capstract:2611.00067v1 [econ.EM] 20 Nov 2026

The Cooling Capacity of a Reporting Clock

Sorell J. Arvenlow
Department of Econometrics, Brackenmere Institute for Applied Measurement

November 20, 2026

Abstract

This paper estimates the mechanical cooling capacity of a clock. The clock supplies timestamps for a facilities contract that infers cooling from the number of warm-room tickets filed by a specified reporting minute. We define a clock's capacity as the reduction in room temperature required to reproduce its ticket count under the reference rate. A randomized grid of mechanical doses identifies a set of equivalent temperatures. We also show that a suitable clock schedule can reproduce the full cumulative ticket curve of a colder room on any bounded reporting window subject to a support condition. In a simulated procurement trial with 1,152 sessions, doubling the reporting rate cuts the acceptance count by 15.2 percentage points and outperforms one degree of mechanical cooling. The projected cooling equivalent spans 1.0 to 2.5 °C conditional on a crossing in the tested range. Audits on a common elapsed clock assign the same device zero capacity.

Keywords: partial identification, duration analysis, clock rates, procurement

1 Introduction

A facilities department can test a cooling installation by comparing the number of warm-room tickets reported within a fixed period with the number received after a reference installation. For the comparison to have units of temperature, the two reporting periods must have units of time in common. We study a procurement rule under which the clock used to close the ticket count is part of the installed equipment. Two devices are available: a machine that cools the room and a clock that completes the reporting period sooner.

The intervention changes the rate of the clock used for the report. Occupants cannot see that clock. The intervention leaves the air handler and the ticket interface operating on elapsed time. A faster reporting clock changes neither the arrival process for tickets nor the thermal conditions that give rise to them. It can change the measured capacity of a room. We express the change in degrees by randomizing mechanical cooling and estimating which dose matches the outcome recorded under the faster clock.

There are two questions. First, how much mechanical capacity can the clock displace under a given reporting rule? A finite grid of reference temperatures generally identifies a range of answers. Second, can a slower sequence of reports, or access to the complete distribution of ticket times, distinguish the clock from refrigeration? On any bounded reporting window a suitable clock schedule can match the full ticket-arrival curve of a colder room. Identification requires information about elapsed time or the mass of occupants who never file. This distinction is related

to the separation of timing and eventual incidence in duration models with a cure fraction (Maller and Zhou, 1992).

We illustrate the problem with a simulated office procurement trial. The simulation combines a randomized temperature schedule with an exact replay of the same records under two rates. A reporting clock running at twice the reference rate clears the one-degree mechanical benchmark on the primary ledger. On a synchronized ledger, the same clock has zero capacity at every reporting horizon. The calculation can be made with the ticket register, six reference temperatures, and no assumptions about how office workers respond to the speed of clocks.

2 Model and design

2.1 Two clocks

Write $T_i(c)$ for occupant i's elapsed time to a first ticket requesting cooler air when the room is cooled by c °C relative to a common reference condition. We set $T_i(c) = \infty$ for occupants who never file. In a room cooled by c, the cumulative incidence in elapsed time is

$$F_c(t) = \mathbb{P}\{T_i(c) \le t\}. \tag{1}$$

The reporting clock's indicated age is s. Let $h(s)$ be the elapsed age at which the reporting clock reaches s. The ledger reports $F_c(h(s))$. For a clock that advances at constant rate q relative to elapsed time, $h(s) = s/q$. We take the clock to operate on the reporting system only. If enabling the clock also shortens breaks, closes the ticket interface or changes when maintenance staff respond, this exclusion fails. In the replay design the reporting system reads the underlying register and creates no feedback to it.

At reporting age s define $c^*(s, q)$ by

$$F_{c^*(s,q)}(s) = F_0(s/q), \qquad c^*(s,q) \in [0, 2.5]. \tag{2}$$

The right-hand side is the rate in the uncooled reference room counted by the faster clock. The left-hand side uses the reference clock. For a fixed reference temperature the degree value of a twofold increase in clock rate need not be constant over s. If both curves are differentiable at an interior crossing and the cooling derivative is negative, implicit differentiation gives

$$\frac{\partial c^*(s,q)}{\partial \log q} = \frac{-(s/q) f_0(s/q)}{\partial F_c(s)/\partial c \,\big|_{c=c^*(s,q)}}. \tag{3}$$

The numerator is the density of first tickets at the true cutoff, scaled by that cutoff. Increasing the clock rate is most productive where many reference-room tickets are about to arrive and actual cooling has little effect on the specified outcome.

2.2 Randomization and recording

The trial assigns six doses, from zero through 2.5 °C in half-degree steps. A block contains six sessions of eight occupants. The six dose assignments are uniformly permuted within block. Sessions in a block share ambient conditions. Treatments can affect interactions within a session. We assume independence across blocks and no interference across sessions. The reference temperature is 25.5 °C. An independent probe supplies the session temperature. The private comfort prompt at

elapsed minute 20 does not enter the ticket ledger. Participants can submit one ticket at any elapsed time during the 120-minute session and receive no response during that period.

The contractor sees the dashboard at reporting minute 20. A separate audit records tickets accumulated through elapsed minute 120 and responses to the private question. In the reference room we render two dashboards from the same register. With the reference clock, a ticket filed at elapsed minute 13 enters the contract count. At twice the rate, the same ticket enters at reporting minute 26. The reporting minute is changed by the replay. The source event is unchanged. The second replay requires no additional occupants.

3 Identification and estimation

3.1 Capacity on a finite grid

For the primary outcome define $D(c) = F_c(20) - F_0(10)$. At a tested dose the sign of $D(c)$ ranks the two devices by the criterion in the contract. A positive sign means the faster clock records fewer tickets than the mechanical device at c. In contrast, solving Eq. (2) for a continuous dose needs a model between tested values.

Proposition 1 (Grid identification)**.** *Suppose $D(c)$ is continuous and strictly decreasing on the design range. A zero at a tested dose uniquely identifies the equivalent. If consecutive grid values $D(c_j)$ and $D(c_{j+1})$ have opposite signs, the identified set is (c_j, c_{j+1}). If both endpoints are strictly of the same sign the design range contains no equivalent.*

Proof. Continuity gives existence and strict decrease gives uniqueness when the signs differ. For any candidate point between opposite signs, insert a zero at that point and interpolate linearly between all grid points and the inserted zero. The resulting curve satisfies the restrictions and reproduces the observed grid means. When the endpoints share a strict sign, the intermediate values do as well. ∎

For the audit outcomes we form corresponding contrasts against the reference room using elapsed minutes for both arms. This gives exact zeros at $c = 0$ for the fast-versus-reference comparison of elapsed ticket incidence and for the private prompt, because both clocks operate on the same records and neither enters these quantities. With strict decrease, zero is the design's unique equivalent under each audit outcome. The estimated contrasts at positive doses quantify how the devices compare at those doses.

We compute contrasts within block, then average them over blocks. This retains the covariance from sharing the same reference room across doses and from measuring both clock policies on the same records. Standard errors use the sample variance of the block differences divided by the number of blocks. We use Bonferroni normal intervals across 18 outcome-by-dose contrasts, including the two degenerate zeros, and report the interval at a zero-variance structural contrast as the singleton zero. The critical value is $\Phi^{-1}(1 - 0.05/36)$. A positive lower limit excludes crossings at or below its dose. A negative upper limit excludes crossings at or above its dose. We report the closure of the retained interval and indicate whether opposite signs are established. This projection covers the whole identified set asymptotically, conditional on the continuous monotone model. See Athey and Imbens (2017) on inference for randomized experiments and Imbens and Manski (2004) on the distinction between covering a particular parameter and its identified set.

For comparison we fit an unrestricted line to the six grid means and use the delta method for the fitted crossing. It estimates a continuous parameter under linearity, requires a slope separated

from zero for the usual normal approximation, and permits estimates outside the range. The calculation retains the covariance of intercept and slope using their block-level analogues. We use a nominal 95% interval for the inverse of the fitted line.

3.2 Fitting the full register

A procurement revision might require the contractor to submit cumulative ticket rates at every reporting minute. The next result gives the clock that matches all such rates simultaneously.

Proposition 2 (Equivalence of ticket histories)**.** *Suppose F_0 is continuous with $F_0(0) = 0$ and strictly increasing wherever $F_0(t) < p_0$, where $p_0 = \lim_{t\to\infty} F_0(t)$. Suppose F_c is continuous and strictly increasing on $[0, S]$, with $F_c(0) = 0$ and $F_c(S) < p_0$. Then the clock schedule*

$$h_c(s) = F_0^{-1}(F_c(s)), \qquad 0 \le s \le S, \tag{4}$$

matches the cumulative incidence of c degrees of mechanical cooling throughout the reporting window. It is unique on that window. If $\lim_{s\to\infty} F_c(s) < p_0$, no unbounded schedule can extend the match to every reporting age.

Proof. The inverse exists on $[0, p_0)$ by the continuity and strict increase of F_0 below p_0. Both functions in Eq. (4) are strictly increasing on the stated range, so the schedule is an admissible clock and substitution gives the equality. Inversion gives uniqueness. Under any unbounded increasing h, the limit of $F_0(h(s))$ is p_0. If the cooled room has lower eventual incidence the limits differ. ∎

When $F_c(s) \le F_0(s)$, this construction counts at most s elapsed minutes by reporting age s. Over a finite reporting window it can be extended past the deadline as an ordinary running clock. When the two eventual rates differ, extending the match indefinitely would require every arbitrarily late report to be compiled from records available by a finite elapsed time. More frequent entries from the same reporting system therefore cannot substitute for a timestamp obtained from an independent timebase.

4 Simulation

4.1 Records

The main design uses 192 independent blocks, 1,152 sessions, and 9,216 occupants. It generates the event register before either dashboard is computed. Draw block and session shocks $U_b \sim N(0, 0.4^2)$ and $V_{bi} \sim N(0, 0.2^2)$ independently. Each occupant has a binary warmth indicator with conditional probability

$$\operatorname{logit}^{-1}(0.28 + U_b + V_{bi} - 0.75c).$$

Warm occupants file a ticket by the end of the session with probability 0.85, and other occupants file with probability 0.10. Conditional on filing, the elapsed arrival distribution is piecewise linear through cumulative probabilities $(0, 0.35, 0.63, 0.87, 1)$ at $(0, 10, 20, 60, 120)$ minutes. Cooling thus changes eventual filing through warmth. The probability law of filing time conditional on filing is common to all arms. The probe measures $25.5 - c + 0.12U_b + e_{bi}$ with independent $e_{bi} \sim N(0, 0.07^2)$. Thermal-control experiments motivate measuring comfort independently of the re-

ported operating condition (Luo et al., 2016). The simulation parameters specify a counterexample rather than estimates from office buildings.

4.2 Acceptance test

Table 1 reports outcome means. In the reference room, doubling the clock rate reduces the share counted by reporting minute 20 by 15.2 percentage points (simultaneous 95% CI 12.5 to 18.0). The desk temperature, private prompt and synchronized ticket count are identical across the two rows for that room. None of the events omitted at the faster cutoff is removed from the source register.

Table 1: Simulated acceptance trial. Each mechanical dose contains 192 sessions of eight occupants. The twice-rate row reuses the reference-room tickets.

Cooling (°C)	Clock rate	Probe (°C)	Ticket, reporting min 20 (%)	Ticket, elapsed min 120 (%)	Too warm (%)
0.0	1×	25.50	34.2	51.9	55.7
0.5	1×	25.00	32.4	48.8	49.4
1.0	1×	24.51	25.6	39.8	39.3
1.5	1×	24.00	21.0	33.3	30.7
2.0	1×	23.50	16.5	25.8	22.1
2.5	1×	23.00	16.0	25.2	20.2
0.0	2×	25.50	18.9	51.9	55.7

The estimated early-ticket contrast changes sign between 1.5 and 2.0 °C of mechanical cooling. The simultaneous projected interval is 1.0 to 2.5 °C under an in-range crossing. The lower endpoint is excluded. The upper endpoint is the edge of the design, and the data do not establish that an in-range crossing exists. The dashboard at one degree gives a 6.6-point advantage to the faster clock (simultaneous 95% CI 1.8 to 11.5). The synchronized ledger at elapsed minute 120 favors the cooled room by 12.1 points. Figure 1 shows the comparisons against every tested dose, with the last two panels evaluated at common elapsed times.

Figure 2 shows what happens when the acceptance deadline varies. A constant twofold rate and a mechanical reduction do not give the same ticket curve throughout. Even so, a procurement test that accepts a clock schedule faces a stronger problem. Figure 3 applies Eq. (4) to the one-degree benchmark under the generating model. The schedule is calibrated once and then used for the full report. The schedule changes rate at breaks in the simulated arrival distribution. By reporting minute 120 the schedule has counted only 39.6 elapsed minutes. Access to the full curve on that interval leaves the two devices observationally equivalent.

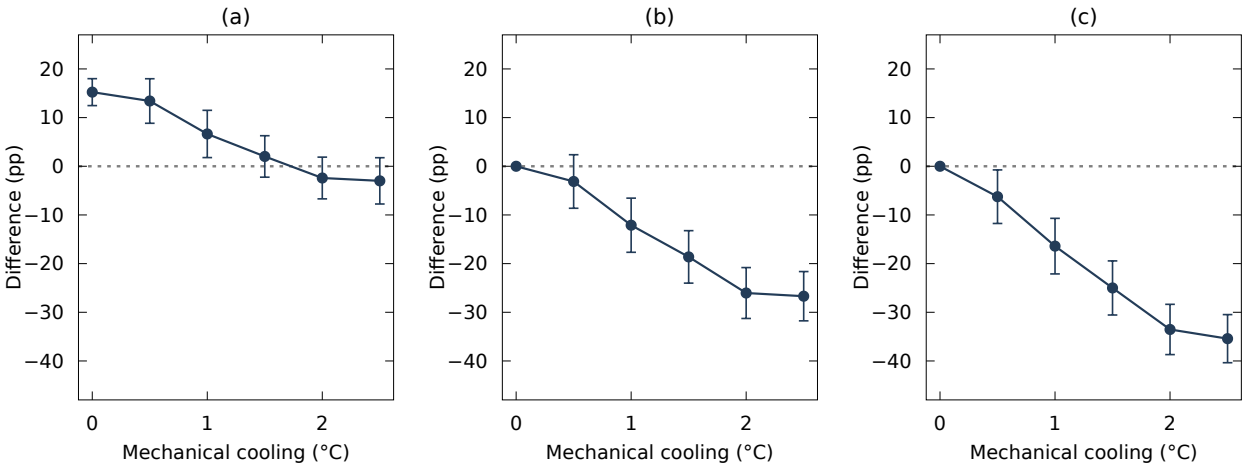


Figure 1: Comparison of mechanical cooling with a doubled reporting-clock rate in the reference room. Panel (a) uses tickets by reporting minute 20. Panels (b) and (c) use tickets by elapsed minute 120 and privately reported discomfort at elapsed minute 20. Whiskers are simultaneous 95% intervals across the 18 contrasts.

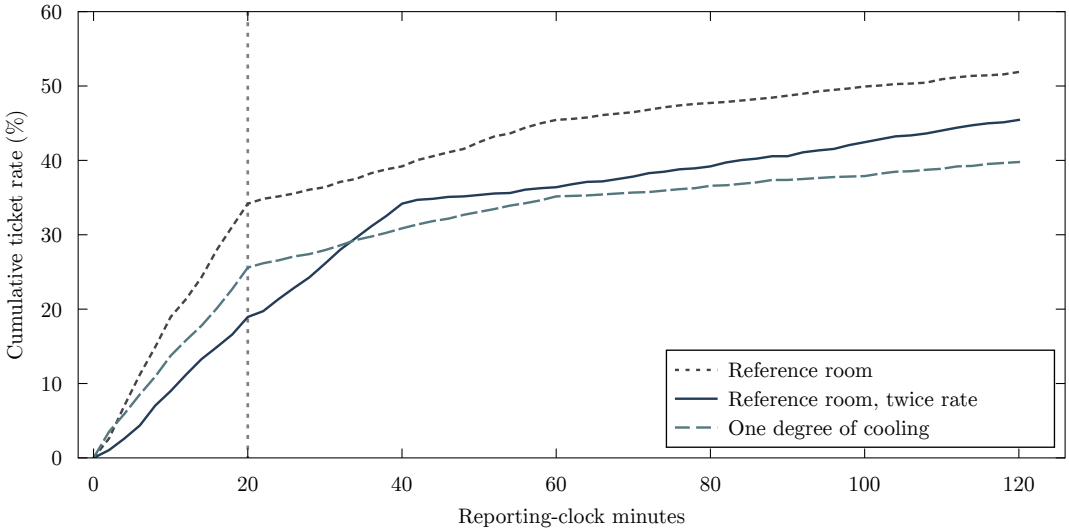


Figure 2: Cumulative incidence against indicated reporting time. The first two series use exactly the same ticket records. The vertical line is the acceptance deadline.

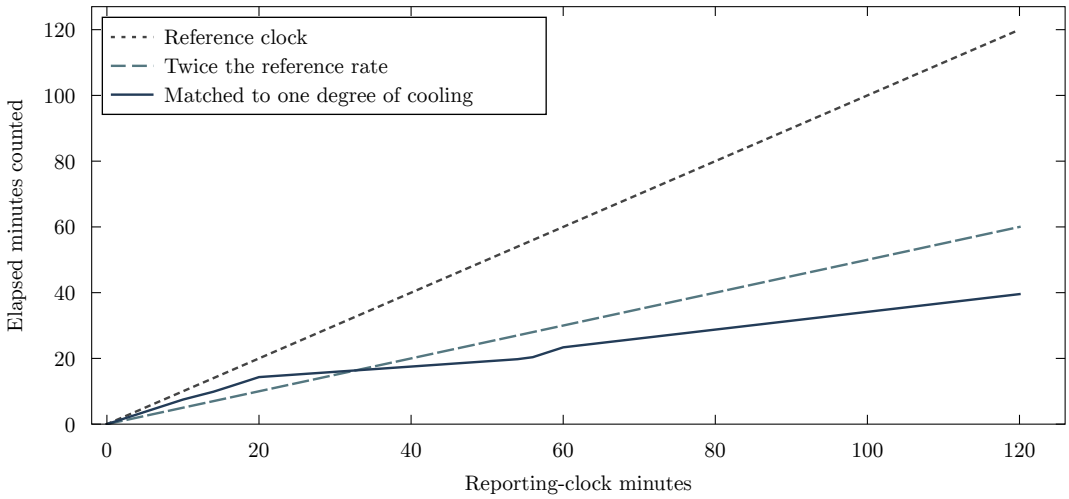


Figure 3: Elapsed minutes included by each reporting clock. The matched schedule reproduces the population ticket curve of a room cooled by 1 °C at every indicated minute through minute 120.

4.3 Repeated samples

Table 2 uses 1,200 samples per specification. The curved benchmark replaces the linear index penalty $0.75c$ with $0.15c + 0.35c^2$. The weak benchmark sets the penalty to $0.10c$ and uses a clock rate of $56/53$. The modest increase in rate in the weak design keeps the population crossing within the tested range. Population crossings are calculated by integrating over the normal shocks. Coverage for projection refers to the entire identified set. Coverage for the linear interval refers to the generating crossing.

Table 2: Repeated-sample performance for the clock's cooling equivalent at reporting minute 20. Widths are medians. Full means the projection retains the entire tested range.

Model	True (°C)	Set cov. (%)	Linear cov. (%)	Set width (°C)	Linear width (°C)	Full (%)
Baseline	1.86	100.0	95.2	1.50	0.59	0.0
Curved	1.80	100.0	96.0	1.50	0.46	0.0
Weak	1.41	99.8	93.8	2.50	3.92	97.4

The weak benchmark can yield a large positive capacity estimate even when the effect of small changes to the clock is precisely measured in ticket units. The poorly determined conversion is the mechanical temperature response. In the curved design the linear estimator has a median interval width of less than half a degree. Its coverage here depends on how closely the projection of the curve onto a line approximates the crossing.

5 Capacity procurement

Suppose the contract values a reduction in ticket probability at v per occupant and lists device costs per occupant as K_q and K_c. It selects the clock whenever $K_q - K_c < vD(c)$. At equal costs and a one-degree benchmark the acceptance ledger and the synchronized audit select different de-

vices. The first fits the measured capacity target with a change to the rate at which the observation period expires. The second measures how many people report a warm room during an equally long exposure. Costs and decisions in this example follow from the specified contract. A comfort valuation would require an outcome and a price for comfort.

The minimum attainable ticket rate depends on the rates permitted by the contract. In the model, $F_0(s/q)$ tends to zero as q grows without bound, for any finite s. Consequently a ticket ceiling with any strictly positive allowable rate can be met without changing the reference room. Requiring records at additional finite horizons continues to admit a sufficiently large rate. Requiring independent elapsed times changes the constraint. Under that requirement a clock purchased separately from the plant has a measured mechanical capacity of zero.

Data and code availability

The event register, tables and figure values can be regenerated using the supplied Python program. The audit uses seed 413867 and the repeated samples use 831163. All occupant records are simulated.

References

Susan Athey and Guido W. Imbens. The Econometrics of Randomized Experiments. In *Handbook of Economic Field Experiments*, volume 1, pages 73–140. North-Holland, 2017. doi: 10.1016/bs.hefe.2016.10.003.

Guido W. Imbens and Charles F. Manski. Confidence Intervals for Partially Identified Parameters. *Econometrica*, 72(6):1845–1857, 2004. doi: 10.1111/j.1468-0262.2004.00555.x.

Maohui Luo, Bin Cao, Wenjie Ji, Qin Ouyang, Borong Lin, and Yingxin Zhu. The Underlying Linkage Between Personal Control and Thermal Comfort: Psychological or Physical Effects? *Energy and Buildings*, 111:56–63, 2016. doi: 10.1016/j.enbuild.2015.11.004.

R. A. Maller and S. Zhou. Estimating the Proportion of Immunes in a Censored Sample. *Biometrika*, 79(4):731–739, 1992. doi: 10.1093/biomet/79.4.731.