

Heat Market: Methodology

Who works and lives in dangerous heat, now and to 2090

heat-market project

September 2026

Abstract

This document describes how the heat-market dashboard estimates how many people, and which customer segments, face a given number of dangerously hot days a year, for the recent climate and for every decade to 2090 under two emissions pathways. It covers the climate data and heat metrics, the station check and calibration, where people are placed, the population projections, the worker, earnings and poverty segments for the United States and for every other country, the work-capacity and wages-at-risk calculation, how the dashboard interpolates between decades, and the main limitations. Every number on the dashboard can be traced to a step described here and to the code in the `heatmarket` package.

Contents

1	What the dashboard reports	2
2	Climate data	2
2.1	Source	2
2.2	Periods and sampling	2
2.3	Afternoon heat	3
2.4	Heat metrics	3
2.5	Day counts and caching	3
3	Checking against stations, and calibration	3
4	Where people live	4
4.1	Population grids	4
4.2	From polygons to climate cells	4
4.3	Rural and urban people	4
4.4	Boundaries	4
5	Population projections	5
6	Workers, earnings and poverty	5
6.1	United States	5
6.2	Every other country	5
6.3	India	6
6.4	Europe	6
7	Work lost to heat, and wages at risk	7
8	The world map and years in between	7
9	Limitations	7

1 What the dashboard reports

The core quantity follows the FEWS NET heat exposure projections: for a place (a county, a district, a country, the world), a period and a heat threshold, the number of people in each *days-per-year* class

$$\{< 1, 1-9, 10-29, 30-59, 60-89, 90+\}$$

of days on which the afternoon heat reaches the threshold. The headline is the number of people at or above a chosen number of days (1, 10, 30, 60 or 90). This is shown for

- **a world view:** every country, and provinces and districts for 201 countries, with a year slider from the 1995–2014 baseline to 2090;
- **detail views** in the format of the United States county view, for the United States, India, Europe (NUTS 3 regions), China, Indonesia, Nigeria, Pakistan, Bangladesh, Brazil, Vietnam and the Philippines, with segments by industry, income, age and poverty, and work hours and wages at risk.

Segments are subsets of the population: everyone, outdoor workers (optionally by industry and minimum income), all workers, people aged 60+, children under five, and people in poverty (a proxy for those most likely to need heat shelter).

2 Climate data

2.1 Source

Daily data come from NASA’s NEX-GDDP-CMIP6 [1], a statistically downscaled and bias-corrected version of CMIP6 on a global 0.25° grid (about 25 km), read directly from the public Amazon S3 bucket. Each file holds one variable for one model, scenario and year. The code reads only the grid cells and days it needs, in 31-day slabs.

Region	Configuration
Contiguous US	5 models (GFDL-ESM4, MRI-ESM2-0, MPI-ESM1-2-HR, ACCESS-CM2, NorESM2-MM); April–October; every second year of each 20-year window; windows 1995–2014, 2021–40, 2041–60.
Contiguous US, Liljegen	2 models (GFDL-ESM4, MRI-ESM2-0); as above, with sunlight and wind.
World	3 models (GFDL-ESM4, MRI-ESM2-0, MPI-ESM1-2-HR); whole year; every fourth year counted from 2021; windows centred on every decade 2030–2090 plus 1995–2014.

Table 1: Climate runs.

Two emissions pathways are used: SSP2-4.5 (intermediate) and SSP5-8.5 (very high). The historical runs supply the 1995–2014 baseline.

2.2 Periods and sampling

A single year’s weather is dominated by natural variability, so every future point is a 20-year window: the window for decade D covers years $D - 9$ to $D + 10$ (2050 is 2041–2060; 2090 is 2081–2100, where the data end). For the world, years are sampled every fourth year counted from 2021 (2021, 2025, ..., 2097). Neighbouring windows therefore share their sampled years, and each holds five. For the US, every second year is used.

2.3 Afternoon heat

For each day, heat is evaluated at the daily maximum temperature T (`tasmax`) with the day’s specific humidity q (`huss`). Specific humidity changes much less over a day than relative humidity, so pairing $T_{\max}$ with daily-mean q approximates afternoon conditions. Pairing $T_{\max}$ with daily-mean relative humidity would overstate humidity at the peak. Vapour pressure and relative humidity are

$$e = \frac{qp}{0.622 + 0.378q}, \quad p = 1013.25 \text{ hPa}, \quad (1)$$

$$e_s(T) = 6.112 \exp\left(\frac{17.67T}{T + 243.5}\right) \text{ hPa (Bolton 1980)}, \quad \text{RH} = 100 \min(e, e_s)/e_s. \quad (2)$$

2.4 Heat metrics

Wet-bulb temperature (Stull 2011 [2]):

$$T_w = T \arctan(0.151977\sqrt{\text{RH} + 8.313659}) + \arctan(T + \text{RH}) - \arctan(\text{RH} - 1.676331) + 0.00391838 \text{RH}^{3/2} \arctan(0.023101 \text{RH}) - 4.686035. \quad (3)$$

Thresholds: 24, 26, 28, 30 and 31 °C. Survivability research points to limits near 31 °C for healthy adults at rest [3], lower for older people.

Simplified WBGT (sWBGT) , the Australian Bureau of Meteorology’s approximation for full sun and light wind:

$$\text{sWBGT} = 0.567T + 0.393e + 3.94. \quad (4)$$

Thresholds: 28, 30, 32 and 34 °C. It assumes full sun and reads about 2 °C higher than a shaded WBGT, so it overstates heat for shaded or indoor work.

WBGT in sun and shade (Liljegren et al. 2008 [4], US only), the model behind most occupational WBGT estimates, run with the reference C implementation (`pywbgt`). Inputs are placed at 2 pm solar time:

- the daily maximum temperature, and dew point from daily specific humidity;
- surface pressure from each cell’s mean elevation (Copernicus GLO-90 DEM);
- daily-mean 10 m wind, which the model brings down to 2 m;
- incoming shortwave radiation scaled from the daily mean by the sun’s height at 2 pm, which assumes the day’s cloud is spread evenly.

The shade version sets radiation to zero. Thresholds: 28, 30 and 32 °C in sun; 26, 28 and 30 °C in shade.

2.5 Day counts and caching

For every cell and year, the pipeline stores a histogram of days in 0.5 °C steps (wet-bulb 20–36 °C; WBGT 24–44 °C). The number of days at or above any threshold on that grid is then a sum over the histogram. This lets thresholds and calibration offsets change without re-reading the climate data. For each model and period, days per year are averaged over the window’s sampled years. The ensemble value is the mean over models. The per-model values give the model range shown as whiskers.

3 Checking against stations, and calibration

The modelled baseline was checked against NOAA GHCN-Daily records (daily maximum temperature and average dew point) at 20 US airports, April–October 1995–2014, using years with at least 90% coverage. Station wet-bulb and sWBGT were computed with the same formulas.

The models' temperatures were close, but their humidity ran 2–5 °C of dew point too high in the humid East. As a result they produced roughly three times too many humid-heat days. For each metric family, a threshold offset δ_s was fitted at each station: the shift, in 0.5 °C steps, that best matches observed day counts. The offsets are interpolated between stations by inverse-distance weighting. Days are then counted at or above threshold + $\delta(\text{cell})$.

Evaluated by leave-one-station-out (sum of absolute log ratios of day counts), the error falls

- from 56.1 to 19.1 for wet-bulb, and
- from 40.4 to 24.1 for sWBGT.

Liljegren WBGT takes 0.7 of the sWBGT offset, its measured sensitivity to dew point. Outside the US there are no dew-point stations in this build. There the US-wide mean offsets (wet-bulb +2.0 °C, sWBGT +1.5 °C) are applied everywhere as a stopgap.

4 Where people live

4.1 Population grids

- **Meta High Resolution Settlement Layer** (HRSL, about 30 m) [6], total and its 60-and-over and under-5 layers, read from Meta's public cloud-optimised GeoTIFFs and summed onto a 0.025° grid, which nests in the 0.25° climate grid.
- **Kontur Population** (2023, H3 resolution-6 hexagons of about 36 km²) [7] for the roughly fifty countries HRSL does not cover (among them Australia, Iran, China, Russia, Israel, Syria and Yemen). Each hexagon's population is placed at its centre on the same grid.

For a set of areas in one country, HRSL is used unless its total is under half of Kontur's. In that case every area in the set uses Kontur. Choosing one source per country avoids shifting people between small neighbouring districts, which Kontur's centre-point placement would otherwise do. If neither source reaches half of Natural Earth's population estimate for a country, that estimate is spread evenly over the land. Age layers exist only for HRSL, and an area that touches a missing age tile gets no age figure.

4.2 From polygons to climate cells

Each area (country, province, district, county) is rasterised onto the 0.025° grid. Its people are counted in the climate cells where they live. People in coastal cells that the climate data mask as sea are assigned to the nearest land cell. For area a with people w_{ac} in climate cell c and $W_a = \sum_c w_{ac}$, the share of people in days-class k is

$$s_{ak} = \frac{1}{W_a} \sum_c w_{ac} \mathbf{1}[\text{days}_c \in k], \quad (5)$$

computed for the ensemble mean and for each model. The mean days shown on maps is $\sum_c w_{ac} \text{days}_c / W_a$.

4.3 Rural and urban people

A fine cell is rural when its surrounding 0.1° block (about 11 km) holds under 300 people per km², the GHSL line between rural areas and urban clusters [8]. Averaging over the block keeps Kontur's hexagon-centre points from reading as towns. Rural population is used to place farm jobs and to apply rural and urban poverty rates.

4.4 Boundaries

- Countries: Natural Earth 1:50m.
- Europe: Eurostat NUTS 2021 (NUTS 2 as provinces, NUTS 3 as districts).
- United States: Census county polygons with FIPS codes; states are their union.

- India detail view: the 640 districts of the 2011 Census (datameet boundaries).
- Everywhere else: geoBoundaries gbOpen, ADM1 and ADM2 [9].

A geoBoundaries area is dropped when more than half of it lies in another country on the world map. For China, Taiwan, Hong Kong and Macau are also dropped by name, because Natural Earth’s outlines for the last two are too coarse for the overlap test. This way places shown as their own entries are not counted twice.

5 Population projections

Future population comes from the IIASA SSP database v3.2 (2025), the Wittgenstein Centre’s projections by country, age, sex and education [10, 11]. These are read from the copy bundled with the open GCAM model’s input data [12]. Each climate pathway is paired with its own population pathway: SSP2-4.5 with SSP2 and SSP5-8.5 with SSP5. Four groups are projected: everyone, ages 60+, ages 0–4, and working age (15–64), which the worker segments follow.

Today’s population (HRSL, about 2020) is grown by its country’s factor

$$g_{i,G}(y) = \frac{P_{i,G}(y)}{P_{i,G}(2020)}, \quad (6)$$

interpolated linearly between the database’s five-year steps. Here i is the country, G the group and y the year (the window’s centre). A segment S in area a of country i is therefore $S_a(y) = S_a(2020) g_{i,G(S)}(y)$. Every area in a country grows at the same rate: sub-national shifts, such as urbanisation, are not modelled. The baseline climate is paired with today’s population. Places the database does not project take the growth of the country they sit in (for example Somaliland from Somalia) or hold today’s population (microstates). The dashboard has a switch to hold population at today’s level, which isolates the effect of climate.

With these projections, world population goes from 7.8 billion (2020) to 9.6 billion (2050) and 9.9 billion (2100) under SSP2, and to 9.2 and 8.1 billion under SSP5. India’s 60+ population grows 2.4 times by 2050 under SSP2; Nigeria’s total population roughly doubles.

For the US county view, which has its own ACS-based baseline, counties are grown by the US national factors. A county-level SSP projection file (Hauer 2019) is used instead when present.

6 Workers, earnings and poverty

6.1 United States

- **Workers:** ACS PUMS 2023 person records for people at work, weighted, by industry (NAICS sectors) and earnings band (wages plus self-employment income: under \$25k, \$25–50k, \$50–75k, \$75–100k, \$100k+).
- **Outdoor workers:** each worker is weighted by the share of their occupation’s O*NET 29.3 Work Context respondents who work “outdoors, exposed to weather” every day (element 4.C.2.a.1.c, category 5).
- **Counties:** PUMS areas (PUMAs) are split across counties by 2020 census-tract population.
- **Poverty:** people below the poverty line, and those aged 65+ below it (ACS 5-year B17001), as the shelter proxy.

6.2 Every other country

1. **Jobs by industry:** ILO modelled estimates (November 2024) of employment in 14 ISIC Rev.4 sections, each country’s latest year (2023 for nearly all) [13].

2. **Placing jobs:** within a country, farm jobs (section A) are placed in proportion to each area’s rural population, and every other section in proportion to total population:

$$J_{a,\text{farm}} = E_{\text{farm}} \frac{R_a}{\sum_b R_b}, \quad J_{a,s} = E_s \frac{P_a}{\sum_b P_b} \quad (s \neq \text{farm}). \quad (7)$$

3. **Outdoor share:** each section takes the pooled US outdoor share of the matching NAICS sectors (Table 2). The same occupational mix within an industry is assumed everywhere. Farm and building work is, if anything, more outdoor in lower-income countries, so this errs low there.
4. **Earnings:** labour income per worker from the Penn World Table 11 [14]: the labour share times output-side GDP divided by employment, in 2021 PPP dollars. This is one national average for every industry. Where PWT has no labour share for a country, the median share across countries is used.
5. **Poverty:** the UNDP/OPHI global Multidimensional Poverty Index 2022 [15]. The country’s rural and urban headcount ratios are applied to each area’s rural and urban people. It covers about 110 low- and middle-income countries and no EU country.

ISIC section	US sectors pooled	Outdoor share
A Agriculture, forestry, fishing	Agriculture	0.35
B Mining	Mining, oil & gas	0.35
C Manufacturing	Manufacturing	0.12
D–E Utilities	Utilities	0.27
F Construction	Construction	0.46
G Trade	Wholesale, retail	0.18
H, J Transport, storage, communication	Transportation, information	0.37
I Accommodation & food	Accommodation & food	0.08
K Finance	Finance & insurance	0.06
L–N Real estate, business services	Real estate, professional, management, admin. & waste	0.14
O Public administration	Public administration	0.21
P Education	Education	0.10
Q Health	Health care	0.05
R–U Other services	Arts & recreation, other services	0.13

Table 2: Outdoor share of workers by industry, from ACS PUMS $\times$ O*NET (US, 2023).

6.3 India

The India detail view uses the 640 districts of the 2011 Census, whose district tables supply the placement of the ILO’s 2023 national totals:

- farm jobs follow cultivators plus agricultural labourers;
- manufacturing follows household-industry plus other workers;
- every other section follows other workers.

Each district’s jobs are split across its households’ annual income bands (the Census “power parity” bands, from under Rs 45,000 to over Rs 545,000). This stands in for workers’ own pay and drives the minimum-income filter. The poverty segment is people in households under Rs 45,000 a year. The world map uses today’s districts (geoBoundaries) with the global method.

6.4 Europe

The Europe detail view covers NUTS3 regions of every NUTS country, mainland and nearby islands. Jobs follow the global method country by country. Eurostat’s regional labour-force and poverty tables would sharpen this but were not reachable from the build environment.

7 Work lost to heat, and wages at risk

Heavy-work capacity falls with WBGT following Dunne et al. (2013) [5]:

$$\text{loss}(\text{WBGT}) = \min(1, 0.25 \max(\text{WBGT} - 25, 0)^{2/3}). \quad (8)$$

Each day’s afternoon WBGT is applied to the 4 hottest hours of a shift, 5 days a week, over the counted season. The US uses April–October; everywhere else uses the whole year. Hours lost per worker per year are

$$H = 4 \times \frac{5}{7} \sum_{d \in \text{season}} \text{loss}(\text{WBGT}_d), \quad (9)$$

averaged over the model years and over where the segment’s people live. Wages at risk are

$$\text{Wages at risk} = H \times \text{workers} \times \frac{\text{annual earnings}}{2000 \text{ hours}}. \quad (10)$$

In the US, earnings come from PUMS by county, industry and band. Elsewhere it is the national labour income per worker (PPP). The US offers WBGT in shade, in sun and sWBGT. Other countries use sWBGT, which assumes full sun, so their hours lost are upper bounds for most jobs.

8 The world map and years in between

Any year. For a year y between two window centres $y_0 < y < y_1$ (the baseline sits at 2005), every quantity is a straight-line blend,

$$x(y) = x(y_0) + \frac{y - y_0}{y_1 - y_0} (x(y_1) - x(y_0)), \quad (11)$$

applied to days-class shares, hours lost and heat-day maps, while population uses its own five-yearly projection. This describes the climate around year y , not a forecast of that year’s weather.

Layers. The *heat days* layer is the ensemble-mean days per 0.25° cell. The *people* layer is population per km² on a 0.1° grid, from HRSL, or Kontur in any 1° block where HRSL holds under half of Kontur’s people. It is grown by each cell’s country projection through a country-index raster. The *people* × *heat* layer is people per km² times the heat days of the climate cell underneath, which gives person-days of exposure per km². The *share of people* layer colours each country or area by the share of the chosen segment at or above the headline.

Encoding. Days-class shares are stored in thousandths, as base-36 text per period. Heat-day rasters are stacked grey PNGs (grey = days × 255/366) that the page blends and colours. Data files are zipped JSON.

9 Limitations

- **Daily resolution.** Afternoon heat is approximated from daily maximum temperature and daily-mean humidity. The Liljegren inputs use daily-mean sunlight and wind.
- **Humidity bias outside the US.** The station calibration is US-only; the uniform offsets elsewhere are a stopgap. GHCN has usable dew-point records in India, the Gulf and Australia, which would allow regional calibration.
- **Few models.** Three models globally (five in the US) under-sample the CMIP6 spread. The whiskers show only the range across the models used.
- **Population.** Growth is national: urbanisation and internal migration are not modelled. HRSL’s 60+ layer undercounts (about 18% of the US against the Census’s 23%). Kontur has no age groups.

- **Workers outside the US.** Industry mixes are national (except India’s farm, household and other split). Outdoor shares are US shares. Earnings are one national average. Informal work is included in the ILO totals but not distinguished.
- **Poverty.** The MPI is a multidimensional deprivation measure, not an income line. It is national with a rural/urban split, and it is absent for high-income countries.
- **Adaptation.** No change in work hours, shade, cooling or acclimatisation is assumed. Hours lost are a hazard measure, not a forecast of lost output.

10 Reproducing the results

```
python -m heatmarket heat                # US heat days (5 models)
python -m heatmarket liljegren           # US WBGT in sun and shade
python -m heatmarket heat --region global \
    --models GFDL-ESM4 MRI-ESM2-0 MPI-ESM1-2-HR
python -m heatmarket segments            # US workers, earnings, poverty
python -m heatmarket validate            # NOAA stations and calibration
python -m heatmarket atlas --fetch       # world, provinces, districts, detail tabs
python -m heatmarket build               # US county view
python -m heatmarket stage <dir>        # a copy for text-only hosts
```

Every download step caches its results and resumes where it stopped.

References

- [1] Thrasher, B. et al. (2022). NASA Global Daily Downscaled Projections, CMIP6. *Scientific Data* 9, 262.
- [2] Stull, R. (2011). Wet-bulb temperature from relative humidity and air temperature. *J. Appl. Meteor. Climatol.* 50, 2267–2269.
- [3] Vecellio, D. J. et al. (2022). Evaluating the 35 °C wet-bulb temperature adaptability threshold for young, healthy subjects (PSU HEAT Project). *J. Appl. Physiol.* 132, 340–345.
- [4] Liljegren, J. C. et al. (2008). Modeling the wet bulb globe temperature using standard meteorological measurements. *J. Occup. Environ. Hyg.* 5, 645–655.
- [5] Dunne, J. P., Stouffer, R. J., John, J. G. (2013). Reductions in labour capacity from heat stress under climate warming. *Nature Climate Change* 3, 563–566.
- [6] Meta and CIESIN (2016–2022). High Resolution Settlement Layer. Data for Good at Meta.
- [7] Kontur (2023). Kontur Population: global population density for 400 m H3 hexagons, release 2023-11-01.
- [8] Schiavina, M. et al. (2023). GHS-SMOD R2023A: GHS settlement layers, application of the Degree of Urbanisation methodology. European Commission, JRC.
- [9] Runfola, D. et al. (2020). geoBoundaries: a global database of political administrative boundaries. *PLoS ONE* 15(4), e0231866.
- [10] KC, S. et al. (2024). Updating the Shared Socioeconomic Pathways (SSPs) global population and human capital projections. IIASA Working Paper WP-24-003.
- [11] IIASA (2025). SSP Scenario Database, version 3.2.
- [12] Joint Global Change Research Institute. GCAM v8 input data, `gcamdata/inst/extdata/socioeconomics`.
- [13] International Labour Organization (2024). ILO modelled estimates: employment by sex and economic activity (ILOEST, November 2024).

- [14] Feenstra, R. C., Inklaar, R., Timmer, M. P. (2015). The next generation of the Penn World Table. *American Economic Review* 105(10), 3150–3182; PWT version 11.0.
- [15] UNDP and OPHI (2022). Global Multidimensional Poverty Index 2022.
- [16] FEWS NET (2024). Heat exposure projections.