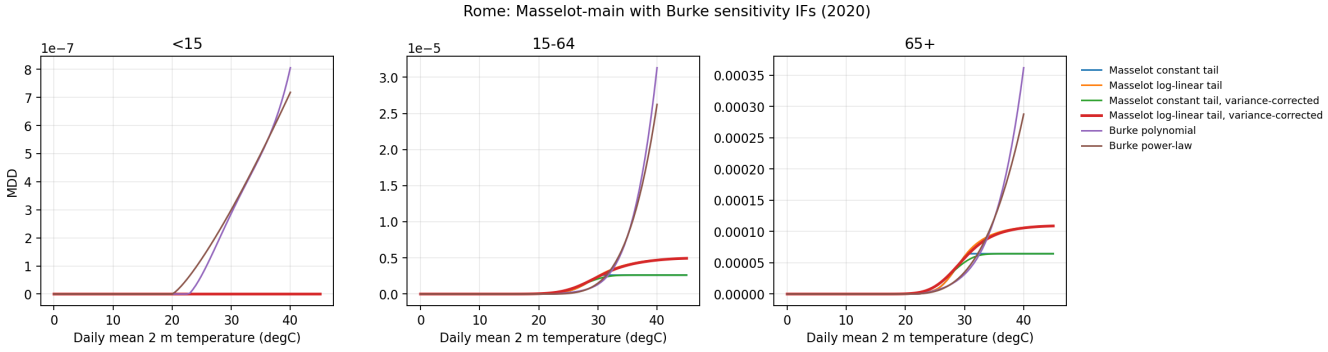


Figure A1. Four-family heat-mortality impact function comparison for Rome at the 2020 reference year, by age group. Each panel shows the mean damage degree (MDD) as a function of daily mean 2 m temperature for the four impact-function specifications carried through the framework: the canonical Masselot log-linear-tail variant (heavy line; main IF), the Masselot constant-tail variant (in-headline extrapolation sensitivity), the Burke polynomial form, and the Burke power-law form (the two Burke variants are sampled as IF-conditional uncertainty bands in the uncertainty layer, Sect. 3.11, and do not enter the deterministic canonical run). The < 15 panel shows the Masselot MDD identically at zero by construction (Masselot excludes ages below 20), while the two Burke variants retain a small positive response in this age band reflecting the structural difference between literature anchors. The 65+ panel illustrates the principal divergence between IF sources: Masselot's B-spline curves are city-specific and centred on the per-city minimum-mortality temperature (here $\sim 20.5\,^{\circ}$C for Rome 65+), whereas the Burke functional forms share a fixed reference temperature of $20\,^{\circ}$C across all cities and grow unbounded with $T - T_{\mathrm{ref}}$. Above each city's historical maximum temperature, the Masselot constant-tail and log-linear-tail variants bracket the structural uncertainty of out-of-sample extrapolation. The figure is generated as part of the NB04 masselot-main workflow (file `rome_IF_masselot_main_burke_sensitivity_comparison.png`); equivalent panels for Athens, Lisbon, and Copenhagen are produced by the corresponding city pipelines.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

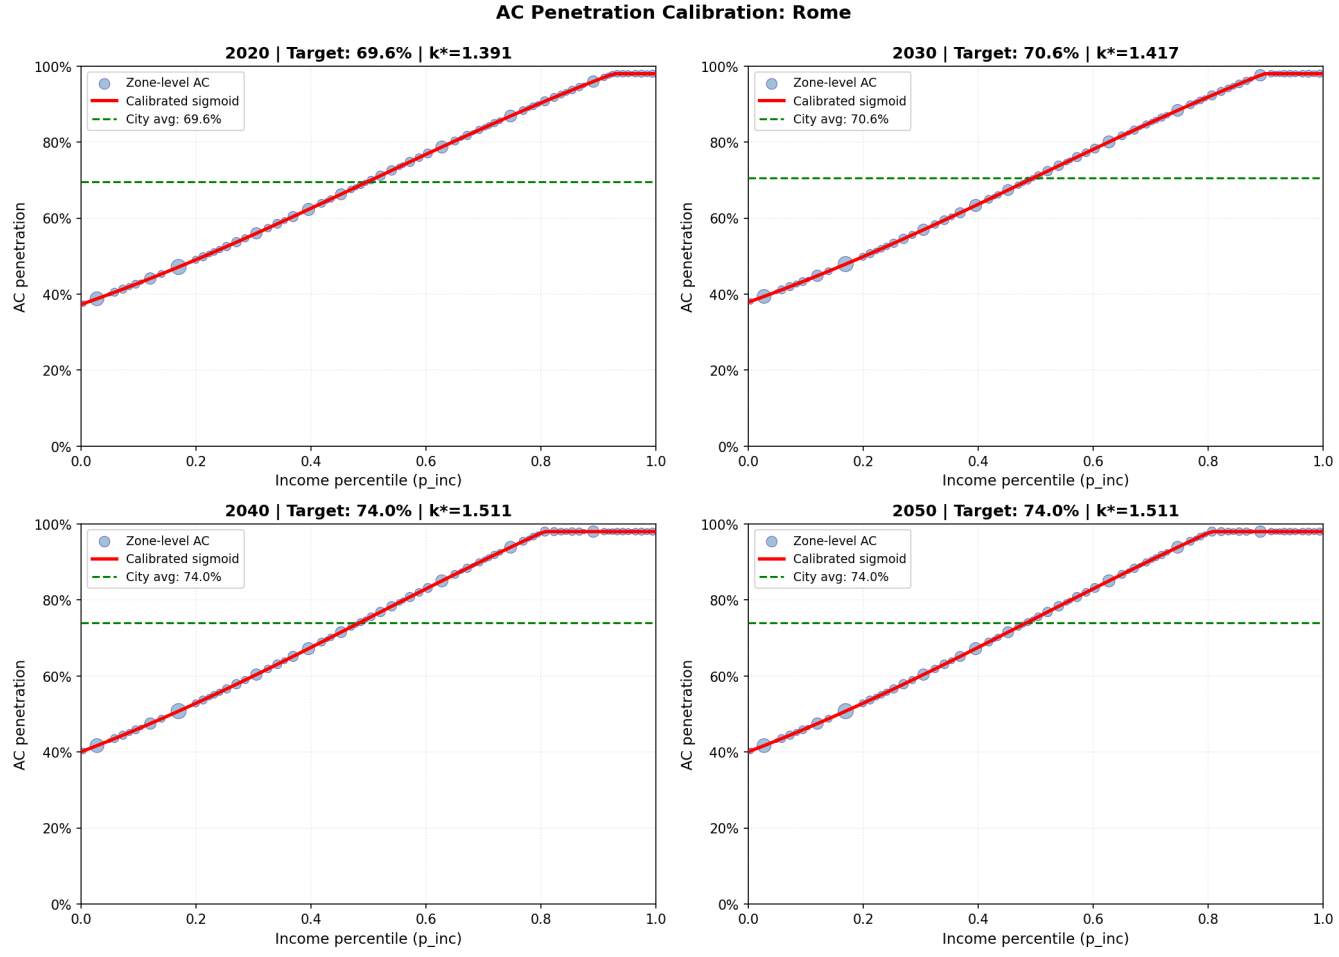


Figure A2. Logistic income-rank downscaling of AC penetration in Rome at the four target years. Each panel plots the model-assigned zone-level AC coverage against within-city income percentile. The calibrated logistic allocation curve is constrained so that its population-weighted mean exactly reproduces the corresponding NUTS3 penetration target shown by the green dashed line. The year-specific calibration factor $k^{\star}$ rescales the fixed logistic profile; it is not an independently estimated shape parameter or evidence of a changing causal income gradient. The construction assigns lower baseline coverage to lower-income zones, providing the reference against which the income-targeted catch-up policy in Sect. 3.9.1 is evaluated.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

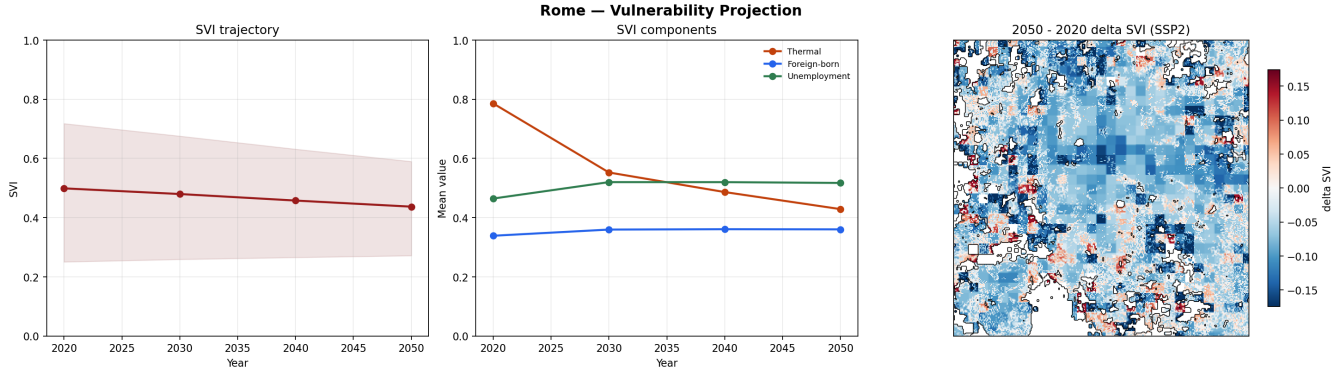


Figure A3. Dynamic projection of the Social Vulnerability Index for Rome over 2020-2050 under SSP2. *Left*: pixel-level SVI distributions at the four target years, showing a shift towards lower composite SVI. *Centre*: city-mean SVI and component trajectories; the shaded areas represent the spatial P10-P90 spread across grid cells, not uncertainty across SSPs. In this configuration, the decline in composite SVI is driven principally by the thermal component, while the foreign-born and non-employment component means are comparatively stable or increase slightly. *Right*: spatially explicit change in SVI between 2020 and 2050; blue indicates declining modelled vulnerability and red indicates increasing modelled vulnerability.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

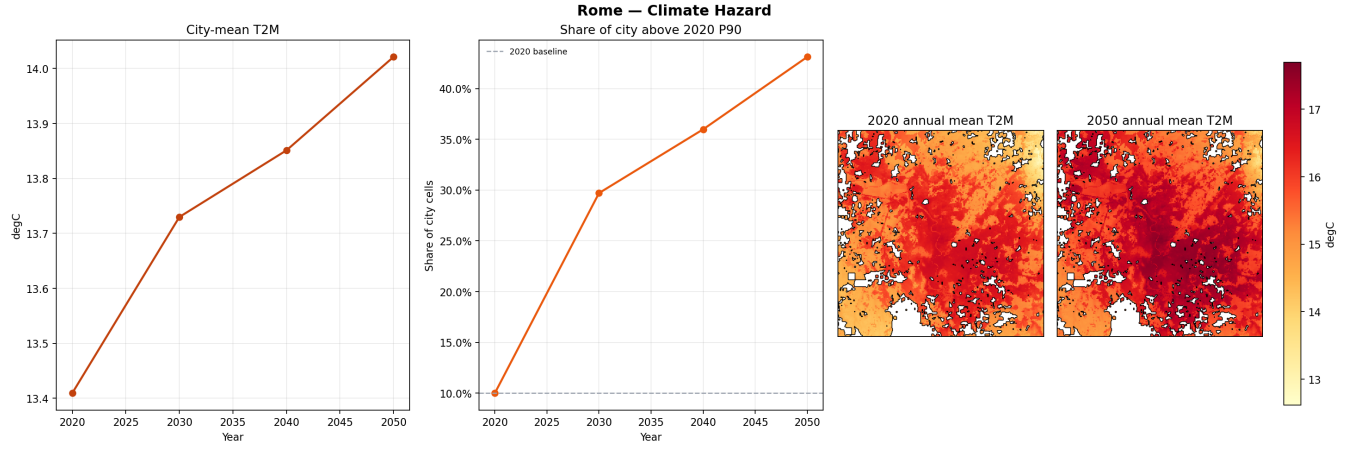


Figure A4. Climate-hazard characterisation for Rome under the central CurPol scenario. *Left*: city-mask mean of the annual-mean T2M field at the four target years, increasing by approximately 0.71 °C between 2020 and 2050. The corresponding changes across the four pilot-city annual-mean maps range from approximately 0.65 to 0.79 °C. *Centre*: share of FUA grid cells whose annual-mean T2M equals or exceeds the spatial 90th percentile of the 2020 annual-mean T2M map, increasing from approximately 10 % in 2020 to 43 % in 2050. This is a spatial annual-mean-map diagnostic, not a daily-temperature exceedance metric. *Right*: annual-mean T2M fields for 2020 and 2050 on the common 100 m grid.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Appendix B: Supplementary Tables

Table B1: Positioning of URBADAPT-HEAT v1.0 against the methodological frontier identified by Liu et al. Liu et al. (2026). The six domains correspond to the review's Section 3 framework; the five challenges to their Section 4 outlook. Ratings reflect how the current implementation addresses each identified gap relative to the 110-study corpus.

Domain	Frontier gap	Rating	URBADAPT-HEAT implementation and residual gaps
Section 3 — Methodological domains			
3.1 Climate scenarios	Sub-km urban data; across-GCM spread	●●●●○	UrbClim 100 m resolves intra-urban UHI; PROVIDE CMIP6 low/central/high GCM bands. *Gap:* delta method preserves baseline variability unchanged; no dynamic UHI evolution.
3.2 Exposure–response	Age-stratified ERFs; spatial heterogeneity; acclimatisation	●●●●○	City- and age-specific quadratic B-spline DLNM curves (Masselot et al., 2023), centred on per-city minimum-mortality temperatures rather than a common reference temperature; polynomial and power-law forms (Burke et al., 2025) are retained as conditional structural sensitivities, alongside sensitivities on level, tail extrapolation and mortality displacement. *Gap:* ERF spatially uniform within each city; no acclimatisation.
3.3 Demography & socioeconomics	SSP-consistent age projections; dynamic vulnerability	●●●●●	WorldPop + WCDE age-structured projections; composite SVI projected via DRMKC + GDL with city-specific k, φ. *Gap:* uniform national scalars; parameters set, not calibrated.
3.4 Urban characteristics	LCZ; green infrastructure; building stock	●●●○○	LCZ conditions cooling coefficients; GVI as adaptation lever; GHS-AGE for thermal vulnerability. *Gap:* static urban form; no LULC change modelling.
3.5 Adaptation & CBA	Multi-pathway with interactions; cost-effectiveness	●●●●●	AC + EWS + trees with waste-heat feedback, AC–EWS overlap, veg–electricity co-benefit; NPV, Pareto, equity analysis. *Gap:* stylised mechanisms; no behavioural model.
3.6 Uncertainty	Full-chain UQ; global sensitivity	●●●●●	512 LHS samples over 57 declared input coordinates, some conditional on other sampled choices and inactive for particular branches, with PAWN global sensitivity. *Gap:* limited structural uncertainty; no deep uncertainty.
Section 4 — Challenges and outlook			
4.1 High-res. climate	Ensemble urban climate projections	●●●●○	UrbClim + PROVIDE ensemble + delta bands. See 3.1.
4.2 Consistent projections	Coupled population–vulnerability scenarios	●●●●●	Dynamic SVI + SSP population. See 3.3.

Continued on next page

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B1: Positioning of URBADAPT-HEAT v1.0 against the methodological frontier identified by Liu et al. Liu et al. (2026) (continued).

Domain	Frontier gap	Rating	URBADAPT-HEAT implementation and residual gaps
4.3 Intra-urban vulnerability	Spatially explicit social vulnerability	●●●●○	Zone-level SVI with income, foreign-born, unemployment, building age. *Gap:* SVI modulates exposure, not dose–response.
4.4 Urban dynamics	Expansion, land-use change, building evolution	●○○○○	Not modelled. Static LCZ and footprint. Defensible for European cities; critical gap for Global South.
4.5 Holistic adaptation	Behavioural, institutional, cross-sectoral	●●○○○	Three physical pathways with CBA but no behavioural compliance, morbidity, or cross-sectoral synergies.

Rating scale: ●●●●● = frontier (among first implementations); ●●●●○ = strong; ●●●○○ = above average; ●●○○○ = partial; ●○○○○ = gap.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B2: Early-warning-system (EWS) maturity taxonomy and its treatment in URBADAPT-HEAT. Countries are grouped by the maturity of their national heat–health warning arrangements; the *URBADAPT-HEAT treatment* row beneath each class states the corresponding modelling assumptions and limitations.

EWS class	**Definition**	**Countries and official sources**
Mature HHWS, epidemiologically based, with surveillance (real-time mortality data)	Operational heat–health warning & action system with an organised health-system response and surveillance, where alert thresholds are derived from local temperature–mortality (epidemiological) relationships.	Italy: Italian Ministry of Health (2026) France: Santé publique France (2026) Spain: Spanish Ministry of Health (2026) Portugal: Instituto Nacional de Saúde Doutor Ricardo Jorge (2026)
URBADAPT-HEAT treatment: URBADAPT-HEAT treats this class as an already-operational mature system (`mature_hhwws`, `marginal` interpretation). The adaptation is the *added* protection on an existing system—a marginal improvement or maintained protection, not deployment from zero—so setup CAPEX is set to zero and costs enter only through per-warning-day and fixed annual operating costs. Efficacy is applied as a marginal reduction in heat-attributable mortality on warning days, age-differentiated and calibrated to on-system estimates deDonato2018,Pavanello2025Dec. **Limitations:** the model does not explicitly simulate local epidemiological thresholds, surveillance protocols, communication quality, public compliance or health-system capacity; these are treated as homogeneous across cities.		
HHWS, meteorologically based or with historical epidemiological grounding (no real-time mortality data)	National operational heat plan with an organised response, but alert thresholds are meteorological (they can be epidemiologically calibrated but are not based on real-time epidemiological data).	Germany: Deutscher Wetterdienst (2026) Netherlands: Rijksinstituut voor Volksgezondheid en Milieu (2026) Belgium: Belgian Federal Public Service Health, Food Chain Safety and Environment (2026) Austria: Austrian Federal Ministry of Social Affairs, Health, Care and Consumer Protection (2026) Hungary: Nemzeti Népegészségügyi Központ (2026) Romania: Administraţia Naţională de Meteorologie (2026) Sweden: Sveriges meteorologiska och hydrologiska institut (2026) Finland: Ilmatieteen laitos (2026)
URBADAPT-HEAT treatment: URBADAPT-HEAT treats this class as an intermediate-maturity system (`meteo_hhwws`, `intermediate` interpretation): a national heat-warning or heat-action framework exists, but triggers are mainly meteorological, biometeorological or historically calibrated rather than based on real-time mortality surveillance, so the adaptation is an *upgrade* toward a more integrated heat–health system rather than either a marginal improvement or deployment from zero. Efficacy is the age-differentiated midpoint of the marginal and counterfactual efficacies, and setup CAPEX is set midway between the mature-system value (zero) and full deployment. **Limitations:** subnational implementation, health-sector coordination, threshold calibration and follow-up actions differ strongly across these countries and are only approximately represented.		

continued on next page

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B2 continued

EWS class	Definition	Countries and official sources
Meteorological warnings only (no HHWS)	National meteorological service issues heat warnings, but there is no comprehensive national heat–health action plan or coordinated health-system response.	Greece: Hellenic National Meteorological Service (2026) Croatia: Državni hidrometeorološki zavod (2026) Slovenia: Slovenian Environment Agency (2026) Bulgaria: Burgas Media (2026) Slovakia: Slovenský hydrometeorologický ústav (2026) Czech Republic: Český hydrometeorologický ústav (2026) Poland: Instytut Meteorologii i Gospodarki Wodnej (2026) Denmark: Danmarks Meteorologiske Institut (2026) Ireland: Met Éireann (2026) Estonia: Estonian Environment Agency (2026) Latvia: Latvijas Vides, ģeoloģijas un meteoroloģijas centrs (2026) Lithuania: Lietuvos hidrometeorologijos tarnyba (2026)
URBADAPT-HEAT treatment: URBADAPT-HEAT treats this class as heat alerts without a comprehensive heat–health action plan or coordinated health-system response, so the intervention is the counterfactual deployment of a full HHWS/HHAP from a low baseline (`weak_threshold`, `counterfactual`). The full setup CAPEX and wider efficacy parameters are applied, with the counterfactual efficacy scaled by region (urban2025: 25.2 % overall, Northern Europe 33.2%, Southern Europe 11.9%). Because that evidence shows no statistically significant efficacy difference by development class ($p = 0.19$) or region ($p = 0.44$), the graded efficacy across tiers is treated as a modelling assumption expressed through wide uncertainty ranges rather than a measured gradient. **Limitations (two-directional):** pre-existing traffic-light alerts may already induce some protective behaviour, so ignoring this baseline could overstate incremental benefits; conversely, the absence of an organised health-system response may reduce realised efficacy relative to the mature-system evidence.		

Efficacy figures from Urban et al. (2025); cost, ramp-up and setup-CAPEX parameters from de'Donato et al. (2018), Pavanello and Valenti (2025), Ebi et al. (2004) and World Health Organization Regional Office for Europe (2021); country classification cross-checked against Casanueva et al. (2019). Official-source URLs last accessed July 2026.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.



Parameter	Rome	Athens	Lisbon	Copenhagen
DRMKC NUTS3 region	ITI43	EL303	PT1A0	DK011
k_{oth}	0.85	0.90	0.70	0.75
k_{ne}	0.85	0.90	0.75	0.65
$\phi_{\mathrm{oth},2030}/\phi_{\mathrm{oth},2050}$	0.92/0.75	0.96/0.85	0.88/0.65	0.86/0.60
$\phi_{\mathrm{ne},2030}/\phi_{\mathrm{ne},2050}$	0.90/0.70	0.94/0.80	0.86/0.60	0.85/0.58
Effective thermal-decay rate r_c (yr^{-1})	0.0138	0.0093	0.0139	0.0065
New-build vulnerability z_{new}	0.15	0.15	0.15	0.15
Growth sensitivity	0.70	0.60	0.80	0.75
Maximum new-build share	0.30	0.25	0.35	0.30

Table B3. Dynamic-vulnerability parameters in the current four pilot-city YAML configurations. The k, ϕ, growth-sensitivity, and maximum-share parameters are configurable response assumptions rather than directly observed city characteristics. Here "oth" denotes the Eurostat place-of-birth proxy and "ne" the non-employment proxy.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B4: URBADAPT-HEAT uncertainty quantification matrix: principal continuous and discrete parameters. Baseline values are the reference configuration and Low/High delimit the sampled range; where a multiplier is shown, the coordinate is sampled relative to the calibrated city-specific value. Entries reading *city-specific differ among the pilot cities and are listed per city in the corresponding footnote. This table reports the principal sampled parameters rather than the complete declaration: the canonical daily-mean track declares 57 input coordinates, some of which are conditional on other sampled choices and inactive for particular branches, as described in Sect. 3.11. Generated from the city configuration files by `workflow_tables.make_uncertainty_matrix`; do not edit by hand.

Parameter	Description	Baseline	Low	High	Distribution
Hazard: additional parameters					
Baseline T2M mode[4]	UrbClim 2008–2017 baseline construction	climatology_mean	domain_peak_day	warmest_summer	Discrete (4 modes)
Waste-heat sensitivity (S)[5]	City-wide daily-mean warming per $\mathrm{W\,m^{-2}}$ of AC heat rejected over the built-up area (LCZ 1–10)	0.02	0.01	0.035	Uniform(0.01, 0.035)
AC electricity at full activity (e_{full})[6]	Electricity per AC household on a day at full cooling activity; also sets each draw's AC electricity cost	15	10	20	Uniform(10, 20)
COP degradation sensitivity[7]	Fractional COP drop per °C of outdoor T rise	0.065	0.04	0.09	Discrete {low, central, high}
Dynamic vulnerability (projected SVI)					
Persistence (k)[8]	Spatial persistence of inequality over time; higher k means slower convergence of vulnerability gaps	city-specific	0.55	0.95	Uniform(0.55, 0.95)

continued on next page

[4] NB02; climate.t2m_baseline_mode_options; also pixelwise_doy_max (extreme sensitivity)

[5] $\mathrm{K\,(W\,m^{-2})^{-1}}$; ac.heat_energy.waste_heat.sensitivity_k_per_wm2, default in `cityheat/ac_heat_energy.py`; modelling median 0.017 (IQR 0.008–0.029) (Li et al., 2026); upper bound Phoenix under extreme heat (Salamanca et al., 2014)

[6] kWh per household per day; 2.5 on the first cooling day, activity ramp 18–25 °C; ac.heat_energy.electricity.kwh_per_household_day_full, default in `cityheat/ac_heat_energy.py`

[7] All configs ac.waste_heat.cop_degradation.sensitivity_per_C; Kou et al. (2026) Fig. 2a: 4–9 % per °C

[8] vulnerability.dynamic.k; controls how quickly spatial vulnerability patterns converge toward the city mean. Rome: default 0.85, foreign-born 0.85, unemployment 0.85; Athens: default 0.9, foreign-born 0.9, unemployment 0.9; Lisbon: default 0.72, foreign-born 0.7, unemployment 0.75; Copenhagen: default 0.7, foreign-born 0.75, unemployment 0.65

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B4 (continued)

Parameter	Description	Baseline	Low	High	Distribution
Spatial retention (φ_{2050})[9]	Fraction of original spatial variance retained by 2050; lower means stronger spatial smoothing	city-specific	0.5	0.9	Uniform(0.50, 0.90)
DRMKC scale (foreign-born)[10]	Short-run (≤2030) DRMKC trend multiplier for the foreign-born component	0.04	0.02	0.08	Uniform(0.02, 0.08)
DRMKC scale (unemployment)[11]	Short-run (≤2030) DRMKC trend multiplier for the unemployment component	0.08	0.04	0.16	Uniform(0.04, 0.16)
GDL–GVI scale (foreign-born)[12]	Long-run (>2030) SSP-driven multiplier for the foreign-born component	0.35	0.15	0.55	Uniform(0.15, 0.55)
GDL–GVI scale (unemployment)[13]	Long-run (>2030) SSP-driven multiplier for the unemployment component	0.5	0.25	0.75	Uniform(0.25, 0.75)
Retrofit rate[14]	Annual rate of thermal-efficiency improvement of the building stock	city-specific	0.005	0.02	Uniform(0.005, 0.020)
Impact functions: continuous parameters					
MDD scale (<15)[15]	MDD multiplier for age <15	1	0.8	1.2	Uniform(0.80, 1.20)
MDD scale (15–64)[16]	MDD multiplier for age 15–64	1	0.8	1.2	Uniform(0.80, 1.20)

continued on next page

[9] vulnerability.dynamic.phi.default_2050; linearly interpolated 2030→2050. Rome: 0.75; Athens: 0.85; Lisbon: 0.65; Copenhagen: 0.6
[10] vulnerability.dynamic.foreign_born_projection.drmkc_scale; JRC DRMKC social-indicator trends
[11] vulnerability.dynamic.unemployment_projection.drmkc_scale; JRC DRMKC social-indicator trends
[12] vulnerability.dynamic.foreign_born_projection.gvi_scale; GDL–GVI SSP scenario trajectories (the national socioeconomic vulnerability index, not the Green View Index)
[13] vulnerability.dynamic.unemployment_projection.gvi_scale; GDL–GVI SSP scenario trajectories (the national socioeconomic vulnerability index, not the Green View Index)
[14] vulnerability.dynamic.thermal_projection.retrofit_rate_per_year; reduces the thermal vulnerability component annually
[15] NB09 sampling range, not a city-config value
[16] NB09 sampling range, not a city-config value

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B4 (continued)

Parameter	Description	Baseline	Low	High	Distribution
MDD scale (65+)[17]	MDD multiplier for age 65+	1	0.8	1.2	Uniform(0.80, 1.20)
Mortality displacement[18]	Fraction of deaths shifted in time	0	0	0.3	Uniform(0.00, 0.30)
% pop. affected (PAA)[19]	Fraction of the population exposed to the heat hazard	1	0.8	1	Uniform(0.80, 1.00)
Early warning system (EWS) parameters					
EWS interpretation[20]	Marginal (enhance an existing system) versus counterfactual (deploy from a low baseline)	city-specific	–	–	Discrete: city-dependent
EWS marginal eff. (<15)[21]	Marginal efficacy for the <15 age group	city-specific	–	–	Discrete {low, central, high}
EWS marginal eff. (15–64)[22]	Marginal efficacy for the 15–64 age group	city-specific	–	–	Discrete {low, central, high}
EWS marginal eff. (65+)[23]	Marginal efficacy for the 65+ age group	city-specific	–	–	Discrete {low, central, high}
EWS counterfactual eff.[24]	Full-system efficacy (flat, no age differentiation)	city-specific	–	–	Discrete {low, central, high}
AC–EWS overlap factor[25]	Fraction of AC-prevented deaths that reduces the EWS addressable pool	city-specific	–	–	Discrete {low, central, high}

continued on next page

[17] NB09 sampling range, not a city-config value
[18] Net mortality reduction = $(1 - \text{disp}) \times$ MDD; NB09 sampling range
[19] CLIMADA fraction of affected exposure, not a population-attributable fraction; NB09 sampling range
[20] Determines which efficacy block applies. Marginal in Rome and Lisbon; counterfactual in Athens and Copenhagen
[21] ews.efficacy_marginal, low/central/high. Rome: 0.02/0.05/0.08; Athens: 0.02/0.05/0.08; Lisbon: 0.01/0.03/0.06; Copenhagen: 0.01/0.03/0.05
[22] ews.efficacy_marginal, low/central/high. Rome: 0.03/0.07/0.12; Athens: 0.03/0.07/0.12; Lisbon: 0.02/0.05/0.1; Copenhagen: 0.02/0.05/0.09
[23] ews.efficacy_marginal, low/central/high. Rome: 0.08/0.15/0.25; Athens: 0.05/0.1/0.18; Lisbon: 0.05/0.1/0.18; Copenhagen: 0.04/0.09/0.16
[24] ews.efficacy_counterfactual, low/central/high. Rome: 0.2/0.35/0.5; Athens: 0.1/0.2/0.35; Lisbon: 0.1/0.2/0.35; Copenhagen: 0.08/0.15/0.28
[25] Reflects city-specific AC penetration. Rome: 0.3/0.5/0.7; Athens: 0.15/0.3/0.5; Lisbon: 0.1/0.2/0.4; Copenhagen: 0.1/0.2/0.4

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B4 (continued)

Parameter	Description	Baseline	Low	High	Distribution
EWS displacement (<15)[26]	Harvesting fraction for <15	0.1	0.05	0.2	Discrete {low, central, high}
EWS displacement (15–64)[27]	Harvesting fraction for 15–64	0.1	0.05	0.2	Discrete {low, central, high}
EWS displacement (65+)[28]	Harvesting fraction for 65+	0.25	0.1	0.4	Discrete {low, central, high}
EWS ramp years[29]	Years for the EWS to reach full efficacy	3	–	–	Sampled in NB09
EWS setup CAPEX[30]	One-off deployment expenditure by baseline maturity class	city-specific	–	–	Discrete: by maturity class
EWS cost model[31]	Operating-cost formulation for the EWS	pavanello	–	–	Discrete {pavanello, chiabai}
Tree / vegetation cooling parameters					
Cooling coeff. scale[32]	Multiplier on the GVI→dT regression coefficients	1	0.5	1.5	Uniform(0.50, 1.50)
Cap uplift (GVI)[33]	Maximum GVI index-point uplift from tree planting	0.12	0.06	0.2	Uniform(0.06, 0.20)
Tree ramp years[34]	Years for a tree to reach full cooling maturity	12	8	15	Discrete
Tree start age[35]	Age of the tree at planting	5 yr	0 yr	5 yr	Discrete (2 levels)
Tree lifetime[36]	Asset life over which tree flows are evaluated	25 yr	–	–	Fixed

continued on next page

[26] ews.displacement; identical across the pilot cities
[27] ews.displacement; identical across the pilot cities
[28] ews.displacement; identical across the pilot cities
[29] ews.ramp_years. Identical across the pilot cities: the implementation applies a single low-efficacy period and does not yet differentiate ramp duration by maturity class
[30] ews.capex_setup (€). Rome: €0; Athens: €150,000; Lisbon: €0; Copenhagen: €150,000
[31] ews.cost_model; Pavanello and Valenti (2025) $0.014/capita/day, Chiabai Chiabai et al. (2018) €6 200–14 000/day
[32] NB07; ±50 % brackets regression uncertainty (standard errors available per LCZ and month)
[33] trees.cap_uplift_0_1; depends on policy ambition and available space
[34] trees.ramp_years; reflects Mediterranean tree growth
[35] NB08 sensitivity: age 0 (seedling) versus age 5 (pre-grown)
[36] trees.lifetime_years

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table B4 (continued)

Parameter	Description	Baseline	Low	High	Distribution
CBA / economic parameters (post-impact)					
Discount rate[37]	Real social discount rate	0.03	0.01	0.07	Reported at 1, 3, 5, 7 %
AC CAPEX/user[38]	AC installation cost per additional user	city-specific	×0.8	×1.2	Discrete (3 multipliers)
AC tariff[39]	Electricity price for AC operation	city-specific	–	–	Sampled in NB09
AC lifetime[40]	Replacement cycle for an AC unit	12 yr	9 yr	16 yr	Discrete (3 levels)
Tree CAPEX/tree[41]	Capital cost per tree planted	city-specific	×0.8	×1.2	Discrete (3 multipliers)
Tree O&M/tree/yr[42]	Annual maintenance cost per tree	city-specific	×1	×5	Discrete (base or 5×)
Tree CAPEX/GVI point[43]	Capital cost per unit of summed regional GVI uplift	€10,000,000	–	–	Fixed

[37] cba.discount_rate; the CBA reports discount-rate sensitivity at 1, 3, 5 and 7 %

[38] ac.capex_per_user (€). Rome: €1,498; Athens: €1,349; Lisbon: €1,336; Copenhagen: €2,156. Sampled as a multiplier on the calibrated city value

[39] ac.tariff_eur_per_kwh (€ kWh^{-1}). Rome: 0.17; Athens: 0.14; Lisbon: 0.17; Copenhagen: 0.3

[40] ac.lifetime_years; identical across the pilot cities, sampled over 9–16 yr

[41] trees.capex_per_tree_eur (€). Rome: €680; Athens: €612; Lisbon: €606; Copenhagen: €978. Sampled as a multiplier on the calibrated city value

[42] trees.om_per_tree_per_year_eur (€). Rome: €63; Athens: €57; Lisbon: €56; Copenhagen: €91. NB08 tests a 5× O&M scenario

[43] trees.capex_per_index_pt_eur (€)

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

	Median		P5		P95	
City	N=128	N=64	N=128	N=64	N=128	N=64
Athens	613	609	399	336	1,004	1,307
Rome	488	484	303	262	836	1,136
Lisbon	169	163	110	101	285	324
Copenhagen	30.2	30.8	15.7	14.7	70.3	72.6

Table B5. Sample-size comparison of the pooled annual heat-attributable mortality distributions for the four pilot cities: the headline $N = 128$ Masselot Latin-hypercube ensemble and a separately generated $N = 64$ ensemble using the same engine, input definitions, and random seed. Relative to $N = 128$, the $N = 64$ median differs by $-0.6\,\%$ in Athens, $-0.7\,\%$ in Rome, $-3.5\,\%$ in Lisbon, and $+2.0\,\%$ in Copenhagen. P5 estimates differ by at most 15.7 %, while $N = 64$ P95 estimates are 3.2–35.9 % higher across the four cities. These differences describe sensitivity to the two sampling designs; they neither identify the causes of individual upper-tail draws nor establish convergence. The larger $N = 128$ ensemble is used for the reported uncertainty distributions.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.



https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Dataset	Provider, version	Identifier or source	Licence	In input record
UrbClim daily 2 m air temperature, 2008–2017, 100 m	VITO, PROVIDE project	https://provide.marvin.vito.be/ftp/compressed_daily/	none stated	no (listed)
PROVIDE monthly warming deltas by climate model	VITO, PROVIDE project	https://provide.marvin.vito.be/ftp/future_deltas/	none stated	no (listed)
GHS functional urban areas	JRC GHSL, R2019A	https://doi.org/10.2905/347F0337-F2DA-4592-87B3-E25975EC2C95	CC BY 4.0	yes
Global map of Local Climate Zones	Demuzere et al. (2022), v3.0.0	https://doi.org/10.5281/zenodo.8419340	CC BY 4.0	no (own DOI)
WorldPop age structures, 100 m	WorldPop, Global2 R2025A	https://doi.org/10.5258/SOTON/WP00841	CC BY 4.0	no (own DOI)
National population projections by age	Wittgenstein Centre, WCDE v3	https://dataexplorer.wittgensteincentre.org/wcde-v3/	none stated	no (listed)
GHS-AGE, dominant age of the built stock	JRC GHSL, R2025A	https://doi.org/10.2905/JRC.YCNTMNG	CC BY 4.0	no (own DOI)
Census 2021 population grid, 1 km	Eurostat GISCO, release V2.2	https://gisco-services.ec.europa.eu/census/2021/	CC BY 4.0	yes
Local vulnerability index, NUTS3	JRC DRMKC Risk Data Hub	https://doi.org/10.2905/JRC.6TVNR28	CC BY 4.0	yes
Exposure-response coefficients	Masselot et al. (2023), v2.0.0	https://doi.org/10.5281/zenodo.10288665	CC BY 4.0	in code archive
AC penetration and electricity by NUTS region	Authors, from Falchetta et al. (2024)	https://doi.org/10.5281/zenodo.18456618	CC BY 4.0	yes
Street-level Green View Index points	Authors, Falchetta and Hammad (2025)	https://doi.org/10.5281/zenodo.15624204	CC BY 4.0	yes
GVI-LST cooling coefficients by city, LCZ and month	Authors, Falchetta et al. (2026b)	https://doi.org/10.5281/zenodo.17803495	CC BY 4.0	yes
LST-T2M emulator coefficients	Authors (this paper)	fit scripts in the code archive	CC BY 4.0	yes
Rome: income by postcode, tax year 2023	MEF, Dipartimento delle Finanze	https://www1.finanze.gov.it/finanze/analisi_stat/	CC BY 3.0 IT	yes
Rome: postcode (CAP) zones, July 2021	Poste Italiane, CAP Zone	https://business.poste.it/	licence held by the authors	yes
Rome Municipi, Lisbon freguesias	OpenStreetMap contributors	https://www.openstreetmap.org/copyright	ODbL 1.0	no (run time)
Athens: income by postcode, 2013, 100 m grid	Pantazis and Psycharis (2016); Wilkens (2025)	https://doi.org/10.1080/13673882.2016.11760792	permission held by the authors	yes
Athens: districts, neighbourhoods and communities	Municipality of Athens, 2023	http://gis.cityofathens.gr/	CC BY 3.0	yes
Lisbon: declared income by freguesia, 2023	Statistics Portugal (INE)	https://www.ine.pt/xurl/pub/739291089	CC BY 4.0	yes
Copenhagen: income per person by district, 2023	City of Copenhagen, table KKIND3	https://kk.statistikbank.dk/KKIND3	CC BY 4.0	yes
Copenhagen: the ten districts (bydele)	Københavns Kommune	https://www.opendata.dk/city-of-copenhagen/bydele	CC BY	yes

Table B6. Input datasets of the four pilot cities, with provider, version, persistent identifier or source and licence. The last column states whether the files are included in the input record of this paper (https://doi.org/10.5281/zenodo.24163362); datasets that have their own persistent identifier, whose terms state no licence, or that are retrieved at run time are listed here and obtained from their providers. The file `urban-heat/DATA_SOURCES.md` of the code archive gives file paths, access dates and licence references for every entry.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Author contributions. Conceptualization: G.F. and A.A.M.; Framework development: A.A.M. core; G.F. supervision; visual items: equal contribution; funding acquisition: G.F.; manuscript writing and revision: G.F. and A.A.M.

Competing interests. The authors declare no competing interests.

Acknowledgements. The authors are grateful to Massimo Tavoni, Enrica De Cian, Niels Souverijns, Martyn Clark, Ana Oliveira, the participants in EGU 2026 session CL2.3 "Urban Climate Science and Services: Observations, Modelling, Tools, and Action for Resilient Cities" and the GEO Heat Resilience Workshop, April 2026, Geneva. The authors used generative AI tools for language editing and to improve code readability and structure; all scientific content, analysis, and interpretation are the authors' own, and the authors take full responsibility for the published material.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

References

Aboudrar-Méda, A. and Falchetta, G.: URBADAPT-HEAT v1.0.1: source code, configurations and scripts, Zenodo, https://doi.org/10.5281/zenodo.24165564, archived release v1.0.1 of the URBADAPT-HEAT model code, 2026a.

Aboudrar-Méda, A. and Falchetta, G.: URBADAPT-HEAT v1.0.1: input data for four pilot cities (Rome, Athens, Lisbon, Copenhagen), Zenodo, https://doi.org/10.5281/zenodo.24163362, 2026b.

Aboudrar-Méda, A. and Falchetta, G.: URBADAPT-HEAT v1.0.1: output data for four pilot cities (Rome, Athens, Lisbon, Copenhagen), Zenodo, https://doi.org/10.5281/zenodo.24163360, 2026c.

Adams, Q. H., Chan, E., et al.: 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, 955–962, https://doi.org/10.1177/00333549221148174, 2023.

Administraţia Naţională de Meteorologie: National heat warning service, https://www.meteoromania.ro/avertizari/, romania; official information page, last accessed July 2026, 2026.

Antofie, T. et al.: DRMKC Risk Data Hub: datasets and indicators for disaster risk management and climate adaptation, Tech. Rep. JRC132193, European Commission, Joint Research Centre, https://doi.org/10.2760/13421, 2022.

Apreda, C., D'Ambrosio, V., and Di Martino, F.: A climate vulnerability and impact assessment model for complex urban systems, Environmental Science Policy, 93, 11–26, https://doi.org/10.1016/j.envsci.2018.12.016, 2019.

Austrian Federal Ministry of Social Affairs, Health, Care and Consumer Protection: National heat–health warning and action arrangements, https://www.sozialministerium.gv.at/Themen/Gesundheit/Hitze/Nationaler-Hitzeschutzplan.html, austria; official information page, last accessed July 2026, 2026.

Aznar-Siguan, G. and Bresch, D. N.: CLIMADA v1: a global weather and climate risk assessment platform, Geoscientific Model Development, 12, 3085–3097, https://doi.org/10.5194/gmd-12-3085-2019, 2019.

Belgian Federal Public Service Health, Food Chain Safety and Environment: National heat–health warning and action arrangements, https://www.health.belgium.be/fr/themes/sante/aide-medicale-urgente/plans-durgence/plan-fortes-chaleurs-pics-dozone, belgium; official information page, last accessed July 2026, 2026.

Bondarenko, M., Priyatikanto, R., Tejedor Garavito, N., Zhang, W., McKeen, T., Cunningham, A., Woods, T., Hilton, J., Cihan, D., Nosatiuk, B., Brinkhoff, T., Tatem, A., and Sorichetta, A.: Estimates of 2015-2030 total number of people per grid square broken down by gender and age groupings at a resolution of 3 arc (approximately 100m at the equator) R2025A version v1., https://doi.org/10.5258/SOTON/WP00841, 2025.

Bowler, D. E., Buyung-Ali, L., Knight, T. M., and Pullin, A. S.: Urban greening to cool towns and cities: A systematic review of the empirical evidence, Landscape and Urban Planning, 97, 147–155, https://doi.org/10.1016/j.landurbplan.2010.05.006, 2010.

Bresch, D. N. and Aznar-Siguan, G.: CLIMADA v1.4.1: towards a globally consistent adaptation options appraisal tool, Geoscientific Model Development, 14, 351–363, https://doi.org/10.5194/gmd-14-351-2021, 2021.

Brousseau, J. J., Stern, M. J., Pownall, M., and Hansen, L. J.: Understanding how justice is considered in climate adaptation approaches: a qualitative review of climate adaptation plans, Local Environment, 29, 1644–1663, https://doi.org/10.1080/13549839.2024.2386964, 2024.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Burgas Media: National heat warning service, https://burgasmedia.com/news/bg-alert-s-izvanredni-preduprezhdeniya-v-burgas-zaradi-opasni-temperaturi?lang=en, bulgaria; official information page, last accessed July 2026, 2026.

Burke, M., Wilson, A., Avirmed, T., Wallstein, J., Martins, M. C., Behrer, P., Callahan, C., Childs, M., Choi, J., French, K., Gould, C., Heft-Neal, S., Jing, R., Qiu, M., Rennels, L., and Southworth, E. K.: Understanding and Addressing Temperature Impacts on Mortality, Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w34313, 2025.

Buzási, A., Simoes, S. G., Salvia, M., Eckersley, P., Geneletti, D., Pietrapertosa, F., Olazabal, M., Wejs, A., De Gregorio Hurtado, S., Spyridaki, N.-A., Szalmáné Csete, M., Torres, E. F., Rižnar, K., Heidrich, O., Grafakos, S., and Reckien, D.: European patterns of local adaptation planning—a regional analysis, Regional Environmental Change, 24, https://doi.org/10.1007/s10113-024-02211-w, 2024.

Casanueva, A., Burgstall, A., Kotlarski, S., Messeri, A., Morabito, M., Flouris, A. D., Nybo, L., Spirig, C., and Schwierz, C.: Overview of Existing Heat-Health Warning Systems in Europe, International Journal of Environmental Research and Public Health, 16, 2657, https://doi.org/10.3390/ijerph16152657, 2019.

Chiabai, A., Spadaro, J. V., and Neumann, M. B.: Valuing deaths or years of life lost? Economic benefits of avoided mortality from early heat warning systems, Mitigation and Adaptation Strategies for Global Change, 23, 1159–1176, https://doi.org/10.1007/s11027-017-9778-4, 2018.

Chu, L., Akkiraju, K., Rao, N. D., Dubrow, R., and Chen, K.: Impact of household air conditioning usage on heat-related mortality in the USA, Nature Health, https://doi.org/10.1038/s44360-026-00184-0, 2026.

Chua, P. L. C., Takane, Y., Ng, C. F. S., Oka, K., Honda, Y., Kim, Y., and Hashizume, M.: Net impact of air conditioning on heat-related mortality in Japanese cities, Environment International, 181, 108 310, https://doi.org/10.1016/j.envint.2023.108310, 2023.

Creutzig, F., McPhearson, T., Bardhan, R., Belmin, C., Chow, W. T., Garschagen, M., Hsu, A., Kılkış, Ş., Islam, S. T., Milojevic-Dupont, N., et al.: Bridging the scale between the local particular and the global universal in climate change assessments of cities, Nature Cities, pp. 1–10, 2025.

Danmarks Meteorologiske Institut: National heat warning service, https://www.dmi.dk/varsler/warnings/, denmark; official information page, last accessed July 2026, 2026.

de Munck, C., Pigeon, G., Masson, V., Meunier, F., Bousquet, P., Tréméac, B., Merchat, M., Poeuf, P., and Marchadier, C.: How much can air conditioning increase air temperatures for a city like Paris, France?, International Journal of Climatology, 33, 210–227, https://doi.org/10.1002/joc.3415, 2013.

De Ridder, K., Lauwaet, D., and Maiheu, B.: UrbClim: A fast urban boundary layer climate model, Urban Climate, 12, 21–48, https://doi.org/10.1016/j.uclim.2015.01.001, 2015.

de'Donato, F., Scortichini, M., De Sario, M., de Martino, A., and Michelozzi, P.: Temporal variation in the effect of heat and the role of the Italian heat prevention plan, Public Health, 161, 154–162, https://doi.org/10.1016/j.puhe.2018.03.030, 2018.

Demuzere, M., Kittner, J., Martilli, A., Mills, G., Moede, C., Stewart, I. D., van Vliet, J., and Bechtel, B.: A global map of local climate zones to support earth system modelling and urban-scale environmental science, Earth System Science Data, 14, 3835–3873, https://doi.org/10.5194/essd-14-3835-2022, 2022.

Deutscher Wetterdienst: National heat warning service, https://www.dwd.de/DE/leistungen/hitzewarnung/hitzewarnung.html, germany; official information page, last accessed July 2026, 2026.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Dodman, D., Hayward, B., Pelling, M., Castan Broto, V., Chow, W., Chu, E., Dawson, R., Khirfan, L., McPhearson, T., Prakash, A., Zheng, Y., and Ziervogel, G.: Cities, Settlements and Key Infrastructure, in: Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, edited by Pörtner, H.-O., Roberts, D., Tignor, M., Poloczanska, E., Mintenbeck, K., Alegría, A., Craig, M., Langsdorf, S., Löschke, S., Möller, V., Okem, A., and Rama, B., pp. 907–1040, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, https://www.ipcc.ch/report/ar6/wg2/chapter/chapter-6/, 2022.

Državni hidrometeorološki zavod: National heat warning service, https://meteo.hr/naslovnica-upozorenja.php?lang=en&tab=upozorenja, croatia; official information page, last accessed July 2026, 2026.

Ebi, K. L., Teisberg, T. J., Kalkstein, L. S., Robinson, L., and Weiher, R. F.: Heat Watch/Warning Systems Save Lives: Estimated Costs and Benefits for Philadelphia 1995–98, Bulletin of the American Meteorological Society, 85, 1067–1074, https://doi.org/10.1175/BAMS-85-8-1067, 2004.

Economics of Climate Adaptation Working Group: Shaping climate-resilient development: a framework for decision-making, Tech. rep., ClimateWorks Foundation, Global Environment Facility, European Commission, McKinsey & Company, Rockefeller Foundation, Standard Chartered Bank, Swiss Re, https://climate-adapt.eea.europa.eu/en/metadata/publications/shaping-climate-resilient-development-a-framework-for-decision-making, 2009.

Estonian Environment Agency: National heat warning service, https://www.ilmateenistus.ee/ilm/prognoosid/hoiatused/?lang=en, estonia; official information page, last accessed July 2026, 2026.

EUROCONTROL: Standard Inputs for Economic Analyses, Tech. rep., EUROCONTROL, Performance Review Unit, https://ansperformance.eu/economics/cba/standard-inputs/latest/, edition 10.0.1, 2025.

European Commission, DG MOVE: Handbook on the External Costs of Transport, Tech. rep., Publications Office of the European Union, 2019.

Eurostat: Census 2021 population grid, https://ec.europa.eu/eurostat/web/gisco/geodata/population-distribution/population-grids, 2025.

Falchetta, G. and Hammad, A. T.: Tracking green space along streets of world cities, Environmental Research: Infrastructure and Sustainability, 5, 025 011, https://doi.org/10.1088/2634-4505/add9c4, 2025.

Falchetta, G., De Cian, E., Pavanello, F., and Wing, I. S.: Inequalities in global residential cooling energy use to 2050, Nature Communications, 15, 7874, https://doi.org/10.1038/s41467-024-52028-8, 2024.

Falchetta, G., De Cian, E., and Lunghi, J.: Street green space and electricity demand: Evidence from metered consumption data, Energy Economics, 158, 109 311, https://doi.org/10.1016/j.eneco.2026.109311, 2026a.

Falchetta, G., Lohrey, S., Souverijns, N., Lauwaet, D., Schleussner, C.-F., and Niamir, L.: Street green space is relevant but not sufficient for adapting to growing urban heat in world cities, Environmental Research Letters, 21, 084 012, https://doi.org/10.1088/1748-9326/ae5c20, 2026b.

Fan, J., Liu, B., Lei, T., Sun, Y., Ma, Y., Guo, R., Chen, D., Zhou, K., Li, S., and Gao, X.: Exploring how economic level drives urban flood risk, Nature Communications, 16, https://doi.org/10.1038/s41467-025-60267-6, 2025.

Gallo, E., Quijal-Zamorano, M., Méndez Turrubiates, R. F., Tonne, C., Basagaña, X., Achebak, H., and Ballester, J.: Heat-related mortality in Europe during 2023 and the role of adaptation in protecting health, Nature Medicine, 30, 3101–3105, https://doi.org/10.1038/s41591-024-03186-1, 2024.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., and ter Maat, J.: Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world, Global Environmental Change, 23, 485–498, https://doi.org/10.1016/j.gloenvcha.2012.12.006, 2013.

Han, D., Zhang, T., Qin, Y., Tan, Y., and Liu, J.: A comparative review on the mitigation strategies of urban heat island (UHI): a pathway for sustainable urban development, Climate and Development, 15, 379–403, 2023.

Hellenic National Meteorological Service: National heat warning service, https://emy.gr/en/warnings, greece; official information page, last accessed July 2026, 2026.

Huang, K., Li, X., Liu, X., and Seto, K. C.: Projecting global urban land expansion and heat island intensification through 2050, Environmental Research Letters, 14, 114 037, 2019.

Huang, W. T. K., Masselot, P., Bou-Zeid, E., Fatichi, S., Paschalis, A., Sun, T., Gasparrini, A., and Manoli, G.: Economic valuation of temperature-related mortality attributed to urban heat islands in European cities, Nature Communications, 14, 7438, https://doi.org/10.1038/s41467-023-43135-z, 2023.

Huisman, J., Martyr, R., Rott, R., and Smits, J.: Projections of climate change vulnerability along the Shared Socioeconomic Pathways 2020–2100, Scientific Data, 12, 1527, https://doi.org/10.1038/s41597-025-05732-z, 2025.

Ilmatieteen laitos: National heat warning service, https://www.ilmatieteenlaitos.fi/tietoa-helle-ja-pakkasvaroitukset, finland; official information page, last accessed July 2026, 2026.

Instituto Nacional de Saúde Doutor Ricardo Jorge: National heat–health warning and action arrangements, https://www.insa.min-saude.pt/conheca-o-sistema-de-monitorizacao-e-vigilancia-icaro/, portugal; official information page, last accessed July 2026, 2026.

Instytut Meteorologii i Gospodarki Wodnej: National heat warning service, https://meteo.imgw.pl, poland; official information page, last accessed July 2026, 2026.

Italian Ministry of Health: National heat–health warning and action arrangements, https://www.salute.gov.it/portale/caldo/dettaglioContenutiCaldo.jsp?lingua=italiano&id=408&area=emergenzaCaldo&menu=vuoto, italy; official information page, last accessed July 2026, 2026.

Iungman, T., Cirach, M., Marando, F., Pereira Barboza, E., Khomenko, S., Masselot, P., Quijal-Zamorano, M., Mueller, N., Gasparrini, A., Urquiza, J., Heris, M., Thondoo, M., and Nieuwenhuijsen, M.: Cooling cities through urban green infrastructure: a health impact assessment of European cities, The Lancet, 401, 577–589, https://doi.org/10.1016/S0140-6736(22)02585-5, 2023.

Jiang, T., Krayenhoff, E. S., Martilli, A., Nazarian, N., Stone, Jr., B., and Voogt, J. A.: Prioritizing urban heat adaptation infrastructure based on multiple outcomes: Comfort, health, and energy, Proceedings of the National Academy of Sciences, 122, https://doi.org/10.1073/pnas.2411144122, 2025.

Jin, L., Schubert, S., Hefny Salim, M., and Schneider, C.: Impact of air conditioning systems on the outdoor thermal environment during summer in Berlin, Germany, International Journal of Environmental Research and Public Health, 17, 4645, https://doi.org/10.3390/ijerph17134645, 2020.

Johar, H., Abdulsalam, F. I., Guo, Y., Baernighausen, T., Jahan, N. K., Watterson, J., Leder, K., Gouwanda, D., Ramanathan, G. L., Lee, K. K. C., et al.: Community-based heat adaptation interventions for improving heat literacy, behaviours, and health outcomes: a systematic review, The Lancet Planetary Health, 9, 2025.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Johnson, D., See, L., Oswald, S. M., Prokop, G., and Krisztin, T.: A cost–benefit analysis of implementing urban heat island adaptation measures in small- and medium-sized cities in Austria, Environment and Planning B: Urban Analytics and City Science, 48, 2326–2345, https://doi.org/10.1177/2399808320974689, 2021.

Karlický, J., Bareš, J., and Huszár, P.: Sensitivity of modeled urban climate to urban canopy parameters over central Europe, Environmental Research: Climate, 5, 015 007, https://doi.org/10.1088/2752-5295/ae228c, 2026.

KC, S., Moradhvaj, Potancokova, M., Adhikari, S., Yildiz, D., Mamolo, M., Sobotka, T., Zeman, K., Abel, G., Lutz, W., and Goujon, A.: Wittgenstein Center (WIC) Population and Human Capital Projections - 2023, https://doi.org/10.5281/ZENODO.14718294, 2024.

Khosla, R., Miranda, N. D., Trotter, P. A., Mazzone, A., Renaldi, R., McElroy, C., Cohen, F., Jani, A., Perera-Salazar, R., and McCulloch, M.: Cooling for sustainable development, Nature Sustainability, 4, 201–208, https://doi.org/10.1038/s41893-020-00627-w, 2021.

Kikegawa, Y., Tanaka, A., Ohashi, Y., Ihara, T., and Shigeta, Y.: Observed and simulated sensitivities of summertime urban surface air temperatures to anthropogenic heat in downtown areas of two Japanese major cities, Tokyo and Osaka, Theoretical and Applied Climatology, 117, 175–193, https://doi.org/10.1007/s00704-013-0996-8, 2014.

Kou, X., Poredoš, P., and Wang, R. Z.: Urban air conditioning is combating self-made island effects, Science Bulletin, https://doi.org/10.1016/j.scib.2026.03.039, 2026.

Kriit, H. K., Chen-Xu, J., Semenza, J. C., Heiliger, H., Markandya, A., Dasandi, N., Jankin, S., van Daalen, K. R., Achebak, H., Alari, A., et al.: The 2026 Europe report of the Lancet Countdown on health and climate change: narrowing window for decisive health action, The Lancet Public Health, 11, e386–e407, 2026.

Latvijas Vides, ģeoloģijas un meteoroloģijas centrs: National heat warning service, https://bridinajumi.meteo.lv/, latvia; official information page, last accessed July 2026, 2026.

Li, D., Varquez, A. C. G., Sun, T., Takane, Y., Yang, J., Ding, M., and Sailor, D.: Anthropogenic heat in urban climate systems: forcing, sensitivity, and feedback, arXiv, https://arxiv.org/abs/2608.17782, preprint, arXiv:2608.17782, 2026.

Li, X., Zhang, C., Li, W., Ricard, R., Meng, Q., and Zhang, W.: Assessing street-level urban greenery using Google Street View and a modified green view index, Urban Forestry & Urban Greening, 14, 675–685, https://doi.org/10.1016/j.ufug.2015.06.006, 2015.

Li, Y., Svenning, J.-C., Zhou, W., Zhu, K., Abrams, J. F., Lenton, T. M., Ripple, W. J., Yu, Z., Teng, S. N., Dunn, R. R., and Xu, C.: Green spaces provide substantial but unequal urban cooling globally, Nature Communications, 15, 7108, https://doi.org/10.1038/s41467-024-51355-0, 2024.

Lietuvos hidrometeorologijos tarnyba: National heat warning service, https://www.meteo.lt/prognozes/pavojingi-reiskiniai/, lithuania; official information page, last accessed July 2026, 2026.

Liu, K., Wang, R., Tukker, A., KC, S., and Hoekstra, R.: Projecting temperature-related mortality impacts at urban scale under climate scenarios: A methodological review, Sustainable Cities and Society, 145, 107 476, https://doi.org/10.1016/j.scs.2026.107476, 2026.

Malloy, J. T. and Ashcraft, C. M.: A framework for implementing socially just climate adaptation, Climatic Change, 160, 1–14, https://doi.org/10.1007/s10584-020-02705-6, 2020.

Markanday, A., Galarraga, I., and Markandya, A.: A Critical Review of Cost-Benefit Analysis for Climate Change Adaptation in Cities, Climate Change Economics, 10, 1950 014, https://doi.org/10.1142/S2010007819500143, 2019.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Masselot, P., Mistry, M. N., Vanoli, J., et al.: Excess mortality attributed to heat and cold: a health impact assessment study in 854 cities in Europe, The Lancet Planetary Health, 7, e271–e281, https://doi.org/10.1016/S2542-5196(23)00023-2, 2023.

Masselot, P., Mistry, M. N., Rao, S., et al.: Estimating future heat-related and cold-related mortality under climate change, demographic and adaptation scenarios in 854 European cities, Nature Medicine, 31, 1294–1302, https://doi.org/10.1038/s41591-024-03452-2, 2025.

Matthews, T., Raymond, C., Foster, J., Baldwin, J. W., Ivanovich, C., Kong, Q., Kinney, P., and Horton, R. M.: Mortality impacts of the most extreme heat events, Nature Reviews Earth & Environment, 6, 193–210, 2025.

Mazzone, A., De Cian, E., Falchetta, G., Jani, A., Mistry, M. N., and Khosla, R.: Understanding systemic cooling poverty, Nature Sustainability, 6, 1533–1541, https://doi.org/10.1038/s41893-023-01221-6, 2023.

Met Éireann: National heat warning service, https://www.met.ie/weather-warnings, ireland; official information page, last accessed July 2026, 2026.

Nemzeti Népegészségügyi Központ: National heat–health warning and action arrangements, https://nnk.gov.hu/index.php/kozegeszsegugyi-hatosagi-ugyek/telepules-egeszsegugyi-klimavaltozas-es-kornyezeti-egeszseghatas-elemzo-osztaly/temainak/hosegriasztas.html, hungary; official information page, last accessed July 2026, 2026.

Norton, B. A., Coutts, A. M., Livesley, S. J., Harris, R. J., Hunter, A. M., and Williams, N. S. G.: Planning for cooler cities: A framework to prioritise green infrastructure to mitigate high temperatures in urban landscapes, Landscape and Urban Planning, 134, 127–138, https://doi.org/10.1016/j.landurbplan.2014.10.018, 2015.

Pavanello, F. and Valenti, G.: Mortality, Temperature, and Public Adaptation Policy: Evidence from Italy, https://doi.org/10.2139/ssrn.6034714, fEEM Working Paper No. 35-2025, 2025.

Rijksinstituut voor Volksgezondheid en Milieu: National heat–health warning and action arrangements, https://www.rivm.nl/en/heat/national-heatwave-plan, netherlands; official information page, last accessed July 2026, 2026.

Salamanca, F., Georgescu, M., Mahalov, A., Moustaoui, M., and Wang, M.: Anthropogenic heating of the urban environment due to air conditioning, Journal of Geophysical Research: Atmospheres, 119, 5949–5965, https://doi.org/10.1002/2013JD021225, 2014.

Santamouris, M. and Vasilakopoulou, K.: Passive cooling for the built environment, Nature Reviews Clean Technology, pp. 1–17, 2026.

Santé publique France: National heat–health warning and action arrangements, https://www.santepubliquefrance.fr/climat/fortes-chaleurs-canicule, france; official information page, last accessed July 2026, 2026.

Sera, F., Hashizume, M., Honda, Y., Lavigne, E., Schwartz, J., Zanobetti, A., Tobias, A., Iñiguez, C., Vicedo-Cabrera, A. M., Blangiardo, M., Armstrong, B., and Gasparrini, A.: Air Conditioning and Heat-related Mortality: A Multi-country Longitudinal Study, Epidemiology, 31, 779–787, https://doi.org/10.1097/EDE.0000000000001241, 2020.

Sharifi, A., Pathak, M., Joshi, C., and He, B.-J.: A systematic review of the health co-benefits of urban climate change adaptation, Sustainable Cities and Society, 74, 103 190, https://doi.org/10.1016/j.scs.2021.103190, 2021.

Simpson, N. P., Simpson, K. J., Ferreira, A. T., Constable, A., Glavovic, B., Eriksen, S. E. H., Ley, D., Solecki, W., Rodríguez, R. S., and Stringer, L. C.: Climate-resilient development planning for cities: progress from Cape Town, Npj Urban Sustainability, 3, https://doi.org/10.1038/s42949-023-00089-x, 2023.

Slovenian Environment Agency: National heat warning service, https://meteo.arso.gov.si/met/en/warning/, slovenia; official information page, last accessed July 2026, 2026.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Slovenský hydrometeorologický ústav: National heat warning service, https://www.shmu.sk/en/?page=987, slovakia; official information page, last accessed July 2026, 2026.

Souverijns, N., Lauwaet, D., Lejeune, Q., Kropf, C. M., Yeung, K. L., Nath, S., and Schleussner, C. F.: 100m climate and heat stress data up to 2100 for 142 cities around the globe, Data in Brief, p. 112497, 2026.

Spanish Ministry of Health: National heat–health warning and action arrangements, https://www.sanidad.gob.es/areas/sanidadAmbiental/riesgosAmbientales/calorExtremo/home.htm, spain; official information page, last accessed July 2026, 2026.

Stalhandske, Z., Nesa, V., Zumwald, M., Ragettli, M. S., Galimshina, A., Holthausen, N., Röösli, M., and Bresch, D. N.: Projected impact of heat on mortality and labour productivity under climate change in Switzerland, Natural Hazards and Earth System Sciences, 22, 2531–2541, https://doi.org/10.5194/nhess-22-2531-2022, 2022.

Stewart, I. D. and Oke, T. R.: Local Climate Zones for Urban Temperature Studies, Bulletin of the American Meteorological Society, 93, 1879–1900, https://doi.org/10.1175/BAMS-D-11-00019.1, 2012.

Sveriges meteorologiska och hydrologiska institut: National heat warning service, https://www.smhi.se/kunskapsbanken/meteorologi/varningar-och-meddelanden/varning-for-hoga-temperaturer, sweden; official information page, last accessed July 2026, 2026.

Toloo, G., FitzGerald, G., Aitken, P., Verrall, K., and Tong, S.: Evaluating the effectiveness of heat warning systems: systematic review of epidemiological evidence, International Journal of Public Health, 58, 667–681, https://doi.org/10.1007/s00038-013-0465-2, 2013.

Tompkins, E. L. and Eakin, H.: Managing private and public adaptation to climate change, Global Environmental Change, 22, 3–11, https://doi.org/10.1016/j.gloenvcha.2011.09.010, 2012.

Tréméac, B., Bousquet, P., de Munck, C., Pigeon, G., Masson, V., Marchadier, C., Merchat, M., Poeuf, P., and Meunier, F.: Influence of air conditioning management on heat island in Paris air street temperatures, Applied Energy, 95, 102–110, https://doi.org/10.1016/j.apenergy.2012.02.015, 2012.

Uhl, J. H., Politis, P., and Pesaresi, M.: GHS-AGE R2025A: Global gridded estimates of the dominant age of the built stock (1975–2020), https://doi.org/10.2905/d503bb56-9884-4e4d-bb8f-d86711d9f749, 2026.

Urban, A., Huber, V., Henry, S., Plaza, N. P., Tušlová, L., Dasgupta, S., Masselot, P., Cvijanovic, I., Mistry, M., Pascal, M., et al.: The effectiveness of heat prevention plans in reducing heat-related mortality across Europe, Environmental Research Letters, 20, 124 071, https://doi.org/10.1088/1748-9326/ae2775, 2025.

Český hydrometeorologický ústav: National heat warning service, https://vystrahy-cr.chmi.cz/, czech Republic; official information page, last accessed July 2026, 2026.

Walker, S. E., Smith, E. A., Bennett, N., Bannister, E., Narayana, A., Nuckols, T., Pineda Velez, K., Wrigley, J., and Bailey, K. M.: Defining and conceptualizing equity and justice in climate adaptation, Global Environmental Change, 87, 102 885, https://doi.org/10.1016/j.gloenvcha.2024.102885, 2024.

Wang, L., Sun, T., Zhou, W., Liu, M., and Li, D.: Deciphering the sensitivity of urban canopy air temperature to anthropogenic heat flux with a forcing-feedback framework, Environmental Research Letters, 18, 094 005, https://doi.org/10.1088/1748-9326/ace7e0, 2023.

https://doi.org/10.5194/egusphere-2026-5357
Preprint. Discussion started: 16 November 2026
© Author(s) 2026. CC BY 4.0 License.

Wang, S., Zhan, W., Zhou, B., Tong, S., Chakraborty, T. C., Wang, Z., Huang, K., Du, H., Middel, A., Li, J., Liu, Z., Li, L., Huang, F., and Li, M.: Dual impact of global urban overheating on mortality, Nat. Clim. Change, 15, 497–504, https://doi.org/10.1038/s41558-025-02303-3, 2025.

Wetterhall, F., Polster, C., Vuckovic, M., and Shtivelman, A.: The CLIMAAX handbook - an online tool for regional climate risk assessments, https://doi.org/10.5194/ems2025-685, eGU General Assembly 2025, abstract EMS2025-685, 2025.

World Health Organization Regional Office for Europe: Heat and health in the WHO European Region: updated evidence for effective prevention, WHO Regional Office for Europe, Copenhagen, ISBN 9789289055406, https://iris.who.int/handle/10665/339462, 2021.

Yang, Q., Xu, Y., Chakraborty, T., Du, M., Hu, T., Zhang, L., Liu, Y., Yao, R., Yang, J., Chen, S., et al.: A global urban heat island intensity dataset: Generation, comparison, and analysis, Remote Sensing of Environment, 312, 114 343, 2024.