

Two Minutes Is Not a Duration: The Announced Remainder as a Denomination of Outstanding Work, with a Discharge Ratio Estimated from 1,904,772 Countdowns

Bodil Ravnsborg¹ and Wenying Tao²

¹Office of Weights and Measures, Aalborg

²Institute for Service-Time Statistics, Leeds

Abstract

The washing machine says two minutes. Ten minutes later it still says two minutes, and the machine has not made a mistake. An announced remaining time is not a prediction of elapsed time at all, but an exact statement of outstanding work, denominated in minutes and not measured in them. On this reading the residual that the estimation literature has spent eighty years trying to eliminate is not error but an exchange rate. Instrumenting eleven classes of domestic appliance and timing completion from the mains current trace rather than from the display, we record 1,904,772 countdowns and estimate a single discharge ratio, $\delta = 3.72$ (95% CI 3.68–3.78): one announced minute redeems for δ minutes of clock. The ratio is invariant across manufacturer, task and operator, and across ten of the eleven appliance classes; it survives a concealed-announcement trial in which the display is suppressed altogether; and it holds unchanged for contact-centre wait announcements, construction contracts, and the intraoperative assurance that the procedure is nearly over. Two consequences follow. Completion is uniform over the discharge window, which is the exact content of the phrase *any minute now*; and the folk instruction to double a stated estimate is not folklore but a rounding of $\delta/2$. The estimate was never wrong. The unit was.

Keywords: announced remainder; discharge ratio; denomination; uniform completion; estimation residual

1 Introduction

A machine that announces a remaining time will overrun it. This is among the most widely replicated findings in applied psychology and among the least useful, because eight decades of replication have produced no correction. The literature’s explanation is that estimators are optimistic (Brandt, 1946; Østergaard, 1974; Shim, 2003). If that explanation were right, feedback would

have eroded the error. A householder observes the same appliance several thousand times. A contractor prices several hundred jobs. Neither converges. A bias that survives unlimited feedback at zero cost of observation is not a bias.

The question is therefore not how to estimate better. It is what the display is reporting. Three readings are available, and two of them are traps.

Reading one: the display predicts elapsed time and is wrong. This is the standard reading. It requires the residual $W - R$ to shrink under feedback, which it does not, and it requires the residual to depend on the difficulty of the particular task. It does not.

Reading two: the display predicts elapsed time under idealised conditions and the world intrudes. This reading survives the first objection by making the residual a sum of unmodelled obstacles. It fails a sharper test: obstacles do not compound into a ratio that is constant to three significant figures across tumble dryers, telephone queues, and load-bearing masonry.

Reading three: the display is exact, and the reader is mis-unitising it. Announcement and elapsed time are not commensurable. The announcement quantifies outstanding work; the clock quantifies time; the number relating them is a conversion factor, not an error term. On this reading the estimation literature has spent eighty years compiling a currency table without noticing.

We adopt reading three and take it literally. Section 3 defines the unit — the *discharge-minute*, dm — and derives the law of completion from two axioms. Section 4 describes an instrumented panel in which announcement and completion are measured by independent apparatus, the display never being used to time itself. Sections 5 and 6 estimate the discharge ratio δ , establish its invariance across substrate, and report a randomised trial in which the announcement is concealed. Section 7 states what the result costs the bias literature, and reports one anomaly we cannot account for.

Contributions. We (i) define the discharge-minute and give the conversion $1 \text{ dm} = \delta \text{ minutes}$; (ii) prove that completion is uniform on the discharge window, and that the hazard of completion diverges at its edge; (iii) estimate δ from 1,904,772 instrumented countdowns and show it invariant across ten of the eleven appliance classes and four institutional substrates, the exception being reported and left unexplained; (iv) demonstrate by randomisation that the announcement is causally inert with respect to the ratio; and (v) show that reference-class forecasting, the standard remedy, buys nothing over multiplication by $\delta/2$.

2 Related Work

The estimation residual. Brandt (1946) first quantified the gap between announced and realised completion in industrial scheduling, and attributed it to the planner’s disposition. The attribution has proved durable (Østergaard, 1974; Shim, 2003; Delacroix et al., 2016). Reference-class forecasting (Shim, 2003) is the standard remedy: replace the plan with the outcome distribution of comparable prior plans. It is a good procedure applied to the wrong object. We show in §6 that its advantage over a fixed multiplier does not grow with the size of the reference class, which is not what a debiasing procedure does.

Announced delay. The queueing literature has long studied announced waiting times as a control input, asking whether the announcement changes the customer’s behaviour (Quillan, 1998; Marchetti and Osei, 2011). That work measures the announcement’s effect on abandonment; it does not ask what the announcement is a measurement *of*. Marchetti and Osei (2011) report a mean realised-to-announced ratio of 3.7 across nineteen contact centres and treat it as a calibration defect specific to their sample. To our knowledge it is the earliest published estimate of δ .

Units and their definition. That a quantity may be reported in a unit other than the one its reader assumes is ordinary in metrology (Fournier, 1954; Abegunde and Roth, 2019). The nautical mile, the troy ounce and the imperial gallon all survive as denominations whose numerals invite the wrong conversion. Fournier (1954) sets out the criteria under which a reported figure should be treated as a denomination rather than a measurement; the announced remainder satisfies all four.

3 The Discharge Model

Definition 1 (Announced remainder). At the instant of announcement a device or agent emits a numeral

$R > 0$. We take R to be a quantity of outstanding work expressed in *discharge-minutes* (dm), a unit distinct from the minute of time.

Definition 2 (Discharge window). The *discharge window* is the interval $(0, \Lambda]$ of wall-clock time within which the outstanding work is redeemed, where $\Lambda = \delta R$ and δ carries units min dm^{-1} .

Figure 2 shows the geometry.

Axiom 1 (Denomination). R is exact. The announcement contains no error term.

Axiom 2 (Uniform discharge). Conditional on R , the realised completion time W satisfies $W \mid R \sim \text{Unif}(0, \delta R]$.

By Axiom 1 the announcement contributes no error term, so the only constraint available on the law of W is its support. Axiom 2 is therefore not an assumption of convenience. It is forced (Halloway, 1979).

Theorem 1 (Idiom exactness). *Let W be supported on $(0, \Lambda]$ and let R be a sufficient statistic for W , so that the announcement conveys the extent of the window and nothing further. Then the distribution of W maximising differential entropy is uniform on $(0, \Lambda]$, and it is the unique such maximiser.*

Proof. Let f be any density on $(0, \Lambda]$ and $u = 1/\Lambda$ the uniform density. Then

$$h(u) - h(f) = \int_0^\Lambda f(w) \log \frac{f(w)}{u(w)} dw = D_{\text{KL}}(f \| u) \geq 0,$$

with equality if and only if $f = u$ almost everywhere. Sufficiency of R forbids any constraint on f beyond its support, so no further moment condition is available to displace the maximiser. $\square$

Theorem 1 has a consequence outside mathematics. The vernacular formula for an imminent completion of unknown timing is *any minute now*. Under Axiom 2 the phrase is not a hedge and not an approximation: it asserts that each remaining minute of the window carries equal probability, which is precisely the content of the uniform law. The usage is attested from 1748 (Merrivale, 1994), predating the formal treatment of the distribution it names.

Proposition 1 (Hazard). *Under Axiom 2 the survival function is $S(w) = (\Lambda - w)/\Lambda$ and the hazard is*

$$h(w) = \frac{f(w)}{S(w)} = \frac{1}{\Lambda - w}, \quad 0 < w < \Lambda, \quad (1)$$

which diverges as $w \rightarrow \Lambda^-$. The expected residual satisfies $\mathbb{E}[\Lambda - W \mid W > w] = (\Lambda - w)/2$.




Figure 1: The hazard of completion, (1), for three announced remainders. Each curve diverges at its own window edge $\Lambda = \delta R$.

Equation (1) is the formal content of impatience, and Figure 1 plots it for three announced remainders. The instantaneous rate of completion grows without bound as the window closes, while the expected wait remaining is always half of what is left. Both statements are true simultaneously, and the second is what the observer feels.

Proposition 2 (On-time probability). *The probability that completion occurs within R minutes of the clock — that is, that the display appears to have been right — is*

$$\Pr[W \leq R] = \frac{R}{\delta R} = \delta^{-1}. \quad (2)$$

At $\delta = 3.72$ this is 0.269. Roughly one countdown in four completes inside its announced numeral, which is the observed rate and which no account based on systematic optimism predicts.

Theorem 2 (Conversion is linear). *Let δ be the unique dimensional constant linking dm to min, and let no further dimensional quantity enter the relation between R and Λ . Then $\Pi = \Lambda/(\delta R)$ is the sole dimensionless group, and the Buckingham π -theorem requires Π to be a pure number: $\Lambda = c\delta R$. Normalising $c = 1$ fixes the scale of δ . No nonlinear correction in R is admissible.*

Theorem 2 is what distinguishes the present account from reading two of §1. A sum of unmodelled obstacles would introduce a further dimensional quantity — a rate, a scale, a difficulty — and would license curvature of Λ in R . None is observed (§6).

4 Data and Instrumentation

4.1 Study 1: the domestic panel

We instrumented 1412 households across four countries between March 2019 and November 2025. Two independent instruments were used, and the display was never permitted to time itself.




Figure 2: The discharge window, in wall-clock time w . The announcement fixes $\Lambda = \delta R$; completion W falls uniformly within it. The hatched region is the portion of the window that lies inside the announced numeral, so that a fraction δ^{-1} of countdowns appear, to a reader converting at unity, to have been correct.

Completion was taken from a clamp-on Hall-effect current logger at the consumer unit, sampling at 50 Hz, and cycle boundaries were scored by the non-intrusive load-monitoring procedure of Anwar and Leconte (2017). A cycle was scored complete at the last sample exceeding 4% of the cycle’s peak draw, a threshold fixed by pilot data before the main collection. The logger has no access to the panel and cannot be influenced by it.

Announcement was taken from a fixed camera imaging the appliance panel at 1 Hz, with the numeral recovered by optical character recognition. Against 9600 manually transcribed frames the recogniser achieved 99.94% agreement; disagreements were resolved to the manual reading.

Eleven appliance classes contributed: washing machine, tumble dryer, dishwasher, combination oven, bread maker, electric pressure cooker, ultrasonic cleaner, domestic 3D printer, air purifier cycle timer, water softener regeneration controller, and microwave oven. After the exclusions below, 1,904,772 countdowns entered the analysis.

4.2 Exclusions

Exclusions were fixed in a pre-registered protocol. We removed cycles interrupted by the operator (31,204), cycles subject to a supply interruption (1877), and cycles in which the panel was obscured for more than 10% of frames (8043). We did *not* exclude cycles that completed early; their retention is essential, since a design that records only overruns would recover δ from a truncated sample and inflate it.

4.3 Study 2: the concealed-announcement trial

To test whether the announcement causes the ratio it reports, we obtained firmware cooperation from three manufacturers and randomised 4190 machines

into three arms. ARM A displayed the remainder normally. ARM B suppressed the display entirely, substituting a rotating token that carried no numeral. ARM C displayed a sham remainder drawn from the empirical distribution of a different machine of the same class. In all three arms the internal remainder was logged over the service bus.

4.4 Study 3: substrate invariance

Four institutional substrates were added: 39 contact centres (17.2 million announced waits); 214 construction contracts with announced completion dates; 7910 time-stamped intraoperative utterances of the form *nearly done*, recovered from theatre audio under existing recording consent; and 2301 software installation progress indicators instrumented at the operating-system level.

4.5 Study 4: reference-class calibration

On held-out countdowns we ran three forecasters: (a) trust the announcement; (b) reference-class forecasting trained on the household’s own prior cycles for that appliance (Shim, 2003); (c) multiply the announcement by $\delta/2$. Forecasters were scored by mean absolute error and by pinball loss at $\tau = 0.9$, the quantile at which a uniform law and a point forecast diverge most sharply, and re-scored as a function of the number of prior cycles available for training. The mean announced remainder in the held-out set was 20.0 dm.

4.6 Study 5: metrological audit

The timebase of the current loggers and of the camera trigger was referred to a national standards laboratory and compared against a caesium reference. The purpose was elimination, not confirmation: if the wall-clock leg of the ratio were itself in error, δ would be an artefact of the instrument rather than a property of the unit.

5 Estimation

Write $U_i = W_i/R_i$. Under Axiom 2, $U_i \sim \text{Unif}(0, \delta]$. The likelihood is maximised at $\hat{\delta}_{\text{ML}} = \max_i U_i$, which is consistent but is determined entirely by the sample extremum and is unusable in the presence of any surviving interrupted cycle. We therefore estimate δ by the method of moments,

$$\hat{\delta} = \frac{2}{n} \sum_{i=1}^n U_i, \quad \text{Var}(\hat{\delta}) = \frac{\delta^2}{3n}, \quad (3)$$

which is unbiased and, at the $n = 1,722,647$ countdowns entering the pooled estimate, has a nominal standard error of 0.0016.

Table 1: Discharge ratio by appliance class, Study 1. D is the Kolmogorov–Smirnov statistic of W/Λ against uniformity.

Class	n	$\hat{\delta}$	95% CI	D
Washing machine	412,908	3.74	3.69–3.80	0.0014
Dishwasher	356,022	3.71	3.66–3.77	0.0014
Tumble dryer	288,140	3.69	3.63–3.75	0.0016
Combination oven	191,447	3.75	3.68–3.82	0.0022
Pressure cooker	118,690	3.73	3.65–3.81	0.0023
Air purifier timer	96,551	3.72	3.63–3.81	0.0034
Water softener	88,213	3.76	3.67–3.86	0.0031
Bread maker	74,315	3.68	3.58–3.78	0.0037
Domestic 3D printer	52,384	3.70	3.58–3.82	0.0042
Ultrasonic cleaner	43,977	3.77	3.64–3.90	0.0042
Microwave oven	182,125	1.00	0.99–1.01	0.0023
Pooled (excl. microwave)	1,722,647	3.72	3.68–3.78	0.00070

That standard error is not what we report. Households differ, and the sampling error at this n is negligible beside the between-household dispersion. We fit a two-level model with household and class random effects,

$$U_{ijc} \sim \text{Unif}(0, \delta_{jc}], \quad \log \delta_{jc} = \log \delta + b_c + e_j, \quad (4)$$

with $b_c \sim \mathcal{N}(0, \tau_c^2)$ across classes and $e_j \sim \mathcal{N}(0, \tau_h^2)$ across households, and report the posterior interval for δ . All intervals below are of this kind.

6 Results

6.1 The discharge ratio

The pooled estimate is $\hat{\delta} = 3.72$, 95% CI 3.68–3.78. Estimates by appliance class are given in Table 1. Ten of the eleven classes are mutually compatible; the eleventh is the microwave oven, discussed in §7.

The variance decomposition of (4) attributes 0.4% of total variation to appliance class (microwave excluded) and 2.1% to household, the remainder being within-household. Whatever δ is a property of, it is not a property of the machine, the task, or the operator.

6.2 Uniformity

Figure 3 plots the departure of the empirical distribution of W/Λ from the uniform diagonal. Plotting the deviation rather than the distribution itself is necessary: at this scale the two curves are indistinguishable when superimposed. A Kolmogorov–Smirnov test against uniformity returns $D = 0.00070$ pooled, against a 5% critical value of 0.00104. At $n = 1,722,647$ the test has power 0.9 against a departure of $D = 0.0016$; it detects none.




Figure 3: Departure from uniformity of the discharge fraction W/Λ , normalised by $\sqrt{n}$ so that one band serves every class. The shaded region is the 95% Kolmogorov–Smirnov envelope at ± 1.358 . Four representative classes are shown together with the pooled curve (heavy line); no class leaves the envelope.

Table 2: Concealed-announcement trial (Study 2). In every arm δ is computed against the machine’s internal remainder, logged over the service bus.

Arm	Panel shows	machines	$\hat{\delta}$	Δ vs. A (95% CI)
A	true remainder	1397	3.72	—
B	rotating token	1401	3.71	−0.01 (−0.09, 0.07)
C	sham remainder	1392	3.73	+0.01 (−0.07, 0.10)

6.3 Invariance across substrate

Figure 4 shows $\hat{\delta}$ for every substrate. The contact centres give 3.71; construction contracts 3.74; intraoperative utterances 3.69; software installers 3.75. None differs from the domestic pooled estimate at the 5% level. The masonry and the tumble dryer agree to two significant figures, and we can offer no property they share except the formation of a remainder. Study 2 shows that announcing it is optional.

6.4 The announcement is inert

Table 2 reports the concealed-announcement trial. $\hat{\delta}$ was 3.72 in ARM A, 3.71 in ARM B, and 3.73 in ARM C. Suppressing the display changes nothing; falsifying it changes nothing. The announcement neither produces the ratio nor is corrupted by being seen. This eliminates the remaining behavioural reading, on which the numeral functions as a target that the operator or the machine then fails to meet.

6.5 Reference-class forecasting buys nothing

Forecaster (c) — multiply by $\delta/2$ — achieved a mean absolute error of 18.6 min, which is the floor $\Lambda/4$ that




Figure 4: Discharge ratio by substrate. Rows 1–10 are the conforming domestic appliance classes; rows 11–14 are institutional substrates — contact-centre wait announcements, construction contracts, intraoperative “nearly done” utterances, and software installers — in which the announcement is made by a person or an organisation rather than a panel. Bars are 95% intervals from the hierarchical model of (4). The dashed line is the pooled domestic estimate. The microwave oven is shown last and is discussed in §7.

Axiom 2 implies. Reference-class forecasting achieved 19.4 min, and trusting the announcement 22.6 min. The decisive result is not the ordering but its behaviour as the training set grows: forecaster (b) is flat in the number of prior cycles from the tenth cycle onward (Figure 5).

A debiasing procedure improves with evidence. This one does not, because there is no bias to remove. There is a conversion factor, and it is already known.

6.6 The clock is exact

The audit returned agreement with the caesium reference to 1 part in 10^9 over the collection window. The wall-clock leg of the ratio is therefore not an instrument artefact. The residual is in the announcement, and Studies 2 and 4 exclude the two readings that would keep it there as error.

7 Discussion

7.1 What this costs the bias literature

If δ is a conversion factor, then the planning fallacy is not a fallacy. The estimator who says two weeks and takes four has reported the work correctly in dm and




Figure 5: Forecast error against the size of the reference class (Study 4). Reference-class forecasting improves until roughly the tenth prior cycle and then stops, settling above a fixed multiplier that uses no history at all and sits on the floor $\Lambda/4$. A debiasing procedure does not have this shape.

been read in minutes by an audience that never checked the unit. The remedy is not training, and not humility. It is a conversion table.

This creates a practical asymmetry. An estimator who has internalised δ and pre-multiplies before speaking is now reporting in minutes while being read in minutes, and is correct — but is also, by the same act, no longer reporting the work. The two conventions cannot coexist in one organisation, which is the strongest argument for standardising rather than for individual correction.

7.2 An anomaly we cannot explain

The microwave oven does not obey the model. Its estimate is $\hat{\delta} = 1.00$, 95% CI 0.99–1.01: one announced minute redeems for one minute of clock. We have measured no property of the microwave — draw, duty cycle, control architecture, manufacturer, panel type, the presence or absence of a mechanical dial — that distinguishes it from the ten classes that conform. The instrument is the same instrument. The anomaly is reported here and is not accounted for by anything in this paper.

7.3 Scope

Satellite navigation systems are outside the model’s scope, and correctly so: they revise the announcement continuously against observed progress, so that each successive announcement is conditioned on the discharge already achieved. A construction date, by contrast, is announced once and never re-derived. Their exclusion is a restriction on where the account applies, not evidence for it. Any device that re-derives its announcement should be expected to conform to the clock, and

if such a device were nevertheless found to overrun by δ , the model would be in serious trouble.

7.4 What would falsify δ

Three observations would do it: a substrate in which $\hat{\delta}$ depends on the magnitude of R , which Theorem 2 forbids; a population in which reference-class forecasting overtakes the fixed multiplier and continues to improve with history; or a drift in $\hat{\delta}$ over calendar time, which would make it a habit rather than a unit. Our series covers nearly seven years and shows no drift beyond sampling error.

8 Conclusion

There is one recommendation, and it is administrative. Announced remainders should be relabelled in discharge-minutes at the point of manufacture, and the conversion $1 \text{ dm} = 3.72 \text{ minutes}$ should be carried on the appliance in the manner of a voltage rating; the drafting conventions already exist (Bureau of Legal Denominations, 2021). Nothing about the machine needs to change. Only the label does.

The risk is worth naming. Standardisation invites the manufacturer to compete on the numeral rather than on the cycle, and a fleet that reports in dm while a rival reports in minutes will appear slower while being identical. That failure mode is familiar from every previous unit reform and is soluble by the same means: dual marking, mandated, for one product generation.

The signal to watch is the microwave. It is the one device that already discharges at unity, and if further classes migrate toward $\delta = 1$ without a change in their mechanism, then δ is not a unit after all and this account should be abandoned rather than patched.

Until then: the estimate was never wrong. The unit was.

A The forecast-error floor

The three forecasters of §6 are separated by an amount that Axiom 2 fixes in advance, so we record the derivation rather than the fitted numbers.

Let $W \sim \text{Unif}(0, \Lambda]$ with $\Lambda = \delta R$. For a point forecast $\hat{w}$ the mean absolute error is $\mathbb{E}|W - \hat{w}|$, minimised at the median $\Lambda/2$.

Forecaster (c), $\hat{w} = \delta R/2$. At the median,

$$\mathbb{E}\left|W - \frac{\Lambda}{2}\right| = \frac{1}{\Lambda} \int_0^{\Lambda} \left|w - \frac{\Lambda}{2}\right| dw = \frac{\Lambda}{4} = \frac{\delta}{4} R, \quad (5)$$

which is $0.93 R$ at $\delta = 3.72$. No forecaster using R alone can do better; (5) is a floor, not a fitted value.

Forecaster (a), $\hat{w} = R$. Splitting at R ,

$$\mathbb{E}|W - R| = \frac{1}{\Lambda} \left[\frac{R^2}{2} + \frac{(\Lambda - R)^2}{2} \right] = \frac{1 + (\delta - 1)^2}{2\delta} R, \quad (6)$$

which is $1.129 R$. Dividing (6) by (5),

$$\frac{\text{MAE}(a)}{\text{MAE}(c)} = \frac{2(1 + (\delta - 1)^2)}{\delta^2} = 1.214, \quad (7)$$

independently of the distribution of R . The observed ratio, $22.6/18.6 = 1.215$, is therefore not evidence for the model: it is a consistency check on the instrumentation, and it passes.

Forecaster (b). Reference-class forecasting estimates the conditional median of W given R and the household's history. Under Axiom 2 that median is $\Lambda/2$, so (b) is a consistent but noisy estimator of the quantity (c) already knows. Its excess over the floor is estimation variance and decays as the reference class grows; its *limit* is the floor, which it approaches from above and cannot cross. Figure 5 shows the approach stopping at 19.4 min, 0.8 min above (5), consistent with the residual dispersion of R within a household's own history.

References

- O. Abegunde and L. Roth. Denomination drift and the reader's default conversion. *Metrologia Applicata*, 44(2):111–130, 2019.
- P. Anwar and D. Leconte. Non-intrusive load monitoring for cycle-boundary detection. *Transactions on Domestic Instrumentation*, 9(4):401–423, 2017.
- H. Brandt. Scheduled and realised completion in twelve engineering works. *Journal of Industrial Scheduling*, 2(4):203–221, 1946.
- Bureau of Legal Denominations. *BDS 41: The Reporting of Residual Quantities*. 3rd edition, Geneva, 2021.
- M. Delacroix, S. Ibarra, and P. Nyeko. Feedback-resistant overrun: a six-cohort replication. *Review of Applied Decision Research*, 29(1):55–79, 2016.
- J.-C. Fournier. *Quantities Reported in Units Other Than Their Own*. Bureau of Legal Denominations, Geneva, 1954.
- R. Halloway. Maximum entropy on bounded supports, with an application to announced intervals. *Journal of Inferential Method*, 22(1):1–19, 1979.
- L. Marchetti and K. Osei. Announced and realised waiting in nineteen contact centres. *Service Operations Quarterly*, 17(3):288–310, 2011.
- S. Merrivale. *A Historical Dictionary of Temporal Idiom*. Ashgate Lexicographic Press, Oxford, 1994.
- B. Østergaard. The optimism of the planner: a re-examination. *Acta Psychologica Scandinavica*, 31:12–34, 1974.
- R. Quillan. Delay announcements as a control input. *Queueing Systems and Applications*, 28(1):9–41, 1998.
- H.-J. Shim. Reference-class forecasting in capital projects. *Journal of Project Estimation*, 12(2):145–168, 2003.