UK Heat Split: data sources and estimation methodology

Version 8 · Causeway Energies · 19 August 2026 · site v8.51.0

causewaygt.github.io/uk-heatsplit · contact@causewaygt.com

Abstract

UK Heat Split is a public, automatically updated estimate of the weekly heating and cooling energy balance of Great Britain. Daily gas throughput to the low-pressure distribution network is decomposed into a temperature-sensitive space-heating component and a weather-independent baseline by segmented linear regression on population-weighted degree days, following the established gas-to-heat method of Watson et al. [1,2] and Sansom [3]. The trailing twelve-month space-heating estimate is calibrated against the Department for Energy Security and Net Zero (DESNZ) Energy Consumption in the UK (ECUK) end-use tables [4]. Non-gas fuels and space cooling, which lack daily operational feeds, are held at annual ECUK levels and modulated by degree days. Delivered service is derived from purchased energy using in-situ combustion efficiencies [5] and measured seasonal heat-pump performance factors [6,7]. Cooling demand in the unequipped building stock is estimated independently from overheating degree-hours and dwelling-stock statistics [9,10]. All emissions use DESNZ conversion factors [11]; all monetary values use the prevailing Ofgem price cap [12]. The headline estimates are additionally published as a weekly historical series over the window the live gas feed serves, and as an hourly retrospective tested against two winters of observed system operation. Every figure resting on author judgement is flagged; the complete source code and data are public.

How to read this document. Each section states the method as it currently stands. Where an approach was tried and abandoned, or where a result may be useful to others working on the same problem, a short note follows under *Notes for investigators* — these are set apart because they are not part of the method. The complete record of corrections, including several that changed published figures materially, is consolidated in Section 13.

---

1. Conventions

1.1 Sourced quantities and author estimates

Every figure on the site is either traceable to a published source or is a Causeway Energies estimate. Estimates are marked with a dagger (†) wherever they appear, both here and on the site, and the contact address for challenge is published beside them. The distinction denotes provenance, not doubt: a well-founded estimate still carries the dagger.

1.2 Purchased energy and delivered service

Two quantities are reported throughout and must not be conflated.

Purchased energy is what was bought: gas at the meter, electricity at the grid supply point, oil delivered. It is what appears on a bill.

Delivered service is what the building received: useful heat after combustion losses, useful cooling after the refrigeration cycle. For heat pumps delivered service exceeds purchased energy, because the difference is lifted from the environment rather than bought.

The energy-balance panel shows both. The gap between them is the sum of combustion losses (a loss) and ambient heat harvested (a gain); these are opposite in sign and are never netted into a single "efficiency".

1.3 Three services, not two

Heat, cooling and **ventilation** are reported separately. Ventilation is fan energy: it moves air rather than heat, has no coefficient of performance, and dissipates inside the space as a cooling *load* rather than a cooling output. It is counted in neither the heat nor the cooling total, and no route in the geothermal what-if displaces it.

---

2. Primary heating estimate: gas-to-heat regression

2.1 Data source

The primary feed is the National Gas Transmission operational-data REST API [13], from which the daily aggregate demand of the thirteen Local Distribution Zones (LDZ) is retrieved and summed. Publication identifiers are resolved from the API catalogue at runtime rather than hard-coded, so a renamed publication produces a diagnostic rather than a silent error. Total National Transmission System (NTS) demand is retrieved separately for context.

The LDZ aggregate is used in preference to the NTS total because the latter includes gas-fired power generation and directly connected industry, which are unrelated to building heat and would dominate the signal. The distinction is material: summed LDZ demand is of order 480 GWh/day at summer minimum against an NTS total of 1,000–1,400 GWh/day.

2.2 Temperature index

Outdoor temperature is a population-weighted mean over twelve Great British urban centres (London, Birmingham, Manchester, Leeds, Glasgow, Newcastle, Bristol, Cardiff, Edinburgh, Southampton, Nottingham, Sheffield), weighted by Office for National Statistics mid-year population estimates. Hourly and daily 2 m air temperatures come from the ERA5 reanalysis [14] via the Open-Meteo archive [15].

Heating degree days at base T_e are HDD = $\max(0, T_e - \bar{T})$, with $\bar{T}$ the weighted daily mean; cooling degree days are CDD = $\max(0, \bar{T} - T_c)$. Degree days are fetched at heating bases 14.5, 15.5 and 16.5 °C and cooling bases 15.5, 18.0, 20.0 and 22.0 °C. **CDD is convex in its base and the series cannot be interpolated between them** — averaging two bases overstates the middle.

2.3 Estimator

Daily LDZ demand G is modelled as a segmented (piecewise-linear) function of temperature, equivalent to ordinary least squares on the degree-day transform:

$$G_i = \alpha + \beta \cdot HDD_i + \varepsilon_i$$

where α (GWh/day) is the weather-independent baseline — hot water, cooking and non-heating gas on the distribution network — β (GWh per degree day) is the marginal space-heating response, and ε is a residual. The fitted space-heating component on day i is $\beta \cdot HDD_i$.

This is the decomposition of Watson et al. [1,2], building on Sansom [3], and is standard in the GB heat-demand literature. The heating base is selected on goodness of fit; current fitted parameters are $\alpha \approx 476$ GWh/day, $\beta \approx 130$ GWh/HDD, base 16.5 °C, $R^2 \approx 0.92$, residual standard error 170–180 GWh/day (Figure 1).

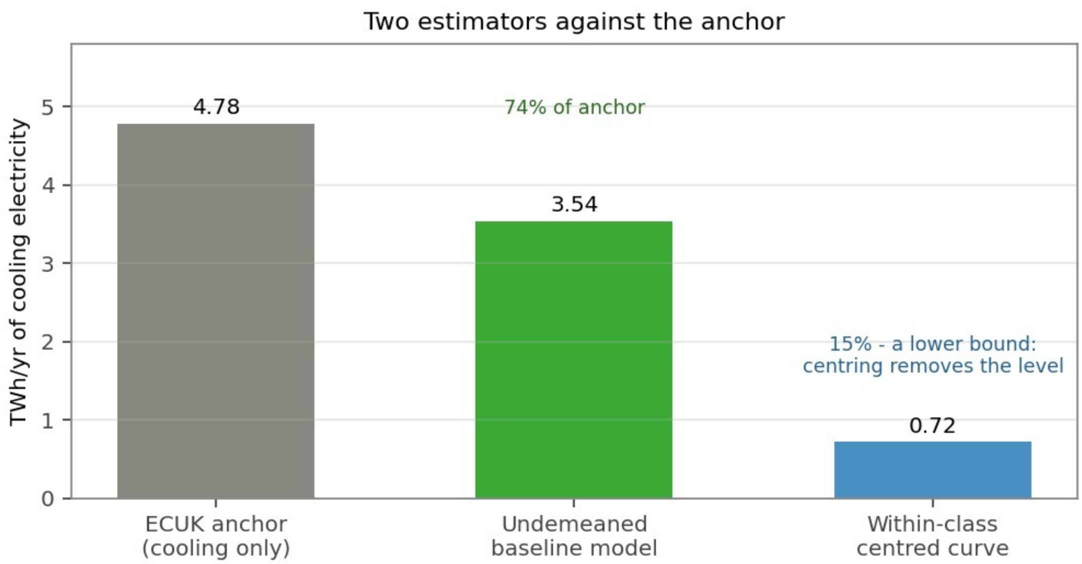

Figure 1

Figure 1. Segmented regression of daily gas demand on temperature. The slope is the marginal space-heating response; the intercept is the weather-independent baseline of hot water, cooking and non-heating gas.

Two constants govern the window and the year, and they are deliberately separate: `WINDOW_DAYS` = 730 for the regression, `ANNUAL_DAYS` = 365 for every trailing-twelve-month sum. A plausibility gate fails the build if trailing-twelve-month heating degree days fall outside 1,200–3,200.

2.4 On the coefficient of determination

$R^2 \approx 0.92$ warrants two caveats. Because HDD and demand share a strong annual cycle, a high R^2 is expected and does not validate the functional form. And R^2 is inflated by the seasonal range: the residual standard error is the more informative figure for daily uncertainty, and is published in the calibration diagnostics.

Notes for investigators

The single fitted base conceals a distribution. The segmented form imposes one balance temperature on the entire building stock. Fitting the same limb with balance points distributed as $T_b \sim N(\mu, \sigma)$ — for which the fleet response has the closed form $E[\max(0, T_b - T)] = \sigma \cdot \varphi(z) + (\mu - T) \cdot \Phi(z)$, $z = (\mu - T)/\sigma$, and which collapses exactly to the kinked model at $\sigma = 0$ — gives **μ = 15.5 °C, σ = 2.00 °C** on 612 days. The improvement is significant: at the observed residual of 169 GWh/day, a true hard kink returned $\sigma \geq 2.0$ in none of 60 simulations and an RSS gain above 0.73% in none of 40.

The consequence is interpretive rather than numerical. **The fitted base of 16.5 °C is roughly the temperature at which 70% of the stock has already stopped heating, not the temperature at which heating ends.** The published estimate continues to use the segmented fit, which is adequate as a shaping device; anyone reading the base as a fleet-average balance point should not.

The estimator is in `analysis/fit_balance_spread.py` with a self-test.

---

3. Calibration to national statistics

The regression yields a self-consistent daily signal but not an absolute national total. The trailing twelve-month integral of the fitted space-heating component is reconciled against ECUK [4], which reports domestic and non-domestic consumption by end use and fuel, derived by DESNZ from the Digest of UK Energy Statistics combined with the Cambridge Housing Model and the National Energy Efficiency Data-Framework.

ECUK space-heating totals are themselves modelled estimates, weather- corrected to a long-run average using the same degree-day apparatus — not metered heat.

The calibration ratio

```
r = (modelled 12-month space heat) / (ECUK anchor, GB-adjusted,
weather-normalised)
```

is monitored continuously, with a publication gate of $|r - 1| \leq 0.10$. The current value is **r ≈ 1.10** — at the edge of tolerance, and reported openly on the site rather than silently rescaled. Candidate contributions to the residual include the treatment of non-domestic gas, the fixed base temperature, and weather-normalisation differences between the trailing window and the ECUK reference year.

The UK-to-GB scaling applied to every ECUK annual is described in Section 5.5.

---

4. Cooling: the equipped fleet (observed)

4.1 Data source and reconstruction

Cooling delivered by installed plant is inferred from national electricity demand. Half-hourly demand comes from National Energy System Operator (NESO) historic demand data and rolling update files [16].

Metered demand understates true underlying demand by the quantity supplied behind the meter, principally embedded solar photovoltaics — which is largest precisely on the hot, bright days when cooling peaks. **Underlying demand is therefore reconstructed by adding NESO estimates of embedded solar and wind to metered demand before analysis.** The correction is material: embedded solar alone reaches 10–16 GW at midday in high summer against a metered demand of 20–30 GW. Future-dated rows and days with other than 48 half-hourly values are removed.

The stored series is already underlying. `demand_GW` in the retrospective store is NESO national demand plus embedded wind *and* embedded solar. Adding embedded generation to it a second time double-counts, and does so in a way that manufactures a temperature signal in daylight hours (see *Notes*).

4.2 Within-class centring

Summer electricity demand carries strong non-cooling structure: it falls during the August holiday period and differs between weekdays and weekends. Binning raw demand against CDD conflates that structure with the cooling response and can yield a spurious negative uplift.

Each daily underlying-demand value d_i is therefore expressed as an anomaly about the mean of its (calendar-month × weekend/weekday) class:

$$a_i = d_i - \bar{d}(m_i, w_i)$$

This is the within transformation of standard panel econometrics — two-way fixed-effects demeaning on calendar month and day type [31] — applied here to isolate a weather signal rather than to estimate a panel coefficient. The term “demeaning” is avoided in public copy in favour of “within-class centring”, to prevent confusion with the ordinary meaning of the word.

The estimator was validated before deployment on synthetic data combining a known cooling slope with an injected holiday depression of comparable magnitude to real August demand (Figure 3). Naive binning reproduced the spurious negative-uplift artefact; within-class centring recovered the injected slope to within sampling error.

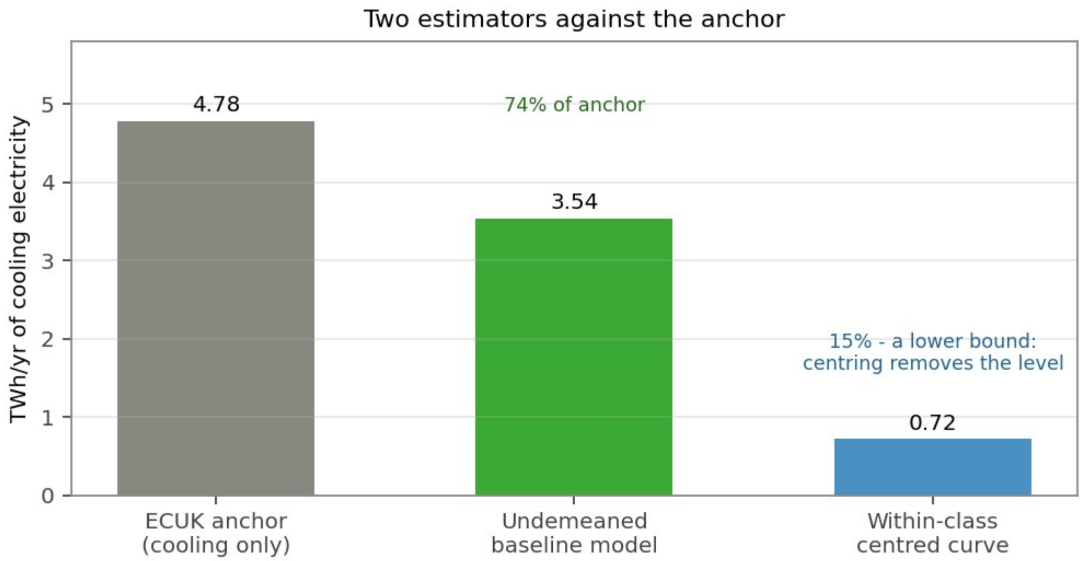

Figure 3

Figure 3. Estimator validation. A synthetic demand series was constructed with a known cooling slope and a holiday depression deliberately injected. Within-class centring recovers the injected slope.

4.3 Response curve

The anomalies are grouped into CDD bins and the response is fitted continuously as

$$E(c) = k \cdot c^2$$

where c is cooling degree days at the base of Section 6.1. The quadratic form is not assumed: it is what an undemeaned regression independently selects (Section 4.5), with an exponent of 2.0 to two significant figures and R^2 0.96 against the bin means. k is published in the payload as `k_quadratic` so the fit can be reproduced, and the **same fit is used for the daily series and the weekly headline** — one estimator, one implementation.

The response is convex, not saturating. It steps about +3 GWh per CDD across the cool bins against +17 across the warm ones. A linear extrapolation from the shallow end therefore sits *below* the observed curve at high CDD.

4.4 What this measures, and what it cannot

The observed response is a lower bound by construction, not an independent estimate of the total. Within-class centring removes the level along with the confound: cooling that runs steadily through July is contained in the July mean and is invisible to the centred estimator. Only day-to-day variation about that mean survives.

Quantitatively, the published Tier 1 recovers about **a third** of the ECUK cooling anchor — 4,681 GWh thermal on the trailing year against 14,352.

This also resolves the tier arithmetic. The measured tier is a **subset** of the anchor, being the weather-varying part of the same equipped stock's consumption. The tiers that lie *outside* the anchor are those the anchor cannot see — buildings that want cooling and do not have it (Section 5), and process refrigeration (Section 4.6) — because a consumption survey cannot count a building that consumes nothing, nor an end use it does not cover.

4.5 The undemeaned reconciliation

Because the centred estimator cannot recover a level, a second estimator is run offline. Daily GB demand, with embedded generation added back once and no centring applied, is regressed on

$$D = \beta_0 + \beta_1 \cdot \text{weekend} + \beta_2 \cdot \text{HDD}(15.5) + \beta_3 \cdot \max(0, \bar{T} - T\_c)^2 + \beta_4 \cdot \text{solar} + \Sigma\, \beta_j \cdot \text{holiday}_j$$

with T_c scanned over 11–20 °C. Because the baseline is modelled rather than removed, the intercept survives and T_c is identified.

On the full 800-day record the base specification gives T_c = 12.0 °C and 3.59 TWh — **75% of the anchor** — and the full specification 11.0 °C and 4.56 TWh, **95%**.

Embedded solar is carried as a control and is the single strongest: it enters at −0.91 with an incremental R^2 of 0.041, and a negative coefficient is what embedded generation displacing metered demand should produce. Calendar terms carry real weight — bank holidays ΔR^2 0.031, the Christmas shutdown 0.028.

Nine regressors reaching 95% is not a validation. Each term is a documented control with a physical rationale rather than a fitted nuisance parameter, and a mis-specified version of the same model returned 164% of the anchor — which is what overfitting looks like. The honest reading is that the undemeaned estimator recovers most of the anchor and is no longer obviously biased.

This estimator is **not** used in the published bill or carbon figures, which remain ECUK-anchored. It is drawn on the site as a comparison line over the Tier 1 bars, with its coefficient pinned to the fit that produced it (`UNDEMEANED_FIT`: k = 0.6109 GWh per °C^2 above 11.0, fitted 2026-08-17, 800 days, nine regressors, R^2 0.918). A coefficient without its vintage is not reproducible, so the fit date, record and resulting share are stored alongside it and updated together.

4.6 Tier 0: process cooling

Process cooling — the food cold chain, data centres and industrial refrigeration — sits **outside** ECUK End Use Table U5, which is a services-buildings table covering comfort cooling and ventilation. It is named on the panel and deliberately not plotted, being several times the other tiers combined.

Components, as electricity:

Component	TWh elec	Source	Data vintage
Retail refrigeration	6.91	Foster, Brown & Evans [32]	2019 (from 2011–16 studies)

Component	TWh elec	Source	Data vintage
Food service / catering	4.16–9.28	Foster, Brown & Evans	2010–2019
Food & drink manufacturing	3.08	Foster, Brown & Evans, from DUKES	2019
Cold stores	1.32	Foster, Brown & Evans; CCF ~1.35	2017–2019
Agriculture & fisheries	0.40	Foster, Brown & Evans	2019
Chemicals & pharmaceuticals	0.88	TICR, IOR/LSBU [33]	2023–24
Data centre **cooling**	1.0–1.5	Derived: DESNZ 4.5 TWh total [34] at PUE 1.5–1.8	2024
GB non-domestic total	**17–22**		
Domestic fridge-freezers — excluded	*10.75*	*Foster, Brown & Evans*	*2019*

The headline figure in the literature — 29.1 TWh for the entire UK food cold chain — includes household fridges, which belong in the residential tier.

As delivered service, converting component by component at its own efficiency (supermarket display ~2–2.5, cold stores and manufacturing ~2.5–3, data centre cooling service approximately the IT load it removes), the total is **25–45 TWh**.

The site reports this as a service figure of order 26 TWh, at the low end of that range. It is not an electricity figure: as electricity the total is 17–22 TWh.

The scale is the substantive point. At 17–22 TWh of electricity, process cooling is **three and a half to four and a half times the entire comfort-cooling anchor** of 4.78 TWh. Britain uses far more electricity keeping product cold than keeping people cool, and none of it is in ECUK's services cooling line.

One overlap risk (†). Data centres are services-sector buildings and their chillers may be counted within BEES's space-cooling line. That is 1–1.5 TWh — about a fifth of the entire Tier 1 anchor — potentially counted in both places. It is not resolved.

This is the stalest evidence on the site. The cold-chain anchor rests on 2019 data, retail refrigeration on studies from 2007–2016, and cold-store guidance was frozen at 1994 values until 2019. Only the data-centre figures are current. The food-service range of 4.16–9.28 TWh is the widest single uncertainty. A conflicting trade-press figure putting UK data centres at 5.8% of national electricity — roughly 15 TWh — is a capacity-based estimate on a different scope; the metered DESNZ figure is used.

Notes for investigators

A double-counted variable cost three published claims. Embedded solar was added back to a series that already contained it, in three separate analyses. Because embedded solar correlates with temperature at r = +0.48 in summer daylight hours, this manufactured a spurious temperature signal in exactly the hours where a cooling signal would appear. It invalidated a diurnal comfort/process separation, a reconciliation figure, and a humidity null. It did **not** affect anything measured at night, where solar is zero. Three defences are worth building in from the start: write into the code what a series contains at the point it is loaded; test any null result against the direction of its own bias; and whenever a control is excluded for having "the wrong sign", check the sign is not an artefact of the dependent variable.

Humidity is a null, obtained on the fourth attempt. Latent load rises with absolute moisture, so moist-air enthalpy was a plausible explanation for the convexity. On the full 800-day record with solar controlled, **dewpoint adds 0.0006 to R^2 and enthalpy days add 0.0000**, and adding enthalpy moves the quadratic coefficient by **+3%** — where a latent-load explanation requires moisture to absorb a material part of it. Three earlier attempts, on contaminated or partial data, moved it +7%, −25% and +27%. The value of the final null lies in the test having demonstrably been capable of failing.

The convexity is the extensive margin — each building is a step function and the fleet is the integral of those steps — on three independent grounds: the response steepens rather than saturating; dividing out the falling COP leaves the bend almost intact; and moisture absorbs almost none of it.

The cooling floor: five routes tried. Recovering the level that centring removes was attempted five ways. Two are dead and are recorded so others need not repeat them.

Route 1, model the baseline (Section 4.5) recovers 75–95% but cannot separate weather-driven from year-round cooling, because a constant load hides in the intercept.

Route 2, separate comfort from process by hour of day. **Dead** — it was the solar artefact. Corrected, the diurnal profile is flat: 0.26 GW/°C at 04:00 against 0.24 at noon, a ratio of 0.93 where 2.59 was claimed.

Route 3, fit a balance point as the vertex of a V. **Dead as an identification method.** Fitting heating and cooling limbs meeting at a vertex is a better functional form than flat-then-steep, and correctly explains the negative low-temperature slopes as a genuine heating limb. But the vertex is not identified: regressed on each hour's own mean temperature it returns a slope of 0.98 with R^2 0.97. **The V-fit places its kink about 1.5 °C above the centre of whatever data it is given.** The general lesson is worth stating: if a fitted parameter correlates with the mean of the data used to fit it, it is not identified — a check that takes two minutes and should be run on any threshold before it is believed.

Route 4, borrow a level from the gas side (Section 2, *Notes*) establishes that balance points are spread across the stock, but yields a switch-on shape without a scale: the cooled stock is a different population from the heated stock, and leaving heating does not imply entering cooling.

Route 5, weekday versus weekend at matched temperature. **Open, and the only route to produce a level.** Comparing weekday with weekend demand at the same temperature, hour and season cancels everything invariant to occupancy — baseload, process refrigeration, the weather itself — leaving load that responds to both temperature and occupancy. On summer afternoons the gap grows **+0.107 GW/°C** (R^2 0.90, six bins), from 2.64 GW at 13.5 °C to 4.27 at 27. Integrated above a 14 °C reference: **467 GWh/yr in the 12–17h band, 1,866 scaled across the day — 10% to 39% of the anchor.**

The placebo is what makes it credible. Repeating the contrast at night in the same season switches the occupancy difference off while leaving the weather unchanged: the night slope is **−0.065 GW/°C** against the day's +0.107. The gap grows with temperature only when buildings are occupied, which a generic weekday/weekend artefact would not produce. The night gap of 0.94 GW is temperature-insensitive — a standing occupancy load that is not cooling.

Three limits: buildings open at weekends appear on both sides and cancel, so this is a lower bound on commercial cooling; ventilation and other occupancy-linked load that tracks temperature are not separated from cooling; and the night placebo rests on three bins over four degrees, because Britain has few hot nights. `analysis/fit_weekend_contrast.py` fails rather than reporting a figure if the placebo does not hold.

Nothing from routes 4 or 5 is published on the site.

---

5. Cooling: the comfort deficit (modelled)

5.1 Threshold and basis

Overheating exposure in buildings with no cooling equipment is quantified as degree-hours above **20 °C on the daily mean of population-weighted outdoor air**, derived from the 2017 Energy Follow-Up Survey indoor/outdoor regression [8a], in which indoor bedroom temperature rises approximately 0.48 °C per 1 °C of outdoor temperature:

$$CDH_{20} = \Sigma_d \, 24 \cdot \max(0, \bar{T}_d - 20)$$

This is not the 15.5 °C cooling shape base of Section 6.1, and answers a different question. 15.5 °C is the balance point at which comfort cooling plant begins to run in buildings that have it; 20 °C is the outdoor daily mean above which buildings that have none become uncomfortable. Both appear on the site and the distinction is stated wherever they do.

A night-time urban heat island uplift is not in the numbers and would add an estimated 10–20% (†).

5.2 Stock and response

Three prevalence scenarios are carried and never averaged: a low case of 11% of dwellings from English Housing Survey self-reported overheating [9]; a central case of 25% (†); and a high case of over one-half, from the Climate Change Committee assessment of dwellings at overheating risk [10]. Non-domestic uncooled comfort space is taken as 190 Mm^2 (†) from BEES floor areas [18].

Thermal response is 0.2 kWh thermal per degree-hour per dwelling (†) — equivalent to an effective cooled-zone UA of about 200 W/K — and 3.0 Wh thermal per m^2 per degree-hour for non-domestic space (†). These are the dominant uncertainties in this tier.

Unit convention. This tier is a **thermal load computed directly from building physics and never divided by an efficiency**. Tier 1 by contrast is measured electricity multiplied *up* by the fleet EER. The two are reported in the same units and rest on entirely different assumptions; the site states this in the chart legend.

5.3 Tier structure

Tier	Content	Relation to the ECUK anchor
1	Weather-varying response of the equipped fleet	**Inside** it — about a third (Section 4.4)
2	Buildings with no cooling equipment at all	**Outside** it — a consumption survey cannot count non-consumption
0	Process cooling: food cold chain, data centres, industrial refrigeration	**Outside** ECUK U5 entirely — 17–22 TWh electricity, 25–45 TWh service (†), Section 4.6

5.4 Delivery efficiency

The comparison route efficiencies are an effective COP of order 20 for passive ground cooling and a seasonal EER of **3.0, range 2.5–4.0**, for air-source.

The range matters. It is the largest single unverified assumption in the cooling chain and propagates one-for-one into every cooling service figure on the site. Legacy oversized plant runs nearer 2; modern commissioned VRF nearer 4.5. BEES excludes fan energy from its space-cooling line, so the boundary here is nearer chiller-plus-condenser than whole-system, which argues for the upper half of the range; UK part-load operation argues for the lower half.

What would resolve it: air-conditioning inspection reports under EPBD carry system type, age and rated efficiency for a large sample of UK non-domestic plant, and are the one dataset capable of producing a fleet-weighted seasonal figure rather than a judgement.

5.5 The UK-to-GB boundary

ECUK reports UK totals; the site reports Great Britain, because the live gas signal comes from the GB distribution zones and Northern Ireland is not on that system. Every annual therefore passes through a UK-to-GB scaling step, **per fuel**:

Fuel	GB share	Basis
Gas	0.988	NISRA 36% of NI homes on gas; GB gas-connected ~23m (†)
Oil	**0.693 → 0.830**	NIHE/BRE reconciled to a trailing basis; NI ~17% of UK

Fuel	GB share	Basis
		oil heat (†)
Solid	0.893	Low as an NI primary method, higher as secondary (†)
All others	0.972	No NI-specific driver, so the household share (†)

Oil is the case that matters. NISRA's Continuous Household Survey 2024/25 reports 61% of NI households on oil central heating as their primary method against roughly 5% in Great Britain, so NI holds a far larger share of UK oil heat than its 2.8% of households implies. Two independent routes agree once put on the same basis: NIHE and BRE [28a] give 10.8 TWh of NI oil heat on a 2016 basis (21% of the UK's 51.0 TWh), and the Irish Heat Split sibling estimates 8.3 TWh on a trailing twelve months (16%). The difference is nine years in the jurisdiction where oil is retreating fastest — the oil share of NI homes fell from 68% to 61% over that period — and 10.8 × 0.90 × 0.93 × 0.97 = 8.8 against 8.3, a 5% residual.

Because the site reports a trailing window, the current-basis figure applies.

All four factors are estimates, and two known biases oppose: the NI oil share is still falling, arguing for a higher GB share; and NIHE covers domestic only, so NI non-domestic oil sits on top, arguing for a lower one.

Notes for investigators

A gas-network factor does not transfer to other fuels. Until August 2026 a single 0.985 — a statement about the GB gas grid — was applied to every fuel. For oil this overstated GB heat by about 19%, because NI's oil penetration is twelve times Britain's. Any UK-to-GB or GB-to-UK scaling should be built fuel by fuel from stock data, not inherited from whichever fuel prompted it.

And a first correction can itself overshoot. The initial per-fuel oil share was derived from household counts alone — 61% of 0.80m NI homes against ~1.1m GB oil homes — giving 0.693. That ignores consumption per household: NI dwellings are smaller and newer than the GB off-grid stock, which skews rural and large. Reconciling against two consumption-based sources moved it to 0.830.

6. Conversion, cost and carbon

6.1 In-situ efficiencies and performance factors

Combustion appliances are derated by in-situ, not nameplate, efficiencies, and space heating and hot water are converted at different ones:

Fuel	Space heating	Hot water	Basis
Gas	0.835	**0.73**	In-situ winter [5]; hot water derived below

Fuel	Space heating	Hot water	Basis
Oil	0.82	**0.71**	As above
Bioenergy	0.70	0.70	Single figure — not covered by SAP Table 4b
Solid	0.55	0.55	As above
Heat networks	1.0	1.0	ECUK "heat" is delivered heat
Resistive electric	1.0	1.0	

Why the two differ. A boiler serving only a cylinder fires in short bursts, cools between them, and returns water above the condensing dewpoint — losing its latent recovery outright. SAP has carried separate winter and summer seasonal efficiencies for this reason since SAP 2009.

The hot-water figures are the mean SAP 2012 Table 4b [5a] summer/winter ratio applied to the in-situ winter anchor, so the sourced winter values are untouched: gas ratios 64/74, 74/84, 65/74, 75/84 → mean 0.879 × 0.835 = 0.73; oil 68/80, 72/84, 73/82 → 0.866 × 0.820 = 0.71. Summer runs 82–89% of winter across the whole table.

The summer figure applies to hot water in every month, not seasonally. SAP is explicit that $\eta_{water} = \eta_{summer}$ for all twelve months: the distinction is between end uses, not between seasons.

Two caveats. Table 4b is SAP's *fallback* for boilers absent from the Product Characteristics Database, so it is conservative and skewed to older plant; a PCDB-weighted fleet figure would be better. And Table 4c deducts five points where a regular boiler has no interlock, which is not applied here.

The effect is that the effective gas efficiency moves with the season — about 0.81 in January and **0.74 in July**, when the load is almost entirely hot water.

Heat pump seasonal performance factors are 2.80 for air-source, from the Energy Systems Catapult Electrification of Heat demonstration [7]; 3.24 for ground-source, from Energy Systems Catapult in-situ monitoring of more than 1,700 installations [6]; and 5.0 for a networked geothermal ambient loop (†), for which no monitored UK fleet exists.

All three are whole-system figures on the same basis as the SEPEMO H4 boundary the two monitored anchors come from: electricity metered at the supply point over heat delivered at the building boundary. For the network route that places inside the boundary the building heat pumps and auxiliaries, district and primary circulation, well pumping where a doublet is used, controls, the round-trip loss on the seasonal store, and — because the denominator is heat delivered at the building — distribution heat losses. Outside it: groundwater treatment where required (not modelled, and would reduce the figure), customer-side emitters, and anything embodied.

This matters because a monitored H4 SPF compared against a plant-only network COP would flatter the network throughout. **The parasitic load is the part most likely to be optimistic.** European ambient-loop and aquifer-storage schemes commonly report 5–15% of delivered heat going to circulation; a plant-level 5.6 becomes a whole-system 5.0 at about 10%, which is the arithmetic this figure assumes. At 15% it would be 4.87 and at 5% 5.33.

These annual factors anchor the model; they are not applied as constants in the hourly or daily engines. See Section 10.3.

6.2 The ECUK cooling and ventilation split

ECUK End Use Table U5 publishes services-sector **"cooling and ventilation" as a single electricity line**, held here at 10.4 TWh/yr. The Building Energy Efficiency Survey [18] — from which ECUK's services end-use split derives — reports the components separately: **fans 6,030 GWh, space cooling 5,090 GWh and humidification 130 GWh**, i.e. **54% ventilation to 46% cooling**. US CBECS 2018, which separates the two natively, gives 18% of commercial electricity to ventilation against 14% to cooling — a 56:44 ratio from a considerably hotter climate, which corroborates the shape.

Ventilation is reported as its own service at unity and excluded from the cooling figure (Section 1.3). It is excluded from the heat figure too: it is neither. It carries its own bar segment, its own line in the bill and carbon tables, and is **carried through the geothermal what-if unchanged**, because a network does not displace fans.

The cooling shape base is 15.5 °C on the daily mean — the balance point at which comfort cooling begins in buildings with high internal gains, the traditional UK degree-day base, and what CIBSE TM46 applies across all sectors. It is corroborated by the undemeaned regression of Section 4.5, which places the best-fit balance point at 11–12 °C. The Met Office cooling-degree-day base of 22 °C daily mean is a climate-comparison convention, not a building balance point, and is not used.

6.3 Cost

Retail costs are priced by sector. Each fuel component is split into its domestic and services shares using the same ECUK U3/U5 tables that anchor the mix; the domestic share is priced at the prevailing Ofgem default-tariff cap unit rates [12], refreshed quarterly, and the services share at non-domestic rates anchored on DESNZ Quarterly Energy Prices [30] (Causeway estimates †, set within the published size-band range). Cooling and ventilation, both ECUK services-only categories, are priced wholly at the non-domestic electricity rate. Standing charges are excluded throughout.

The domestic share differs by END USE, not only by fuel: gas space heating is 73.5% domestic against gas hot water at 85.9% and oil at 53.3%. A single blended share per fuel would misprice both.

The daily spark gap uses wholesale commodity prices only — National Gas System Average Price [13] and the Elexon market index [20] — divided by the respective efficiency or performance factor, and excludes network, policy and supply costs.

6.4 Carbon

Combustion emissions use DESNZ (Defra) greenhouse-gas conversion factors [11] on a gross calorific value basis consistent with UK metered gas billing: natural gas 0.18296 $kgCO_2e/kWh$

(2025), kerosene ~0.247, coal ~0.345. Electric routes use live GB grid carbon intensity from the NESO Carbon Intensity API [21] as a seven-day mean. Bioenergy is counted at zero combustion emissions under the biogenic convention, excluding supply-chain emissions — a stated simplification.

Because combustion factors are fixed by fuel chemistry while the electric factor tracks the decarbonising grid, the heat-pump rows fall over time and the combustion rows do not; the panel is designed to make this divergence legible.

6.5 The heat-price series

A daily series of the commodity cost of delivered heat is published for four routes, on **two bases** and for **three services** — six combinations.

Wholesale. Gas is the National Gas System Average Price divided by the boiler efficiency for the service selected. The three electric routes are the daily electricity market index divided by **each route's own COP at that day's temperature** (Section 10.3), not by a fixed seasonal factor.

At the meter. The same four routes at the Ofgem default-tariff cap in force on each date — a step function of policy decisions rather than a market series. The geothermal network route is priced at the **non-domestic** electricity rate, because a network operator buys on a commercial contract; what is compared on that line is the operator's input cost, not a connected household's bill, which is a network tariff this model does not carry.

Three services:

Service	Flow temperature	Gas efficiency
Space heating	Weather-compensated, 50 °C at −5 falling to 30 °C at +15	0.835
Hot water	**52 °C year-round** (†) — a cylinder needs its temperature whatever the weather	0.73
As delivered	Both, blended by the mix served that day	Both, blended

The blend is a harmonic mean weighted by delivered service, because what is averaged is a cost per unit delivered:

```
effective COP = 1 / (w_space/COP_space + w_dhw/COP_dhw)
```

with w_space from the same HDD-shaped, summer-damped daily weighting the hourly engine uses (Section 10.2) and hot water flat. An arithmetic blend would flatter whichever service is dearer. This is the only view that moves with the load, tracking space heating through the winter and hot water through the summer.

A service-volume band is drawn beneath the chart, on the same window and x axis: delivered GWh per day, space heating stacked with hot water, the selected service solid and the other faded. A price is only as interesting as the service it prices — by July the load is entirely hot water, so a space-heating price that month describes something nobody is buying.

Two electricity means are published per day: the flat mean, which the headline spark-gap ticker uses, and the **heat-weighted mean**, which is what a heat pump would actually pay because its load concentrates on the morning and evening peaks. The difference is the coincidence premium (Section 10.7), approximately 3% across a winter.

Notes for investigators

Hot water sets an honest limit on the heat-pump case. Because the cylinder pins the flow temperature, none of the flow-reduction benefit of Section 10.9 is available to it and the lift never shrinks. The geothermal network's advantage over a gas boiler is about **1.7× in both January and July** on hot water, against 1.9× rising to 4.4× on space heating. The case for heat pumps is strongest exactly where the flow temperature can fall, and weakest on the load that pins it.

Any average of a price series taken over days rather than over heat delivered will flatter whichever route performs best in mild weather. Unweighted, air-source shows a higher mean COP than ground source (4.88 against 4.06) and inverts the site's own SPF ordering; weighted by heat, the ordering is correct (3.49 against 3.66). On the thirty highest-heat days the figures are 2.21 and 3.08. Both weighted and unweighted values are published so the distinction is checkable.

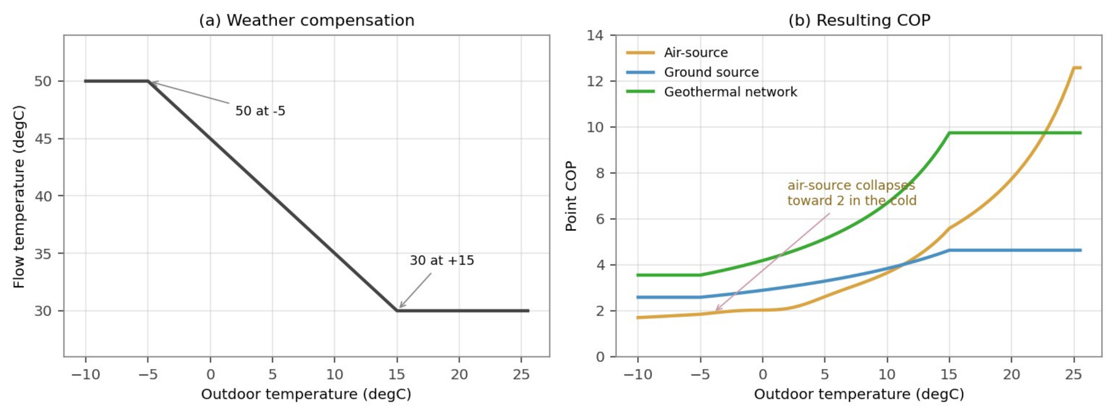

Figure 5

Figure 5. (a) The weather-compensation curve and (b) the point COP it produces for each route. Air-source falls toward 2 in the cold, exactly when heat is most wanted; a ground loop barely moves.

---

7. Indigenous-supply share

The headline indigenous share is computed on a delivered-service basis: each unit of delivered heat or cooling inherits the UK-origin fraction of its energy input — gas ~42% (UK-origin share of the gross supply pool: UKCS production plus grid biomethane over production, imports and stock draw, DUKES 2026 [22]), grid electricity ~75% (†), harvested ambient and ground heat 100%. The last follows the Eurostat and DUKES convention of treating captured ambient heat as indigenous renewable supply [22,23].

Cooling's delivered multiple inherits the origin share of its input electricity; thermodynamic leverage is not treated as an energy origin. Ventilation electricity carries an origin share but is part of neither service.

The purchased-energy-basis share is retained for continuity; it changes little under a geothermal shift because ground heat is never purchased, which is the substantive point the services basis is designed to expose.

---

8. Other panels

8.1 Non-gas fuels

Oil, bioenergy, solid fuel, heat networks and electric resistive heating lack daily operational feeds and are held at ECUK annual levels [4], modulated by degree days. They move slowly and contribute little week-to-week variance. Space and water heating are shaped separately — space on degree days, water flat — and converted at different efficiencies (Section 6.1). A fuel that does both jobs therefore appears twice, on two different curves, and is summed only for display.

8.2 Geothermal

The present geothermal contribution is anchored on the EGEC 2025 UK Country Update [24]: approximately 1.43 TWh/yr from an installed base of order 55,000 ground-source heat-pump units, plus of order 0.07 TWh/yr from deep, mine-water and open-loop district schemes. The stated range is 1.4–2.0 TWh/yr (†); no national register or metering of ground-source output exists, and the Environment Agency's operational-unit estimate is lower than the installed-unit figure, so this is the most weakly constrained quantity on the site.

Forward bars for 2027 and 2031 are trend and scenario projections (†) informed by the EGEC Geothermal Market Report 2025 [25] and the Climate Change Committee Seventh Carbon Budget pathway [26]; the 2050 bar is the Project InnerSpace/REA/ARUP ambition [27].

Two archetypes are modelled, corresponding to the BGS classification: *ground source* is shallow geothermal — open-loop doublets, closed-loop borehole arrays, mine water — modelled as heat-only or heat-dominated at an 8 °C source. *Geothermal network* blends shallow underground thermal energy storage, where summer heat rejected by cooling and waste heat from data centres and other sources is banked in the ground, with intermediate geothermal where the geology suits it. The difference between the two is mostly what the loop is sitting at: 19.6 °C for a charged network loop against 8 °C for a passive array.

8.3 The empty bar: installed hardware vs the what-if

The 20% what-if is also priced in hardware. Serving one-fifth of delivered buildings heat at 2,000 equivalent full-load hours (†) requires installed thermal capacity of order 30–40 GWth;

operating UK ground-source capacity is 861 MWth, with deeper geothermal adding of order 10 MWth (†).

The panel sets these against France (2,293 MWth ground source + 724 deep), the Netherlands (2,486 + 367) and Sweden (8,120 + 47), from the EGC 2025 country update. The per-person framing is the substantive point: Sweden operates roughly 778 installed thermal watts per person against the UK's ~13 (†). A companion calibration restates each fleet as a share of its own country's buildings heat: the UK ~0.5%, France ~2%, the Netherlands ~5%, Sweden ~20% — **Sweden's existing fleet already serves approximately the share the UK what-if proposes.**

For capacity flow, the UK replaces of order 1.6 million boilers a year — roughly 35–40 GW of combustion capacity installed annually (†) — so the what-if corresponds to about one year of the existing replacement flow.

8.4 Northern Ireland

Northern Ireland operates separate gas and electricity systems and the all-island Single Electricity Market, and is not covered by the GB live feeds. Its heat is estimated annually at **14.0 TWh** and modulated by NI degree days. That estimate is corroborated by NIHE and BRE [28a], who model NI domestic heat at 14.2 TWh — 10,700 GWh of primary space heating plus 3,500 GWh of water heating, from the House Condition Survey — within 2% of the carried figure.

Oil heats just over 60% of NI dwellings and has no meter or daily feed; the resulting data gap is itself reported. The same source underpins the per-fuel GB oil share of Section 5.5.

Northern Ireland is covered properly by the companion site, Irish Heat Split, which applies the same estimator family to the island of Ireland on its own feeds. The two share constants where the underlying quantity is shared and diverge where the systems do. **Cooling is deliberately not compared across the two**: the island figure counts the cold economy (data centres, refrigeration, process), the GB figure counts ECUK comfort cooling.

8.5 Whole-economy context (WHY HEAT?)

Four annual pies place heat beside transport and non-heat electricity across delivered services, spend, imported energy and emissions. Sourced anchors are ECUK 2025 final-energy and transport totals [4] and DESNZ provisional greenhouse-gas statistics by sector [29]; the service-level allocations (heat total ~630 TWh, spend, import and emission splits) are author derivations (†), deliberately rounded. Non-energy emissions are excluded; the view is of energy services, not a national inventory.

---

9. Historical series (the trend layer)

9.1 Construction

The four headline estimates — purchased energy, indigenous share, national bill and emissions — and their 20%-geothermal what-if counterparts are published as a weekly historical series, rendered as a sparkline under each headline with a selectable one-week to twenty-four-month window.

The series is built from calendar weeks (Monday–Sunday); a week enters only if the gas feed served all seven of its days, so every published point is live by construction. Each point is computed by the same estimator code path that produces the current-week headline figures — one implementation applied to a different week, rather than a parallel reconstruction that could drift. The series is capped at 120 weekly entries, about twenty-seven months.

Entries are frozen as first published, except the two most recent weeks, which are recomputed each day against upstream feed revisions.

9.2 Restatement

Two markers govern restatement and the distinction matters. `HISTORY_SCHEMA` tracks new FIELDS; `ANCHOR_EPOCH` tracks a changed BASIS — a constant that alters what a past week WOULD have computed. Bumping either restates every stored week the gas feed still serves, reusing each week's stored grid intensity and the tariff in force, so no historical inputs are refetched.

The markers are written to the output **only after the restatement completes**, so a failure is retried rather than silently skipped, and a `FORCE_RESTATE` override exists. Every run prints the schema and epoch transition, the counts of stored, complete and computed weeks, and how many were restated.

9.3 Time-consistent pricing and carbon

Two inputs vary materially within the window and are applied at their historical, not current, values. Monetary figures use the Ofgem cap unit rates in force during that week, from a maintained table of cap periods. Emissions use that week's mean GB grid carbon intensity from the half-hourly NESO actuals [21].

Estimator constants for which no weekly series exists — non-cap fuel prices, non-domestic rates and sector shares, the ECUK anchor vintage, and the indigenous input-origin shares — are today's values applied uniformly across the window. This bounds the interpretation of the bill and indigenous-share histories more than the energy and emissions histories, and means the earliest weeks are not exact reproductions of what the dashboard published at the time.

9.4 The live/modelled boundary

The National Gas feed serves approximately the trailing twenty-five months of daily LDZ demand. The feed's actual served window is recorded at every build in the published data, so the boundary is empirical rather than assumed. The trend layer publishes nothing beyond it.

A modelled back-cast — the same regression applied to archived degree days and calibrated to the ECUK vintages then current — remains a planned second phase and is not built. Live and modelled points will never be blended undistinguished.

Notes for investigators

A 52-week sum of weekly shares is not the annual total. Each week's non-gas components are an annual total multiplied by that week's share of *its own* trailing-twelve-month degree days. Those shares sum to 1.0 only if the denominator is constant across the window, and it is not: in a warming or cooling year the earlier weeks are divided by a different year's total. The trailing-year cooling sum currently runs about 17% above the annual anchor, 2026 having been much hotter than 2025. This is stated rather than corrected, because normalising the sum onto the anchor would erase a real signal — a hotter year *should* show more cooling.

---

10. The hourly retrospective and the system view

10.1 Purpose and construction

The panels of Sections 1–9 describe the transition in weekly aggregate. This section tests it against the electricity system Britain actually operated, hour by hour. A separate engine reconstructs an hourly retrospective: temperature (ERA5, the same twelve population-weighted points), wholesale price (Elexon market index), grid carbon intensity (NESO), observed underlying demand, and observed wind and solar generation.

The store holds two winters — 800 days, backfilled to June 2024 — so the coldest hour, the binding hour and the tightest week are chosen *across* winters rather than describing the only one on record. Every annual statement is computed on the trailing 8,760 hours.

10.2 Hourly shaping of heat demand

Annual useful space heat is distributed to days in proportion to each day's share of heating degree days at the live regression base, then within each day in proportion to heating degree-hours; hot water is a flat hourly rate.

Two conventions govern the result. Days below a threshold of one heating degree day contribute no space heat, with a smooth ramp to full contribution at three (†): raw degree-hour allocation assigns winter-strength heat to cool summer nights, which is behaviourally implausible and inconsistent with the weekly estimator. And the weights are normalised *after* damping, so the trailing year still integrates exactly to the annual anchor — the damping redistributes heat into the heating season rather than discarding it.

The resulting shape is modelled, not metered, and is marked as such wherever it is drawn.

10.3 Coefficient of performance and calibration

For a heat pump lifting heat from a source at T_s to a delivery temperature T_f, both in kelvin, the thermodynamic ceiling is the Carnot coefficient of performance:

$$COP_{Carnot} = T_f / (T_f - T_s) = T_f / \Delta T$$

Delivery temperature over lift. Real machines achieve a fraction of this:

$$COP = \eta \cdot T_f / \Delta T$$

with a floor of 8 K on ΔT (†), because the ratio diverges as the lift closes and no real machine does.

Route sources:

Route	Source temperature
Air-source	Outdoor air less a 3 K evaporator approach, plus a defrost penalty of up to 12% between −5 and +7 °C
Ground source	**8.0 °C** — 11 °C ground less a 3 K brine approach
Geothermal network	**19.6 °C** (†) — a charged ambient loop, UTES or intermediate-doublet

Flow temperature is weather-compensated linearly between endpoints: 50 °C at T_out = −5 °C falling to 30 °C at +15 °C, clamped outside that range (†). Hot water uses 52 °C year-round.

η is solved, not assumed. For each route it is the value that makes the annual heat-weighted mean of the hourly COP reproduce that route's published seasonal performance factor:

Route	η	Anchor SPF
Air-source	0.33198	2.80
Ground source	0.33646	3.24
Geothermal network	0.33436	5.00

The agreement of the three fractions is not independent evidence. For air-source and ground source both the source temperature and the SPF are sourced, so η is genuinely solved and its value is informative. For the network route neither is: the SPF of 5.0 is an estimate, and the source temperature of 19.6 °C was **chosen so that the network reproduces that SPF at the same Carnot fraction as the other two**. The match is therefore how the source was derived, not a test of it.

What the pair does establish is that SPF 5.0 and a 19.6 °C source are mutually consistent under the same Carnot fraction that fits two monitored fleets — a meaningful plausibility statement, and not a trivial one. It does not establish either quantity independently, and a reader should treat them as one estimate with two faces. The honest presentation is a range: at the calibrated fraction, a network source of 15 °C supports an SPF of about 4.3 and 25 °C about 6.2.

The build refuses to publish if the three fractions drift more than 15% apart. That is a **wiring check** — it catches a route whose calibration has failed to converge — and not a validation.

10.4 Sensitivity to flow temperature

Differentiating the Carnot form with respect to flow temperature, and noting that both numerator and denominator depend on T_f:

$$\partial COP/\partial T_f = \eta \cdot (-T_s) / \Delta T^2$$

The flow temperature cancels from the numerator; only the source survives. As a fractional change:

$$(1/COP)\cdot \partial COP/\partial T\_f = -T\_s / (T\_f \cdot \Delta T)$$

η cancels, so the percentage gain is independent of machine quality, and **ΔT remains in the denominator**, so the gain per kelvin is inversely proportional to the lift. Halving the lift doubles the value of each further kelvin.

At a 45 °C flow: ground source at an 8 °C source has ΔT = 37 K and gains 2.4% per kelvin; the network at 19.6 °C has ΔT = 25.4 K and gains 3.6%. The ratio is simply 37/25.4.

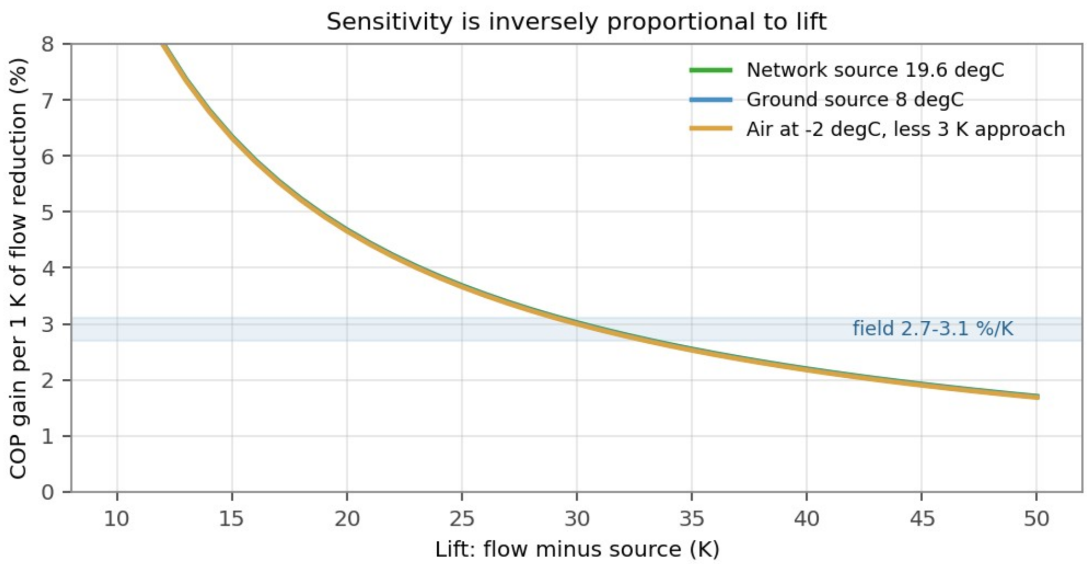

Figure 6

Figure 6. Fractional COP gain per kelvin of flow reduction, against lift. The route with the smallest lift gains most. The shaded band is the HeatpumpMonitor field range.

10.5 Stress, netting and the binding hour

Two hourly statistics are published. The **coldest hour** is the year's highest modelled heat demand, with each route's replacement electricity expressed against that hour's observed demand. It is a gross figure.

The **binding hour** is, for each route, the hour maximising observed demand plus the route's net addition less observed wind and solar, searched across November to March. Net here subtracts the fifth of existing resistive space heating the what-if displaces, shaped like space heat and route-independent.

With a ceiling that varies by winter, the binding hour is the hour of LEAST HEADROOM, not of highest absolute requirement. The two coincide when the store holds a single winter and need not otherwise.

The two statistics need not fall in the same hour, or even the same week, and their ratios differ: the coldest-hour figures are additions alone and separate by the ratio of coefficients of performance, while the binding-hour figures are total system requirements dominated by a common standing demand.

10.6 The capacity line

Requirements are set against a capacity line constructed as a de-rated block of dispatchable and interconnector capacity **plus the wind and solar actually generated in each hour**. Because outturn wind and solar are added hourly, the block must **exclude** de-rated wind, or wind is counted twice.

The block is derived identically for each winter from the NESO Winter Outlook:

```
block = ACS peak demand + de-rated margin − de-rated wind
```

Winter	ACS peak	Margin	Total de-rated	Less wind	**Block**
2024/25	59.8 GW	5.2	65.0	4.4	**60.6 GW**
2025/26	61.0 GW	6.1	67.1	~4.8 †	**62.3 GW**

The 2024/25 wind de-rating is NESO's own (30,532 MW installed de-rated to 4,416 MW). The 2025/26 figure is not published in the summaries and is scaled with installed capacity growth (†).

Each winter is tested against its own fleet.

10.7 What the capacity comparison is not

It is a **historical stress test, not a capacity-adequacy model**. No redispatch, interconnector response, demand-side response or price response is represented; distribution-level constraints are unmodelled; and a capacity expressed in gigawatts cannot represent storage duration — batteries can serve a peak hour but not a fifth consecutive cold day.

Exceedance is reported as hours in which a route's requirement exceeds available headroom under that winter's fleet, never as any statement about supply interruption.

Two winters is better than one and is still a small sample. In 2025-26 the cold snap was windy and the still spell (18–21 March, wind falling to 1.6 GW) was mild, so the coincidence of the two remains a risk the reader may construct but this record does not demonstrate.

10.8 The coincidence premium

Each route's annual replacement electricity is priced twice: at the wholesale index of the hour in which it is drawn, and at the year's flat mean price. The difference is the coincidence premium — the cost of *when* heating draws power rather than how much. It is small and positive for all three routes, largest for air-source, and is a wholesale-basis figure.

10.9 Verification

Four gates run before any hourly output is published. G1 confirms each route's calibrated performance factor reproduces its anchor. G2 reconciles the hourly network-route electricity against the weekly estimator's equivalent figure; it is near-tautological by construction and therefore proves wiring and heat conservation rather than independent physics. G3 confirms cooling electricity is conserved against its annual anchor. A feed gate refuses any series missing more than five per cent of its hours. Failure of any gate leaves the previous hourly output in place and the weekly panels untouched.

Independent corroboration comes from elsewhere: the located worst hour falls inside the peak week the weekly gas estimator identifies separately, and on the same day as the year's record observed demand hour.

Notes for investigators

Flow temperature is the largest single lever, and the field data says the model understates it. HeatpumpMonitor.org [35] publishes, from several hundred MID-metered UK systems at the SEPEMO H4 boundary, the regression

```
SPF_H4 = −0.1054 · (heat-weighted flow − outside) + 6.8762,   R² 0.547
```

This is the strongest correlation in that dataset, and implies **2.7–3.1% per kelvin** over a 25–35 K lift, against this model's 1.9% for air-source at the same condition. **The model therefore understates the effect on the one route for which field data exists.** The most likely cause is that η is held constant: if the Carnot fraction itself improves as the lift shrinks — less superheat, better part-load matching — the field slope contains two effects where this model contains one. That would make the network figure conservative rather than optimistic.

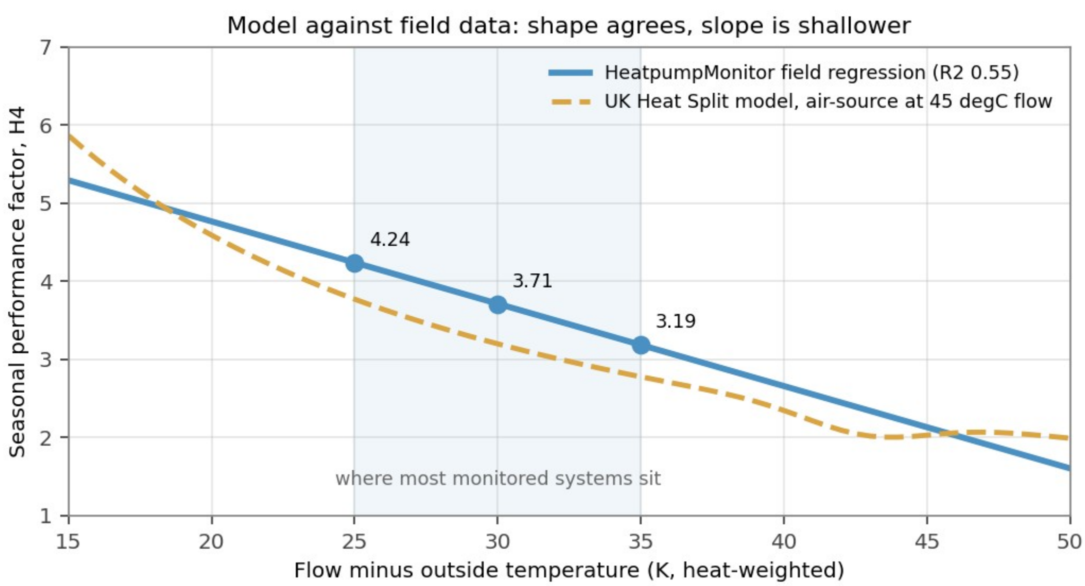

Figure 7

Figure 7. Model against field data: the shape agrees, the model's slope is shallower.

Two related observations from the same source are worth passing on. Systems achieving SPF 4.0 run average flow temperatures of 36.6 °C on the coldest days, against design rules of thumb of 45 and 50 °C — buildings are being designed roughly 10 degrees hotter than they turn out to need. And the same 5 K flow reduction is worth about 7.7% to air-source, 11.8% to ground source and 17.8% to a network: **low-temperature emitters and a warm shared loop are complements, not alternatives.**

The site's own comparison against field data uses different x variables and this should be stated when the two are shown together: the field regression's independent variable is heat-weighted flow minus outside temperature across a whole year, while the model's is a single flow at a single outside temperature. Aggregation over a season is not the same as a point on a curve, and for a convex function it biases the two apart.

---

11. Automation, failure handling and reproducibility

A scheduled GitHub Actions workflow executes daily at 03:43 UTC: fetch, fit, compute, write, commit. The complete source, the published data files and the build history are public.

Four hourly gates are described in Section 10.9. Three whole-build gates also run:

Annual-quantity plausibility. Trailing-twelve-month heating degree days must fall within 1,200–3,200.

Document structure. The published HTML is checked for content before the doctype, balanced script tags, and JavaScript leaking into markup. A syntax check on an extracted script proves the script parses, not that the document is well formed.

Scope. The build script's `main()` is walked for names read but never assigned in that scope and not module-level. A syntax check cannot see an out-of-scope name.

Four test suites run against the built page and payload before release, asserting that published claims match their own figures — including that the frontispiece points quote the panels they summarise, and that the value-for-money case still stands with its cooling component removed.

Notes for investigators

Every one of these gates was written after a defect reached publication, and the pattern in those defects is consistent: a change propagating to one consumer and not another. A correction to the ventilation split reached the useful-energy chain but not the purchased chain; a continuous curve fit reached the daily chart but not the weekly headline; a service toggle reached the wholesale basis but not the retail one. Structural tests do not catch this class, because both consumers remain internally consistent. What catches it is either an explicit invariant asserted at build time, or a reader comparing two published numbers and asking why they differ.

---

12. Principal limitations

— **The ECUK cooling and ventilation split** is sourced (54:46, Section 6.2) but rests on a 2014–15 survey applied to later metered totals. No comprehensive UK non-domestic end-use survey has superseded BEES.

— **The cooling EER of 3.0** (range 2.5–4.0) propagates one-for-one into every cooling service figure and is the largest single unverified assumption in that chain.

— **The comfort-deficit stock fractions and thermal-response coefficients** carry a several-fold range.

— **The observed cooling response is a lower bound**, not an independent estimate, because within-class centring removes the level. The undemeaned reconciliation recovers 75–95% of the anchor but is not used in published figures.

— **Tier 0 rests on evidence dating largely from 2019 and earlier**, and carries an unresolved overlap of 1–1.5 TWh with the services cooling anchor.

— **The geothermal network SCOP of 5.0 and source of 19.6 °C** have no monitored UK fleet behind them, and are one estimate with two faces rather than two that agree. Every network figure inherits this.

— **The flow-temperature sensitivity for the network route** extrapolates beyond any monitored system, and the model is known to understate the effect on the one route where field data exists.

— **Geothermal output is not nationally metered**; the present-day bar remains the most weakly constrained figure on the site.

— **The capacity line is a static overlay** on two winters of observed weather, with per-winter blocks that must be revised each autumn. Two is not a distribution.

— **Annual anchors lag** by roughly nine to fifteen months.

— **Hourly heat shape is modelled from weather, not metered.**

— **The value-for-money appraisal is one counterfactual**, not a full Green Book appraisal (Section 9 of the site; see the panel's own scope statement).

---

13. Corrections declared

The site's practice is to declare corrections rather than silently restate. Nineteen were found and fixed between v7.1 and v8; the first nine arose from an external peer review of 17 August 2026 and are marked *(PR)*.

#	Defect	Effect	Resolution
1	Regression window used as if it were a	Every annual quantity doubled;	`ANNUAL_DAYS` separated from

#	Defect	Effect	Resolution
	year	electrification limit halved for one day	`WINDOW_DAYS`; plausibility gate added
2	Comfort-deficit threshold applied an indoor criterion to outdoor air	Tier understated roughly five-fold	Rebased to 20 °C daily mean from EFUS
3	Binned response curve imposed a ceiling	Hot days understated; annual total 26% low	Continuous quadratic fit
4	Saturation tier premise contradicted by the data	A published band that was structurally zero	Tier withdrawn; computation kept as a diagnostic
5	Ventilation credited with a coefficient of performance	Cooling service overstated by about one third	Ventilation reported separately at unity
6	Single capacity block applied to two winters	2024/25 tested against 2025/26's fleet	Per-winter blocks
7	Network source stated as 10 °C	Understated the network COP and omitted the storage mechanism	Corrected to 19.6 °C, charged loop
8	One boiler efficiency applied to both end uses	Hot water credited with a condensing recovery it does not achieve	Separate SAP-derived hot-water efficiencies
9	Price-chart service selector short-circuited on basis	On the retail basis the service toggle was ignored	Six explicit basis-service combinations
10 *(PR)*	**Embedded solar added twice in three diagnostics**	The diurnal comfort/process separation was an artefact (2.59× collapses to 0.93); the 74% cooling reconciliation invalid. Night-hour results unaffected	All three call sites corrected; diurnal chart and 74% figure withdrawn
11 *(PR)*	**Network source presented as an independent check**	19.6 °C is derived to make the fractions agree, so the	Claim withdrawn; sensitivity range published instead

#	Defect	Effect	Resolution
	on the Carnot fraction	agreement evidenced nothing	
12 *(PR)*	**Front-end added the air-source approach instead of subtracting it**	Displayed air-source COP too favourable	Sign corrected; fallback given the payload's sign
13 *(PR)*	**"Electrification limit" implied capacity-adequacy modelling**	The calculation is a historical headroom stress test	Renamed; scope stated on-panel
14 *(PR)*	**Humidity null obtained under a bias running toward null**	The original null rested on the contaminated variable; two partial-record reruns then reversed it twice	Rerun on the full 800 days with solar controlled; null properly obtained
15 *(PR)*	**Weekly Tier 1 kept the binned estimator after the daily series moved to a continuous fit**	The weekly bar understated a hot week by about 80%	Fit lifted above both accumulations; one estimator
16 *(PR)*	**Gas-network GB factor applied to every fuel**	GB oil heat overstated by about 19%	Per-fuel shares from NISRA and NIHE household data
17 *(PR)*	**VFM headline claimed a general heat-network result from one counterfactual**	Overstated an illustrative single-counterfactual appraisal	Heading, label and claim rescoped
18 *(PR)*	**Pre-FID development capital excluded from a social appraisal**	A financing argument used to exclude a resource cost	Brought inside at a blended 7.40%; BCR 2.53 → 2.37
19 *(PR)*	**Network SPF 5 system boundary not stated**	A whole-system figure and a plant-only COP are not comparable	Boundary listed explicitly; parasitic sensitivity published
20	Tier 0 published as a "bottom-up estimate" with no bottom-up	A figure several times the other tiers, with no components or sources	Derived component by component (Section 4.6); restated as service

#	Defect	Effect	Resolution
21	Ventilation counted as cooling in the bill, carbon and what-if chains	The 20% what-if displaced a fifth of national fan energy — about 1.09 TWh/yr — with ground cooling	Ventilation made its own component throughout and carried through the what-if unchanged
22	Tier 1's share of the anchor stated as "about a seventh"	The published series had moved to a continuous quadratic; the true share is about a third	Recomputed from the payload
23	Trend-window totals threw on every click	A duplicated block referenced a variable from another scope; no window total updated	Duplicate removed; the 24-month window also gained its missing entry
24	Restatement markers written before the restatement ran	A failure would advance the epoch while the history fell back, and never retry	Markers written only on success; `FORCE_RESTATE` override added

14. What changed since v7.1

Structure. The site was reorganised into six numbered sections with a value-for-money panel at the end.

Cooling was rebuilt from the anchor outward. The ECUK cooling-and- ventilation line is split 54:46 from BEES rather than 50/50 by assumption; ventilation is a separate service throughout; the shape base moved from 18 °C to 15.5 °C; the binned response became a continuous quadratic; the saturation tier was withdrawn; and Tier 0 was derived from source for the first time.

Three things were tested and reported as negative results rather than quietly dropped: humidity, the slope-to-level inversion, and the saturation tier. Two further routes to the cooling floor were tried offline and are recorded as dead.

The hourly store holds two winters, each tested against its own NESO capacity block, and COP is evaluated at each day's temperature rather than as a fixed seasonal factor.

A heat-price series was added for four routes, two bases and three services, with a delivered-volume band beneath it.

Boilers gained a separate hot-water efficiency, and the UK-to-GB step became per fuel.

Pre-FID development capital was brought inside the value-for-money appraisal, and that appraisal was rescoped to what it is: one counterfactual on Green Book conventions.

Twenty-four defects were found and declared, nine of them by external peer review. Six new build gates were added as a direct result.

References

[1] Watson, S.D., Lomas, K.J., Buswell, R.A. (2019) Decarbonising domestic heating: what is the peak GB demand? Energy Policy 126, 533-544.

[2] Watson, S.D., Lomas, K.J., Buswell, R.A. (2021) How will heat pumps alter national half-hourly heat demands? Empirical modelling based on GB field trials. Energy and Buildings 238, 110777.

[3] Sansom, R. (2014) Decarbonising low grade heat for a low carbon future. PhD thesis, Imperial College London.

[4] Department for Energy Security and Net Zero (2025) Energy Consumption in the UK (ECUK) 2025. End-use tables (domestic U3; overall U5). OGL v3.0.

[5] Regulatory Assistance Project / Rosenow, J. et al. (2020) In-situ measured gas-boiler efficiencies; UK field-trial synthesis.

[5a] Building Research Establishment (2012) Standard Assessment Procedure (SAP) 2012, Table 4b: seasonal efficiency for gas and oil boilers, p207.

[6] Energy Systems Catapult (2023) In-Situ Heat Pump Performance, 1,700+ installations 2017-2022: ground-source seasonal performance factor 3.24.

[7] Energy Systems Catapult (2023) Electrification of Heat Demonstration Project, 742 heat pumps: annual SPF 2.80.

[8] Chartered Institution of Building Services Engineers (2013) TM52; CIBSE (2017) TM59.

[8a] Building Research Establishment (2017) Energy Follow-Up Survey: indoor/outdoor temperature regression.

[9] Ministry of Housing, Communities and Local Government (2023) English Housing Survey 2022-23. OGL v3.0.

[10] Climate Change Committee (2025) A Well-Adapted UK: progress in preparing for climate risks.

[11] Department for Energy Security and Net Zero (2025) Greenhouse gas reporting: conversion factors 2025. OGL v3.0.

[12] Office of Gas and Electricity Markets (Ofgem) Energy price cap: default-tariff unit rates.

[13] National Gas Transmission Operational Data REST API: LDZ and NTS demand; System Average Price.

[14] Hersbach, H. et al. (2020) The ERA5 global reanalysis. QJRMS 146, 1999-2049.

[15] Open-Meteo (2023) Historical Weather API (ERA5 archive). CC BY 4.0.

[16] National Energy System Operator Historic Demand Data and demand update files. NESO Open Licence.

[17] Climate Change Committee; National Energy System Operator (2025). UK domestic air-conditioning penetration below 5%.

[18] Department for Business, Energy and Industrial Strategy (2016) Building Energy Efficiency Survey (BEES) 2014-15.

[19] Underground thermal energy storage round-trip efficiency: EGEC [25] and aquifer/borehole thermal-storage literature.

[20] Elexon BMRS Market Index Data. Contains BMRS data (c) Elexon Limited.

[21] National Energy System Operator Carbon Intensity API.

[22] Department for Energy Security and Net Zero (2026) Digest of UK Energy Statistics (DUKES) 2026. OGL v3.0.

[23] Eurostat. Energy statistics methodology: ambient heat captured by heat pumps.

[24] European Geothermal Energy Council (2025) EGEC Geothermal Market Report 2025 - UK Country Update.

[25] European Geothermal Energy Council (2025) EGEC Geothermal Market Report 2025 - Key Findings.

[26] Climate Change Committee (2025) The Seventh Carbon Budget.

[27] Project InnerSpace, REA and ARUP (2026) The Future of Geothermal in the United Kingdom.

[28] Department for the Economy (NI) Energy in Northern Ireland; NISRA Continuous Household Survey 2024/25.

[28a] Northern Ireland Housing Executive and Building Research Establishment, Energy consumption in Northern Ireland’s housing stock: 2016.

[29] Department for Energy Security and Net Zero (2025) 2024 UK Greenhouse Gas Emissions, Provisional Figures.

[30] Department for Energy Security and Net Zero (2026) Quarterly Energy Prices, tables 3.4.1-3.4.2. OGL v3.0.

[31] Wooldridge, J.M. (2010) Econometric Analysis of Cross Section and Panel Data, 2nd ed. MIT Press.

[32] Foster, A., Brown, T., Evans, J. (2023) Carbon emissions from refrigeration used in the UK food industry. International Journal of Refrigeration 150, 297-303.

[33] Institute of Refrigeration / London South Bank University (2025) Transport, Industrial and Commercial Refrigeration (TICR), DESNZ-funded.

[34] Department for Energy Security and Net Zero (2026) Energy Trends special feature: Data centre electricity consumption in Great Britain, 2020 to 2024.

[35] OpenEnergyMonitor / HeatpumpMonitor.org. Monitored UK heat pump performance at the SEPEMO H4 boundary.

[36] Rystad Energy (2025) Geothermal economics: capex composition and well success rates.

Every author estimate on the dashboard carries a dagger and an invitation. Corrections and challenges to any figure, source or method in this document are welcomed at contact@causewaygt.com.

A Causeway Energies public-interest tool — causewaygt.com