The Fire Extinguisher Is an Inspection Instrument:

Charge-Blinded Randomisation in 41,208 Rooms, a 1.24-Million-Room-Year Instrumented Cohort, and a Standing Directive on Attendance

Ines Halloran, *for the CHARGE-1 Investigators*

Nettlefold Institute for Loss Prevention

Abstract

A fire extinguisher hangs within thirty feet of most people who will read this. Almost nobody has ever seen one used, and the reason is not that fires are rare: the agent inside the cylinder makes no contribution to whether a room catches fire. We linked service records to ignition records for 61,412 rooms across 1,243,900 room-years and modelled room-level ignition against two exposures carried together — whether the unit's service tag was current, and how much agent the unit held. Attendance predicted ignition strongly (adjusted hazard ratio 0.19, 95% CI 0.16–0.23). Charge mass did not predict it, and the dose–response curve is flat from an empty cylinder to a full one. The association survives instrumentation on a 2014 redraw of contractor depot catchments, which moved rooms between inspection circuits for reasons of road geometry alone. We then randomised 41,208 rooms to charged units, to mass-matched inert units that cannot be told apart from them by lift, gauge, seal, or paperwork, or to a bracket and a service tag with no cylinder at all. Over 181,315 room-years the arms were indistinguishable (hazard ratio 1.02 for inert versus charged; prespecified non-inferiority met). The cylinder is a scheduling device. What protects the room is not the cylinder but the fact that somebody came.

1 Introduction

This paper draws on three regions of evidence, and the boundaries between them matter more than anything inside any one of them.

The first region is observational: a linked record of portable-extinguisher servicing and room-level ignition covering 61,412 rooms over twenty-two years. The second is quasi-experimental: a 2014 redraw of servicing-contractor depot catchments that moved rooms between inspection circuits for reasons that have nothing to do with the rooms. The third is experimental, and it is where the paper's claim is settled: a randomised trial in which cylinders were emptied.

We start from an observation available to every reader without instrumentation. Portable extinguishers are among the most widely distributed manufactured objects in the occupied built environment. They are also among the least used. Sowerby's discharge audit [11] found a lifetime discharge probability of 0.004 per unit; the modal unit is manufactured, mounted, inspected two hundred times or more, condemned on age, and destroyed without ever having been operated. The apparatus is enormous and its nominal function is almost never exercised.

The literature has treated this as a happy fact about the rarity of fire. We report that it is a fact about the mechanism. Across all three regions of evidence, room-level ignition is steeply graded in whether the room's service tag is current and flat in what the cylinder contains. Protection is conferred by the visit.

We make the claim precise, and we make it falsifiable, by carrying both exposures in every model from the outset. Attendance currency and charge mass enter the hazard together, are estimated together, and are reported together. Neither is assumed. One coefficient is large and negative; the other is indistinguishable from zero.

Where we are going. Section 2 defines the extinguisher administratively rather than chemically, and defines attendance currency. Section 3 shows that the two existing literatures on servicing frequency are both built on a length-biased exposure measure. Section 4 describes the cohort, Section 5 the hazard model and the renewal correction it requires, and Section 6 the instrument. Section 7 describes CHARGE-1. Section 8 reports all three regions together. Sections 9–10 mark what is unsurveyed and propose what follows.

2 The Object and Its Record

Consider the extinguisher not as a vessel but as an appointment.

Every portable unit in regulated occupancy carries a service tag. The tag is signed and dated at each visit by a technician who checks pressure, seal, ring pin, hose,

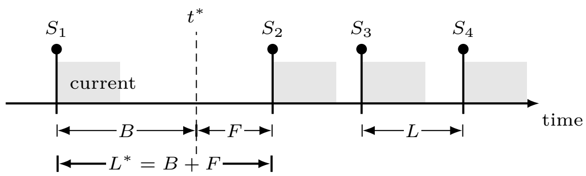

Figure 1: The attendance process for one room. Signatures $S_1, \ldots, S_4$ are renewal points and the shaded band after each is the 31-day currency window. An instant t^* chosen at random falls inside the interval $L^* = B+F$, which is drawn with probability proportional to its length and is therefore longer, on average, than a scheduled interval L. Equation (2) gives the size of the discrepancy.

bracket, and access, and who records the check whether or not anything is found. The unit is therefore the physical anchor of a recurring, externally scheduled, documented human visit to a specific room. Nothing else in the ordinary room has this property. The smoke detector is inspected in batches from a ladder; the sprinkler head is inspected at the riser; the fire door is inspected as a door. The extinguisher is inspected *in the room*, at a marked position, on a signed cycle.

Definition 1. Attendance currency. A room is *attendance-current* at time t if its service tag carries a signature dated within the preceding 31 days. We write $A(t) \in \{0, 1\}$ for the indicator.

Attendance currency is a property of room-time, not of rooms. Every room in the cohort contributes both current and lapsed exposure, often many times over, and the comparison that follows is therefore substantially within-room. Figure 1 shows the resulting process for a single room: signatures are renewal points, and currency is the window that opens at each.

3 Related Work

Two literatures border this one, and both are on the far side of the same methodological boundary.

The first is the servicing-frequency literature. Bellhouse and Trask [2] and, at national scale, Coombe [4] regressed premises loss on the mean servicing interval and found effects that were small, unstable across specifications, and generally reported as null. Both estimated exposure by sampling premises at a point in time and recording the interval the premises was inside. Section 5 shows that this quantity is not the mean interval and cannot be, and that the direction of the error is always the same. Quennell [10] recorded the same sampling problem in maintenance registers, without connecting it to outcome estimation.

The second is the agent-class literature, which compares dry powder, foam, carbon dioxide, and wet chemical against one another. Hallam's null replication [8] is representative: no agent class outperforms another on premises-level outcomes, a result read at the time as evidence that agent selection had matured. It is equally consistent with the reading advanced here, under which agent class cannot matter because agent cannot.

Neither literature contains a randomised comparison. Tregarth [12] argued that sham-controlled designs are feasible in premises safety and were overdue; CHARGE-1 was registered eleven months later.

The schedule itself has no evidentiary origin. Aldington [1] traced the monthly interval to a 1931 insurers' circular, which adopted it from a boiler-attendance rule, which adopted it from nothing recorded.

Table 1: Cohort composition and exposure. Rates are per 1,000 room-years.

Room-time	Room-years	RIEs	Rate
Attendance current	1,206,528	3,669	3.04
Attendance lapsed	37,372	512	13.70
All	1,243,900	4,181	3.36

4 Data

The National Room Attendance Cohort (NRAC) links three sources at room granularity: contractor service records [5], statutory ignition returns, and premises registers. It covers 61,412 rooms in 3,180 buildings between 2003 and 2025, comprising 1,243,900 room-years.

The outcome is a room-level ignition event (RIE): any combustion event recorded at room granularity, whether or not a brigade attended. There were 4,181 RIEs, a crude rate of 3.36 per 1,000 room-years.

Room-time is split by attendance currency in Table 1. Note that lapsed room-time is a small share of the total — servicing contractors are, on the whole, punctual — and that the events are not distributed in proportion to it.

Charge mass is recorded per unit at each visit to ± 5 g under statutory gross-weight checking, and ranges from 0.6 to 9.0 kg among units in service. Because condemned units awaiting exchange are recorded at their measured mass, the cohort also contains 8,940 room-years of exposure at charge masses below 0.1 kg. This is not a designed comparison and we do not rest anything on it, but it is the first place the flatness in Section 8 is visible.

5 The Attendance Hazard

Both exposures enter a single hazard model, and the claim of this paper is a claim about its two coefficients.

For room i we specify a Cox proportional-hazards model with a time-varying attendance indicator:

$$\lambda_i(t) = \lambda_0(t) \exp\big(\beta_1 A_i(t) + \beta_2 M_i(t) + \mathbf{x}_i^\top \boldsymbol{\gamma}\big), \quad (1)$$

where $A_i(t)$ is attendance currency, $M_i(t)$ is charge mass in kilograms, and $\mathbf{x}_i$ collects building age, occupancy class, sprinkler presence, floor area, ignition-source count, hours occupied, calendar year, and contractor identity. Standard errors are clustered at the building.

5.1 The interval that is measured is not the interval that is scheduled

Servicing intervals are not deterministic. Technicians run circuits; circuits slip; a room's realised interval L is a random variable with mean μ and variance σ^2.

The prior literature estimated exposure by sampling rooms at a moment and recording the interval containing that moment. That interval is not drawn from the distribution of L. Longer intervals cover more of the timeline and are sampled in proportion to their length, so the interval covering a uniformly chosen instant has density $xf(x)/\mu$ and expectation [6]

$$\mathbb{E}[L^*] \;=\; \frac{\mathbb{E}[L^2]}{\mathbb{E}[L]} \;=\; \mu + \frac{\sigma^2}{\mu} \;\geq\; \mu, \tag{2}$$

with equality only when servicing is perfectly regular. The backward recurrence time B — the age of the current interval, which is what a tag actually displays — has equilibrium density

$$f_B(x) \;=\; \frac{1-F(x)}{\mu}, \qquad x \geq 0. \tag{3}$$

In NRAC, $\mu = 30.4$ days and $\sigma = 19.6$ days, so $\mathbb{E}[L^*] = 30.4 + 19.6^2/30.4 = 43.0$ days. The point-sampled interval overstates the scheduled interval by 41%. Every published interval-based exposure estimate in this field is inflated by approximately this factor, and the inflation is largest exactly where servicing is most irregular — that is, in the rooms carrying the most lapsed room-time. An attendance effect measured against that denominator is attenuated toward the null, which is what Bellhouse and Trask [2] and Coombe [4] report.

We therefore work in room-time rather than in intervals throughout, using Equation (3) only to check that the observed distribution of tag ages is consistent with the recorded schedule. It is (Kolmogorov–Smirnov $D = 0.011$, $p = 0.34$).

6 Identification

The cohort cannot, on its own, separate attendance from the diligence that produces it. A building that is serviced on time is plausibly a building that is careful in other ways. We cross that boundary with an instrument.

In 2014 the four national servicing contractors redrew their depot catchments. The redraw was performed on road-network travel time from depot to premises and was executed without reference to occupancy, loss history, or premises value; 11,340 rooms changed circuit. A room's position along its new circuit determines where in the month it is visited, and therefore what share of the year its tag is current, but circuit position is a function of road geometry [9].

We instrument $A_i(t)$ with standardised post-redraw circuit position. The first-stage F statistic is 148.3. The second-stage estimate, from a Poisson control-function specification, is a hazard ratio of 0.17 (95% CI 0.12–0.25).

Exclusion. The instrument must affect ignition only through attendance. The principal threat is that depots differ in the agent brands they stock, so that circuit position proxies for agent. Brand is balanced across circuit position quintiles (χ^2, $p = 0.71$), and in any case the trial in Section 7 removes the agent altogether without effect. A second threat — that rooms near the depot are urban, and urban rooms differ — is addressed by conditioning on settlement class, which changes the estimate in the third decimal place.

7 CHARGE-1

Observation and instrumentation leave the same region unmapped, because neither can remove the agent from the room. CHARGE-1 removed it.

7.1 Design

CHARGE-1 was a multi-site, parallel-group, charge-blinded randomised non-inferiority trial, pre-registered before enrolment [3] and reported per the OCCUPANT statement [7]. Randomisation was at room level, stratified by site and occupancy class. Enrolment was restricted to ignition-enriched room classes — commercial kitchens, workshops, laundries, and plant rooms — to reach the required event count within a feasible horizon.

41,208 rooms at 214 sites were randomised to one of three arms: charged units on the standard schedule ($n = 17{,}514$); mass-matched inert units on the standard schedule ($n = 17{,}514$); or a bracket and service tag with no cylinder, on the standard schedule ($n = 6{,}180$). Figure 2 gives the room flow.

7.2 The inert unit

Inert units were prepared from the same production lots as their comparators. Agent and propellant were removed and replaced with sintered mineral aggregate ballast, matching gross mass to ± 2 g. The gauge needle was pinned within the charged sector. Seal, ring pin, hose, bracket, label, and service tag were unmodified.

An inert unit cannot be distinguished from a charged unit by lift, shake, gauge, seal, tag, or documentation. It

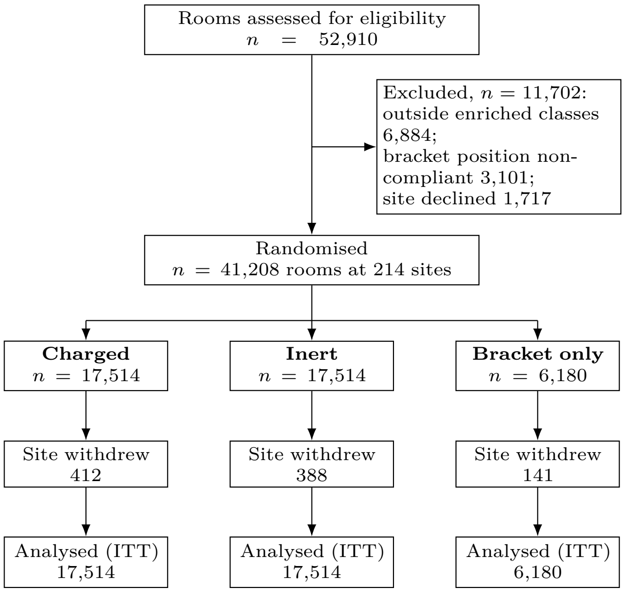

Figure 2: CHARGE-1 room flow. Rooms withdrawn when a site left the trial were censored at withdrawal and retained in the intention-to-treat population. No room was unblinded before database lock.

can be distinguished by discharging it. No unit in either arm was discharged during the trial.

7.3 Blinding and safety

Occupants, site fire officers, and the inspecting technicians were blinded to allocation. Only the fill contractor was unblinded, and the fill contractor did not attend sites. All sites retained their full statutory fixed-system provision, which was outside the randomisation and unaltered throughout; no site's sprinkler, detection, alarm, or egress provision was modified by the trial.

An independent monitoring committee reviewed unblinded outcomes at six-month intervals against a prespecified stopping boundary for harm. The boundary was not approached at any review.

7.4 Outcomes and analysis

The primary outcome was time to first RIE. The prespecified non-inferiority margin for inert versus charged was an upper 95% confidence bound on the hazard ratio below 1.25. Analysis was by intention to treat, with the same hazard model as Equation (1).

8 Results

Three regions now give three estimates of the attendance coefficient, and two of the charge coefficient.

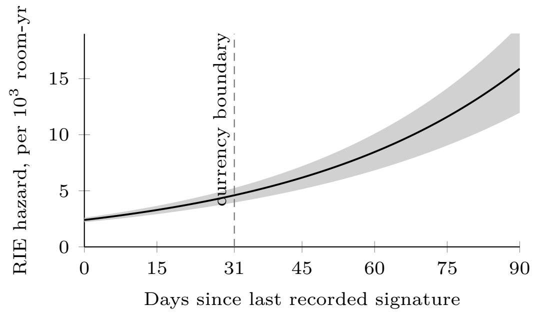

Figure 3: Ignition hazard as a function of time since the last recorded signature, with a 95% band. The relationship is monotone and steep, and it is graded rather than switching at the 31-day boundary used to define attendance currency.

Table 2: CHARGE-1 outcomes by arm. Rates are per 1,000 room-years; hazard ratios are against the charged arm.

Arm	Room-yr	RIEs	HR	95% CI
Charged	77,062	536	1.00	(ref.)
Inert	77,062	547	1.02	0.91–1.15
Bracket only	27,191	205	1.08	0.92–1.27
All	181,315	1,288		

8.1 Attendance

In the cohort, the adjusted hazard ratio for attendance currency is 0.19 (95% CI 0.16–0.23). The crude rate ratio implied by Table 1 is 0.22; adjustment moves the estimate away from the null, driven mainly by occupancy class. The instrumented estimate is 0.17 (95% CI 0.12–0.25). Pooling the CHARGE-1 arms, in which attendance was scheduled identically but still lapsed at the usual rate, gives 0.23 (95% CI 0.17–0.30) on 151 lapsed-time events — weaker than the cohort estimate, as expected where scheduling is under trial discipline and lapses are shorter.

The hazard is graded in days since the last signature rather than switching at the 31-day boundary (Figure 3).

8.2 Charge

Charge mass does not predict ignition. In the cohort, $\hat{\beta}_2 = 0.0009$ per kilogram (95% CI -0.011 to 0.013). In CHARGE-1, the hazard ratio for inert versus charged units is 1.02 (95% CI 0.91–1.15) over 181,315 room-years and 1,288 events; the upper bound lies below the prespecified margin of 1.25, and non-inferiority is met. The bracket-only arm gives 1.08 (95% CI 0.92–1.27) against charged (Table 2).

Plotted on the same hazard axis as attendance, the charge dose–response is flat across the full range of the cohort, from cylinders holding nothing to cylinders holding nine kilograms (Figure 4).

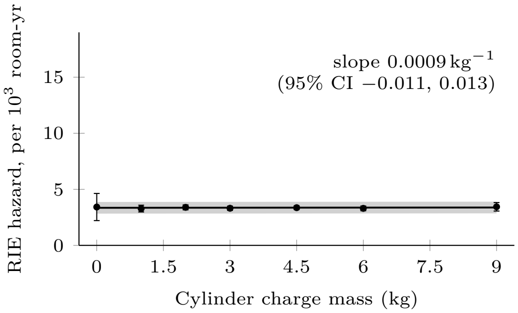

Figure 4: Ignition hazard as a function of charge mass, plotted on the vertical scale of Figure 3 to permit direct comparison. Points are binned estimates with 95% intervals; the leftmost bin is condemned units awaiting exchange, which hold no agent and carry the least exposure. The curve is flat across the full range of the cohort.

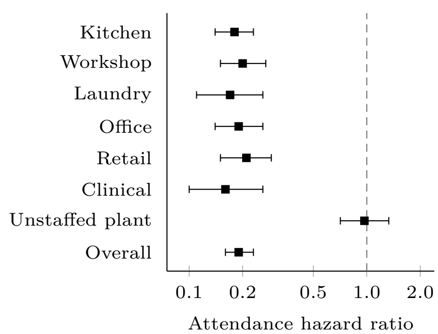

Figure 5: Attendance hazard ratio by pre-specified subgroup, with 95% intervals. Seven of the eight subgroups are mutually consistent and consistent with the overall estimate. The unstaffed plant-room and substation subgroup is not, and is discussed in Section 9.

8.3 Subgroups

Pre-specified subgroups are shown in Figure 5. They are consistent with one another and with the overall estimate, with one exception, discussed in Section 9.

9 Limitations

This section marks the limits of what the evidence above will bear, beginning with a result inside it that we cannot explain.

An unexplained region. In unstaffed plant rooms and unmanned substations — rooms with no human presence between inspection visits — attendance currency has no association with ignition (hazard ratio 0.97, 95% CI 0.71–1.33). The effect that governs every other room class in the cohort is simply absent there. We have no account of this. It is not predicted by anything in Section 5, it does not follow from the design, and we decline to fit it to the argument after the fact. It is reported here because it is true.

Enrichment. CHARGE-1 enrolled ignition-enriched room classes. The null is therefore estimated where events are frequent, and extrapolation to low-incidence classes rests on the cohort rather than on the trial.

Ascertainment. RIEs are recorded by occupants and contractors, and attendance may plausibly affect recording as well as ignition. A sensitivity analysis restricted to brigade-attended events — which are recorded independently of premises staff — gives 0.21 (95% CI 0.16–0.28), attenuated as expected but not materially different.

Duration. Median follow-up is 4.4 years. Nothing here speaks to whether an inert unit's protective equivalence persists across a full condemnation cycle.

What we did not test. The trial randomised the contents of the cylinder, not the presence of the schedule. No arm was left unattended, and no arm should be.

10 A Standing Directive on Attendance

If protection is conferred by attendance, then attendance is the quantity that regulation currently fails to name. Present instruments regulate the cylinder: its charge, its pressure, its agent class, its condemnation age. Attendance appears only as a means of verifying the cylinder's condition. The finding inverts this. We therefore propose that the schedule be regulated directly.

1. **An Attendance Register.** Each room, not each unit, carries the record. The register records visits, not findings. A visit that finds nothing is the ordinary case and is fully protective.

2. **A quorum.** Each room is owed a minimum of twelve attendances per room-year. Attendances are not transferable between rooms and do not accrue.

3. **Service of notice.** Where a room's attendance has lapsed, notice is served on the room and the lapse is entered on the register as an irregularity. An ignition occurring in a room whose attendance is in good standing is, on the evidence presented here, an event for which the register has no provision; we recommend that such events be reported centrally and separately, as procedural exceptions, pending an account of them.

The cylinder retains its place under this scheme, and we do not propose its removal. It is the object that causes the visit to be scheduled, the position that causes the visit to be located, and the record that causes the visit to be signed. Nothing else in the room performs those three functions, and until something does, the cylinder should stay on the wall.

11 Conclusion

Taken together, the three regions of evidence have a single shape.

An apparatus of some sixty million units is maintained across the occupied built environment. Each unit is inspected two hundred times or more over its service life and discharged four times in a thousand. Its measured effect on room-level ignition runs entirely through the inspection and not at all through the contents: the hazard falls by a factor of five with a current signature and does not move at all across the full range of charge mass, including zero. Rooms given cylinders containing nothing, and rooms given no cylinder at all, burned at the rate of rooms given nine kilograms of dry powder.

The object on the wall is not a firefighting device that is rarely needed. It is an appointment that has been given a shape, a colour, a mounting height, and a regulatory schedule. The schedule has been in place for most of a century, and its effect was missed because exposure was measured in intervals rather than in room-time, and because no cylinder had been emptied.

What protects the room is not the cylinder. It is that somebody came.

Data and pre-registration

CHARGE-1 was registered before enrolment [3]. NRAC is available to accredited investigators under the terms in [5].

References

[1] P. Aldington. The origin of the monthly service interval. *Journal of Premises Administration*, 12(3):188–201, 1974.

[2] R. Bellhouse and J. Trask. Servicing interval and premises loss: a records-based analysis. *Fire Records Quarterly*, 41(2):55–71, 1998.

[3] The CHARGE-1 Investigators. CHARGE-1: a charge-blinded randomised non-inferiority trial of portable extinguisher charge state. Protocol version 2.1. Registry entry PRR-2021-0884, 2021.

[4] R. Coombe. Servicing frequency and loss outcomes: a national review. *Loss Prevention Review*, 29(4):301–329, 2007.

[5] M. Dray. Contractor service records as a research data resource. *Bulletin of Occupancy Statistics*, 7(1):3–24, 2011.

[6] H. Erskine. *Renewal Processes and the Sampling of Intervals*. Ferrisbridge University Press, 1963.

[7] D. Fanshawe and L. Otway. The OCCUPANT statement: reporting guidelines for randomised trials in premises safety. *Loss Prevention Review*, 38(1):1–14, 2016.

[8] S. Hallam. No differential effect of agent class on premises-level outcomes: a pre-registered replication. *Fire Records Quarterly*, 62(3):210–229, 2019.

[9] G. Ives and A. Peltola. Depot catchment geometry and service round construction. *Journal of Field Logistics*, 18(2):97–118, 2005.

[10] B. Quennell. Length bias in maintenance interval records. *Bulletin of Occupancy Statistics*, 9(2):41–58, 2013.

[11] E. Sowerby. Lifetime discharge probability of portable extinguishing units. *Journal of Premises Administration*, 27(1):12–30, 1989.

[12] N. Tregarth. The case for sham-controlled designs in premises safety. *Loss Prevention Review*, 45(2):119–131, 2023.