

The House Is Never at Temperature: Comfortable Settings Form a Set of Measure Zero, with 118,414 Instrumented Room-Days and a Crossing Timetable

Wilhelmina Aske

Department of Pure Mathematics, Cawdray College

Nikolaj Sperber

Laboratory for Occupied Buildings,
Institute of Building Physics, Vasterhamm

5 September 2026

Abstract

You turn the heating up because the room is cold, and an hour later you are too warm and turn it back down. The setting you are looking for is not missing from your house: it is missing from the dial. A domestic heating system is a switch with a delay, and a switch with a delay cannot hold a temperature—it can only circle one. Two rhythms govern the circling: the boiler’s own minimum-cycle timer and the thermal lag of the building. Wherever their ratio locks, the delivered swing is wider than the band an occupant will sit still in, which we measure from a held button rather than from a questionnaire. Comfort therefore requires an unlocked ratio. Sweeping a 4,096-position dial in an occupied test dwelling, we find that the unlocked settings are nonempty and have measure zero: a set of box dimension 0.87 (95% CI 0.85–0.89). No dial can select from such a set, and a finer dial does not help—a concealed trial spanning a hundredfold range of setpoint resolution separates its arms by nothing. Households are comfortable for forty-one minutes of an occupied day, and every minute of it is spent on the way between two temperatures the house can hold and nobody will accept. Houses will be sold with a timetable where the thermostat used to be.

Note on terminology. Two words are used here in a narrow sense. A *locked* setting is one at which the boiler’s firing rhythm and the building’s thermal rhythm settle into a fixed ratio of whole numbers; the room then oscillates on a repeating schedule. A *crossing* is an interval during which the room’s temperature passes through the range its occupants accept, without stopping there. The distinction between passing through a range and resting in it is the whole of the measurement problem addressed below.

Keywords: singular sets; mode locking; thermal acceptance; relay control with dead time; domestic heating.

1 Introduction

The mid-market thermostat now quotes hundredths of a degree. Ten years ago it quoted halves, twenty years ago whole degrees, and the trajectory is treated throughout the trade as progress along an obvious axis: the finer the instrument, the closer the delivered room to the commanded one. That reading is about to fail, and it will fail in a way the industry has already paid for without noticing.

The complaint the instrument was built to answer is universal and is almost never investigated. An occupant is cold and raises the setpoint. Some time later—long enough that the two events are rarely connected—the same occupant is too warm and lowers it. The interval between those actions is not spent searching a fine grid for a value that the dial is too coarse to express. We show that it is spent moving between two values that the dwelling can hold, neither of which the occupant will accept, and that the value the occupant wants lies in a set that no dial of any resolution can address.

Before improving a component one establishes that the component is on the critical path; the diagnostic question is whether a hundredfold improvement in the part under consideration would be visible at the output. For domestic temperature control this question appears never to have been asked before the improvement was undertaken. We ask it twice: once analytically, in Section 4, where the geometry of the reachable set is derived, and once empirically, in Section 6, in a concealed trial that varies setpoint resolution across two orders of magnitude and holds everything else fixed. Both answers agree, and both are negative.

Three things follow. We built an instrument that records thermal acceptance as an interval of time rather than as a rating, and it returns a band far narrower than the standards assume (Section 3). We show that under the two-clock model the settings at which a dwelling can hold an acceptable temperature contain no range of any width: between any two settings that work lies one that does not (Section 4). And we show that what cannot be selected can still be scheduled — a next-day crossing forecast, accurate to a few minutes (Sections 6 and 7).

The resolution era is over. It ended in the second heating season of our trial, and no one has announced it.

2 Related work

Thermal comfort has been treated for six decades as a problem of steady states. The dominant framework predicts a percentage of dissatisfied occupants from a vector of environmental and personal variables and, by construction, cannot descend below a fixed floor even at neutrality (Fanger and Olsen, 1972). Standards built on that framework specify acceptable *ranges*, typically one and a half to three degrees wide (Halloran and Vidmar, 2022). The literature says which steady states would be acceptable. It does not ask whether any of them can be reached.

The control literature has long known that they cannot be reached exactly. Relay control with transport delay produces a limit cycle whose amplitude is bounded below by the delay, independently of the switching band (Bexhill and Törnqvist, 2018; Dahlmann, 2004). This result is standard, is taught, and appears in manufacturers' own application notes. Its consequence—that the delivered temperature is an oscillation and never a value—is nonetheless absent from the comfort literature, which continues to model the room as a set point plus noise.

Mode locking in forced oscillators is likewise standard. Where two rhythms compete, the winding number as a function of a control parameter is a devil's staircase; at critical coupling the staircase is complete, the locked intervals carry full Lebesgue measure, and the unlocked complement is a Cantor set of Hausdorff dimension approximately 0.87 (Cadogan, 2011; Ivarsen, 2015). Complete staircases have been demonstrated in convection cells, in cardiac preparations, and in driven junctions (Grieve, 1996). To our knowledge the domestic heating loop, which contains two independent and unsynchronised clocks, has not previously been examined as a member of this class.

Finally, the household-practice literature has documented the adjustment ritual itself—the up-then-down sequence, its frequency, and the friction it generates between occupants—and has interpreted it as a behavioural phenomenon requiring a behavioural explanation (Menteith, 1989; Ferreira and Okpara, 2020). We take the ritual to be correct behaviour under a misdescribed instrument.

3 The acceptance width

3.1 Why acceptance must be measured as an interval

A questionnaire asks whether a respondent is comfortable in a room; it returns a rating attached to a moment, and the moments are chosen by the investigator. Sustained acceptance is a different quantity: the range within which an occupant will remain still, indefinitely, without




Figure 1: The domestic loop. The appliance clock and the fabric clock are marked; neither is referenced to the other.

adjusting clothing, position, or controls. The two are not interchangeable, and the second is what a heating system is asked to deliver.

We therefore instrumented acceptance directly. Each participating room carries a sprung handset (Sperber, 2024); occupants are asked to depress it for as long as they are content to sit as they are, and to release it when they are not. The spring returns the handset on release, so it cannot be left depressed. Presses are logged at 1 Hz against a globe thermometer at seated height, so the recorded variable is operative rather than dry-bulb temperature.

3.2 Result

Across 3,109 dwellings, the pooled acceptance region is symmetric about a person-specific neutral point T^* with half-width

$$w = 0.34^\circ\text{C} \quad (95\% \text{ CI } 0.31\text{--}0.37),$$

estimated by pooled interval regression with a random intercept per occupant. The release hazard is memoryless above the threshold and rises steeply at it, which supports treating acceptance as a set rather than as a preference ordering: within the band, occupants are indifferent; outside it, they act.

This width is between a quarter and a ninth of the ranges written into comfort standards. The discrepancy is not a contradiction. Standards report the range over which a population will, when asked, decline to complain. We report the range over which an individual will not move.

4 Model

4.1 Two clocks

Figure 1 gives the loop. The thermostat compares the sensed operative temperature $T(t)$ against a commanded setpoint s and closes a contact. The contact does not command the burner. It commands an appliance whose own protection logic imposes a minimum off-interval P_a , fixed at manufacture and typically between two and ten minutes. The building answers on a quite different timescale, its first-order response to a step in delivered power having a fabric constant λ of order tens of hours, with an emitter-to-sensor transport delay τ of order minutes.

The loop therefore contains two periodicities that are set independently, are not communicated between components, and are not adjustable by the occupant: the appliance’s permission rhythm P_a and the fabric’s demand rhythm $P_f(s)$, the latter varying with setpoint because the heat loss to outdoors varies with the temperature the dwelling is being held near.

4.2 The amplitude floor

Write r_+ and r_- for the magnitudes of the heating and cooling slopes at the sensor and h for the switching hysteresis.

Proposition 1 (Amplitude floor). *For a relay loop with hysteresis h , dead time $\tau > 0$ and slopes $r_{\pm} > 0$, in which successive switch instants arrive at a fixed phase of the fabric cycle, every attracting limit cycle has peak-to-peak amplitude*

$$A \geq h + \tau(r_+ + r_-),$$

and consequently $\liminf_{h \rightarrow 0} A \geq \tau(r_+ + r_-) > 0$.

Proof sketch. At the switch instant the sensed temperature already leads the commanded threshold by $r_{\pm}\tau$, and the two leads so generated add to the band h traversed between switch points. They add only when consecutive cycles arrive in register, which is the hypothesis of the proposition. The construction is standard and is carried out in Appendix A. $\square$

Proposition 1 already disposes of the naive reading of the resolution trend: the amplitude the occupant experiences is bounded below by a transport delay that no refinement of the commanded value can shorten. It does not yet explain why comfort is observed at all, since $\tau(r_+ + r_-)$ is typically 0.45 °C to 3.1 °C across the cohort, and thus several times the acceptance band $2w$. The resolution of that puzzle is the subject of the next subsection, and it is where the two clocks become essential.

4.3 Locking, and the resonance hypothesis

Let $\rho(s)$ denote the winding number of the loop at setpoint s : the long-run ratio of appliance firings to fabric cycles,

$$\rho(s) = \lim_{t \rightarrow \infty} \frac{N_{\text{fire}}(t)}{N_{\text{fabric}}(t)},$$

which exists for every s by monotonicity and continuity of the lift, and depends continuously on s . Where $\rho(s)$ is rational, the loop is *locked*: firings and fabric cycles repeat on a common period, and the delivered trajectory is periodic with that period. Where $\rho(s)$ is irrational, the loop is *quasi-periodic*: firings never repeat their phase against the fabric cycle, and the delivered trajectory is dense on the invariant two-torus of the underlying flow.

Hypothesis 1 (Resonance). The lower bound of Proposition 1 is attained only under locking. Where $\rho(s) \notin \mathbb{Q}$, successive firings arrive at uniformly distributed phases of the fabric cycle, the overshoots of consecutive cycles do not accumulate in register, and the delivered amplitude falls to $A(s) = h + \tau(r_+ + r_-) \kappa(\rho)$ with $\kappa(\rho) < (2w - h)/\tau(r_+ + r_-)$ for every irrational ρ in the operating range.

Hypothesis 1 is the model's one substantive commitment. It says that the amplitude floor is a resonance phenomenon rather than a kinematic one: that the overshoots of Proposition 1 are additive only when they arrive in phase, and that an unlocked loop distributes them. It is the one place where Proposition 1's register condition is asserted to be necessary rather than merely sufficient. Section 5 tests it directly, by measuring $A(s)$ against $\rho(s)$ at 4,096 dial positions.

Definition 1 (The comfortable set). $\mathcal{C} = \{s \in [s_{\min}, s_{\max}] : A(s) < 2w\}$, the setpoints at which the sustained delivered trajectory remains within the acceptance band $[T^* - w, T^* + w]$.




Figure 2: A single evening, with three setpoint changes. The acceptance band is entered three times; the trajectory does not remain inside it.

4.4 The geometry of $\mathcal{C}$

Theorem 1 (The comfortable set is null). *Assume Hypothesis 1 and let the loop be critically coupled, so that the winding-number map $s \mapsto \rho(s)$ is a complete devil’s staircase. Then $\mathcal{C}$ is nonempty, totally disconnected, contains no interval, and satisfies $\text{Leb}(\mathcal{C}) = 0$; its closure, obtained by adjoining the countable set of plateau endpoints, is a Cantor set of the same measure.*

Proof. By Hypothesis 1, $\mathcal{C} = \rho^{-1}(\mathbb{R} \setminus \mathbb{Q})$ intersected with the dial range. At critical coupling the locked plateaux $\rho^{-1}(p/q)$ form a countable family of closed intervals whose union has full measure in $[s_{\min}, s_{\max}]$; their complement is therefore a set of measure zero. It contains no interval, since the plateaux are dense; and it is nonempty, since ρ is continuous and surjective onto its range while $\mathbb{Q}$ is countable. Total disconnection follows from the absence of intervals. $\mathcal{C}$ is not itself closed — each plateau endpoint is a limit of irrational winding numbers but is rational — so it is the closure $\bar{\mathcal{C}}$, which adjoins those countably many endpoints and is therefore still null, that is closed, bounded, perfect and totally disconnected, and hence a Cantor set. $\square$

Corollary 1 (Resolution is not on the critical path). *Let a dial expose N addressable positions on a grid of any spacing, placed at an offset θ drawn from any distribution absolutely continuous with respect to Lebesgue measure. Then $\Pr[\exists n \leq N : s_n(\theta) \in \mathcal{C}] = 0$, and the bound is independent of N , including $N = \infty$.*

Proof. Each event $\{s_n(\theta) \in \mathcal{C}\}$ is null, since $\theta \mapsto s_n(\theta)$ is a translation and $\text{Leb}(\mathcal{C}) = 0$. By countable subadditivity the union over $n \geq 1$ is null, and a fortiori every finite subunion. The bound is therefore independent of N , which is the content of the corollary. $\square$

Corollary 2 (Comfort is transit). *Let the commanded setpoint lie outside $\mathcal{C}$ and satisfy $|s - T^*| > A(s)/2 + w$, so that the sustained trajectory does not reach the acceptance band. If an occupant then records acceptance on a set of times of positive measure, the trajectory is not stationary on that interval. Every such episode is a crossing of the acceptance band, of finite duration, bounded at both ends by a change of regime.*

Remark 1. Corollary 1 is worth stating plainly because it inverts a forty-year programme. A dial addressing hundredths of a degree does not stand a hundred times the chance of a dial addressing whole degrees. Both stand a chance of zero. The quantity that a manufacturer can profitably move is not the fineness of the grid but the phase of the appliance clock relative to the fabric—and no product on the market exposes it. Figure 2 shows a single evening under this description.




Figure 3: The winding number over the dial range, and what is left when the locked settings are removed. The comb is thresholded to the plateaux the sweep resolves and is drawn at 0.075 of the range; the complement itself, at 0.006, is below the reproduction limit of the page.

5 The sweep

5.1 Apparatus

A stepper-driven thermostat body was installed in an occupied test dwelling (solid-wall terrace, 1904, condensing combination boiler, eight radiators) in place of the resident control. The body exposes 4,096 addressable positions across 12 °C to 28 °C, a step of 3.9 mK. The occupants were not told which position was in force at any time and had no means of reading it.

Each position was held for 115 min — somewhat under two fabric time constants, with the residual trend removed by detrending the final 20 min of each dwell — and the sweep was executed in randomised order across eleven months, discarding positions whose dwell was interrupted. At each position we recorded the winding number $\rho(s)$ from the boiler current clamp and the fabric cycle, and the sustained amplitude $A(s)$ from the globe.

5.2 The staircase

The measured $\rho(s)$ is a complete devil’s staircase (Figure 3). Plateaux account for 0.994 of the swept range, the seven widest corresponding to the lowest-denominator rotation numbers. Within the plateaux the amplitude is uniformly large: median $A = 1.12$ °C (IQR 0.81–1.63), or 1.6 times the acceptance band $2w$. At positions not resolved as locked at the sweep’s spacing, amplitude collapses: median $A = 0.21$ °C, well inside the band, exactly as Hypothesis 1 requires. We record these as candidate rather than confirmed members of $\mathcal{C}$: a finite sweep can exhibit no plateau, and cannot certify the absence of one.

Box-counting on the fitted plateau widths (Yeo and Radulescu, 2013), extrapolated over the three and a half decades of scale the sweep spans, gives a dimension

$$D = 0.87 \quad (95\% \text{ CI } 0.85\text{--}0.89),$$

which coincides with the value reported for the unlocked complement of critical circle maps (Cadogan, 2011; Ivarsen, 2015). We take the coincidence as corroboration that the domestic loop belongs to that universality class.

5.3 Where the dwelling sits

The critical line matters practically. Below it the tongues do not fill the range, the complement has positive measure, and a sufficiently fine dial would find it. All 3,109 field dwellings, and




Figure 4: Locked regions in the dwelling’s parameter plane, with the critical line and the golden-mean rotation point marked. Tongues are rendered for denominators up to sixteen. Cohort operating points lie on the line, within the estimation error of K .

Dwelling class	Dwellings	Room-days	Median λ (h)	Median τ (min)
Solid-wall terrace	941	36,104	61	14
Cavity-wall semi	1,088	41,657	44	11
Post-2005 apartment	702	26,481	27	8
Timber-frame detached	337	12,609	19	6
High-mass underfloor	41	1,563	96	41
Total	3,109	118,414	44	11

Table 1: The instrumented cohort. λ is the fabric time constant, τ the emitter-to-sensor transport delay.

the test dwelling, operate at criticality to within the estimation error of K : the appliance clock is stiff, the fabric is heavy, and the coupling between them is not a design parameter of any product in the cohort. Figure 4 shows the plane. The operating points cluster on the critical line, and the coupling that would move them off it is not a purchasable parameter.

6 Field results

6.1 Cohort

Table 1 describes the instrumented cohort: 3,109 dwellings in four national housing stocks, 118,414 instrumented room-days between 2019 and 2026.

6.2 The profile: where the error actually is

Before testing any intervention we decomposed the variance of the delivered-temperature error into components, by mixed model with dwelling and season as random effects. The result (Figure 5) is that the phase of the appliance clock relative to the fabric carries 61.2% of the variance, fabric lag magnitude 24.7%, sensor siting 8.4%, and setpoint quantisation 1.8%, with 3.9% residual.

Setpoint quantisation is the component the industry has spent four decades improving. Removing it entirely — that is, granting a dial of infinite resolution — would remove 1.8% of the error in delivered temperature. On the profile alone, resolution is not a promising target. Theorem 1 says something stronger: it is not a target at all.



Figure 5: The bottleneck profile. Variance shares of delivered-temperature error; setpoint quantisation is 1.8%.




Figure 6: Comfort delivered against precision commanded, across two decades of setpoint resolution. The shaded band is the pre-registered equivalence margin.

6.3 The resolution trial

We randomised 1,044 dwellings, concealed, to one of three thermostats, identical in appearance, differing only in the quantisation of the commanded setpoint: 1.0 °C, 0.1 °C and 0.01 °C. Two heating seasons. The primary outcome, registered in advance, was comfort-minutes per occupied day from the held-button record.

Comfort-minutes per occupied day were 41.2, 40.9 and 40.6 in the three arms. The slope per decade of refinement is -0.3 minutes (95% CI -1.9 to $+1.3$), against a pre-registered equivalence margin of ± 5 minutes (Figure 6). Two orders of magnitude of instrument precision deliver a difference indistinguishable from zero, which is the difference Corollary 1 predicts exactly.

We note that the occupants of the 0.01 °C arm adjusted their thermostats 2.4 times as often as those of the 1.0 °C arm. Finer resolution did not improve the delivered room. It increased the rate of adjustment.

6.4 The crossing census

Segmenting every occupied day into intervals satisfying $|T(t) - T^*| < w$ gives the census in Table 2. The median household is comfortable for 41 minutes of an occupied day, distributed over 2.7 episodes of median duration 15.2 minutes — close to the traverse time $2w/r$ implied by the sensed slopes of Section 4. No episode outside the underfloor class exceeds 34 minutes, in 118,414 instrumented room-days.

The premise of Corollary 2 holds almost everywhere in the record: the commanded setpoint satisfied $|s - T^*| > A(s)/2 + w$, and the sustained trajectory therefore did not reach the band at all, in 96.2% of occupied hours. In that 96.2%, every episode is bounded at both ends by a change of regime: a setpoint adjustment, a door, a solar step, or an occupancy change. This is

Dwelling class	Comfort min/day	Episodes/day	Median episode (min)	Longest (min)
Solid-wall terrace	38	2.4	15.8	32
Cavity-wall semi	41	2.7	15.2	33
Post-2005 apartment	44	3.1	14.2	34
Timber-frame detached	47	3.5	13.4	34
High-mass underfloor	51	2.9	17.6	402
All except underfloor	41	2.7	15.2	34

Table 2: The crossing census. Three underfloor dwellings contribute every episode longer than an hour and are discussed in Section 8; the class is excluded from the pooled row.

the signature Corollary 2 requires. Comfort in these dwellings is not a condition that obtains; it is an event that occurs, and it occurs while something else is changing.

6.5 The golden-mean trial

$\mathcal{C}$ cannot be reached by search, and it cannot be reached by computation either: a commanded value is a grid position and the set has no width. What computation can locate is the neighbourhood in which the neighbouring plateaux are narrowest — the setpoint whose rotation number is nearest the golden mean, the irrational whose rational approximations converge most slowly, in the sense that the Hurwitz bound $|\rho - p/q| > 1/(\sqrt{5}q^2)$ is sharp there and $\sqrt{5}$ is the largest constant for which such a bound holds of any irrational. A dwelling commanded there does not sit in $\mathcal{C}$. It passes near it more often, and its crossings are longer.

We computed the golden point for each of 214 dwellings from its own fitted tongue diagram — predicted from λ and P_a — and randomised them to that setpoint or to the nearest conventional 0.5°C position. Comfort-minutes per occupied day rose from 41 to 187 (95% CI 171–203): the golden point is held only while outdoor temperature is stationary, so dwellings occupied it for a median of 187 of 960 occupied minutes before drift carried them into the neighbouring plateau. The gain is in any case unavailable to a household turning a dial: the target is 3.9 mK wide and it moves with the weather.

6.6 Crossing forecasts

The final experiment tests the practical consequence. For 486 dwellings we published a next-day crossing schedule — the clock times at which the dwelling’s trajectory would pass through the acceptance band — and asked occupants to be in the room at the stated times, blind to the model and to the schedule’s provenance. Predicted crossings arrived within ± 3.5 minutes in 78% of cases, against 34% for a same-length control window drawn from the previous day’s schedule.

A dwelling cannot be held at an acceptable temperature. It can be met at one.

7 Discussion

The finding is not that domestic heating is badly engineered. It is that it has been engineered against the wrong variable, with great competence, for a long time. The dial is a precision instrument pointed at a set it cannot resolve; every increment of that precision has been real, has been delivered, and has changed nothing that an occupant can feel.

The variable that moves the outcome is phase. The appliance clock is fixed at manufacture and the fabric clock is fixed by the building, and the relationship between them — the only quantity in the loop that determines whether the two rhythms lock together or drift past one another — is not readable from either, and is published by no manufacturer. It could be. A programmable anti-cycle interval is a firmware change to equipment that already contains

a timer, and it converts an inaccessible set into an addressable one. We estimate the *phase budget* of a typical condensing appliance — the range of minimum off-intervals compatible with its protection requirements — at roughly four minutes, which is wider than the widest low-denominator plateau measured in any dwelling class.

The second consequence is a change of instrument. Where the two rhythms cannot be separated, comfort remains a crossing, and crossings are predictable a day ahead to within a few minutes. A dwelling’s schedule of crossings is a publishable quantity, computable from the fabric constant, the appliance period, and a weather forecast. It belongs on the same page as the tide table, and for the same reason: it describes when a thing passes, not where it rests.

Ten years from now the thermostat will occupy the position the barometer occupies today: mounted in the hall, accurate, and no longer consulted before a decision. The dwellings sold in that decade will be sold with two numbers: a fabric constant and an appliance period. From those, a schedule. What should be done now is smaller and available immediately: stop paying for resolution, ask the manufacturer for the timer, and, in the meantime, consult the schedule rather than the dial.

8 Limitations

Three limitations bear on the strength of the claim.

First, the held button measures willingness to remain still. That is the right variable for a heating system and the wrong one for welfare: an occupant absorbed in work will hold the button through conditions they would not otherwise accept, and we have not attempted to separate absorption from acceptance.

Second, the staircase was measured in a single test dwelling. A single sweep occupied that dwelling for eleven months and a second was not funded, so every tongue diagram outside it is fitted rather than swept, and the field cohort constrains $\rho(s)$ only at the positions its occupants happened to command.

Third, and unresolved: in 3 of the 41 high-mass underfloor dwellings, the comfortable set has clearly positive measure, occupying 0.19 of the dial range, and the fitted staircase is incomplete. These dwellings hold an acceptable temperature for hours at a stretch — one for 402 consecutive minutes — and our model does not predict them. Their fabric constants and appliance periods are extreme, but not anomalous within their class: the other 38 underfloor dwellings share them and behave like the rest of the cohort. Their instrumentation was checked twice, and the effect persisted across two seasons and a change of occupants. We do not know what is different about these three houses. A model that cannot account for them is not finished.

9 Conclusion

A domestic heating system cannot hold a temperature; it can only circle one, and the settings at which the circle is small enough to accept form a set of measure zero. Comfort has therefore never been selected in these dwellings. It has been crossed, fifteen minutes at a time, forty-one minutes a day.

The remedy is not a finer instrument. It is a published schedule, and a firmware change to a timer that already exists. Houses will be sold with a timetable where the thermostat used to be.

Data availability

The room-day corpus, the sweep record, and the fitted tongue diagrams for all 3,109 dwellings are deposited with the Institute of Building Physics, Vasterhamn, under embargo until the fitted tongue diagrams are independently replicated.

Funding and ethics

Supported by the Cawdray College Fund for Applied Analysis and by an instrumentation grant from the Institute of Building Physics. The protocol was reviewed by the Vasterhamm Committee on Occupied Research (ref. VOR-2019-114). Occupants of the test dwelling were compensated at the local median rate for inconvenience and retained the right to override the sweep at any time; four positions were lost to that right and are excluded.

A The amplitude floor

Let the sensed temperature obey $\dot{T} = r_+$ while the relay is closed and $\dot{T} = -r_-$ while it is open, with the relay responding to $T(t - \tau)$. Suppose the relay opens when the delayed signal reaches $T^* + h/2$. The sensed temperature at that instant is already $T^* + h/2 + r_+\tau$, which is therefore a lower bound on the peak of that cycle. The symmetric construction on the falling edge gives a trough of at most $T^* - h/2 - r_-\tau$. Subtracting, $A \geq h + \tau(r_+ + r_-)$, as claimed, provided the two extrema belong to the same cycle — that is, provided the switch instants arrive in register with the fabric, which is the hypothesis of Proposition 1. Equality holds when the slopes are constant over the cycle, which is the case to within 4 % in the cohort.

References

- Bexhill, R. and Törnqvist, M. (2018). Relay control with transport delay: amplitude floors in domestic emitters. *Journal of Building Services Engineering*, 41(2), 88–104.
- Cadogan, P. (2011). Winding numbers in forced thermal oscillators. *Nonlinearity in Applied Systems*, 24(6), 1441–1467.
- Dahlmann, E. (2004). Anti-cycle protection and the phase problem in condensing appliances. *Applied Heat*, 19(1), 33–51.
- Fanger, P. O. and Olsen, B. (1972). Steady-state models of thermal acceptability and their dissatisfaction floor. *Archives of Environmental Engineering*, 8(1), 5–28.
- Ferreira, L. and Okpara, N. (2020). The occupant as controller: integral action with delay in domestic thermostats. *Ergonomics of Control*, 12(4), 301–318.
- Grieve, A. (1996). Complete devil’s staircases in dissipative systems: a review. *Reviews of Applied Dynamics*, 7, 1–44.
- Halloran, S. and Vidmar, T. (2022). Setpoint resolution and delivered comfort: a systematic review of forty years of claims. *Energy and Buildings Digest*, 77, 55–70.
- Ivarsen, K. (2015). Badly approximable rotation numbers and robust quasi-periodicity. *Journal of Approximation Dynamics*, 33(2), 145–170.
- Menteith, J. (1989). Thermostat interaction as household ritual. *Journal of Domestic Practice*, 6(1), 12–29.
- Sperber, N. (2024). A held-button telemeter for sustained thermal acceptance. *Instrumentation for Occupied Spaces*, 9(2), 77–92.
- Yeo, H. and Radulescu, D. (2013). Box-counting estimates for parameter-space complements. *Computational Measure*, 21, 405–422.