

Trustworthy decisions in stochastic amplification assays: timing limits, measurement design, and effective capacity

17 September 2026

Abstract

Changing a readout time, adding an identity-sensitive measurement, increasing the reagent reserve, and improving the inference procedure solve different problems, and a laboratory that confuses them can spend its effort on a step that cannot succeed. We separate these problems for a finite pure-birth amplification source with Poisson loading, a latching detection threshold and an effective resource capacity. For the source at threshold five, every measurable randomized classifier of the complete crossing time that keeps blank error at most 1% misses more than 5.55% of loaded reactions; this all-rule obstruction, its concrete source identities and a source-preserving hit-and-identity repair with blank error below 0.03% and miss below 4.22443% are verified in Lean 4. We then prove, with ordinary mathematics, that for any fixed sequential source with shared kinetics a deadline is optimal among all crossing-time classifiers; this converts every all-deadline exclusion into an all-classifier exclusion and sharpens the certified bound to 6.9%. Using the holding-time representation of the crossing time we show that the obstruction is not a small-count artefact: for every integer threshold $h \geq 10^6$ with the resource exhausted at detection, every crossing-time classifier at 1% blank error misses more than 5.63%, an explicit identity measurement misses less than 4.13%, and doubling the capacity restores a usable deadline with errors below 0.9373% and 4.8732%. The proof replaces earlier Chebyshev allowances by gamma Chernoff and weighted Bernstein bounds and reduces the sufficient threshold from 10^8 to 10^6 . For the mechanism question we exhibit two capacities with exactly the same blank endpoint law, prove that calibrated intermediate observations recover capacity through clock-free ordered block ratios with strictly monotone, uniquely invertible ordering probabilities, extend the block theorem to non-exponential factorized holding times, and show that thresholds chosen as fractions of each reaction's own plateau lose the leading capacity information. A signal-to-state interval construction and a censoring-safe confidence procedure that retains every incomplete observation complete an executable analysis. All performance statements are conditional mathematical results; no physical channel, calibration or biochemical resource is asserted to have been measured.

1 Introduction

A digital amplification reaction can cross its signal threshold because it contained target or because background amplification eventually produced the same product [11, 23, 25, 30]. An earlier deadline suppresses background at the price of discarding slowly developing target reactions. Four different remedies are routinely discussed: tune the cutoff, add an identity-sensitive readout such as a melt curve [26], change the reaction so that more resource remains at detection, or analyse the data more carefully. These are not interchangeable, and the purpose of this paper is to give, for one explicitly declared stochastic source, mathematical diagnostics that tell them apart.

A worked decision. Consider a laboratory with a digital assay whose blank wells occasionally cross threshold and whose loaded wells contain on average four target units. The operating limits are a 1% blank error and a 5% miss per reaction. The questions that determine the next experiment are the following.

- (1) *Is there a readout time at which both limits hold?* If the answer is provably no for every deadline, cutoff tuning is futile. Section 3 certifies such an instance.
- (2) *Could a cleverer rule on the same crossing-time datum do better than a deadline?* For the declared source the answer is no: Theorem 3.2 shows that the best measurable randomized function of the complete crossing time is a deadline. A negative answer to (1) is therefore an information obstruction, not a search failure.
- (3) *What would an additional measurement have to achieve?* Proposition 3.4 gives the hit-conditioned sensitivity and false-positive rate that make a hit-and-identity rule meet the limits, and Section 3 exhibits a joint observation law realizing them without changing the count kinetics.
- (4) *Would the conclusion change at realistic thresholds?* Section 4 shows it does not: at every threshold $h \geq 10^6$ with the resource exhausted at detection the obstruction persists with explicit constants, while a doubled capacity restores a deadline. Which of these regimes an assay occupies is decided by its remaining headroom at detection, not by the size of the count.
- (5) *Does a failing timing rule identify the mechanism?* No: Section 5 exhibits two capacities with identical blank endpoint laws, and shows what calibrated intermediate observations do and do not recover.
- (6) *How should incomplete observations be treated?* Section 6 shows that deleting slow reactions reintroduces the very clock dependence the method removes, and gives a confidence procedure that retains every observation and returns an interval, a one-sided bound, an unresolved verdict or a model incompatibility.

Main results in words. The source is a pure-birth chain on $\{0, \dots, h\}$ with rate $(b + gz)(1 - z/R)$ from count z , latched at the threshold h , with Poisson loading of mean λ (Section 2). Writing $a = b/g$, the nominal instance has $a = .01$, $\lambda = 4$ and limits $\alpha = .01$, $\beta = .05$.

- (i) At $R = h = 5$, every measurable randomized classifier of the complete crossing time with blank error at most .01 has miss greater than .0555; a source-preserving hit-and-identity rule at model time seven has blank error below .0003 and miss below .0422443, an advantage of more than 1.32557 percentage points (Theorem 3.1; Lean-verified).
- (ii) For any fixed positive pre-hit rates shared by both classes, and any Poisson (indeed any) loading law, a deadline at the blank α -quantile maximizes power among all measurable randomized crossing-time classifiers (Theorem 3.2; ordinary proof). Combined with the certified separating time at $t = 5$ this sharpens (i) to miss greater than .069, and combined with exact enclosures at threshold twenty it excludes every classifier there with miss greater than .087 (Corollary 3.3).
- (iii) Classwise total-variation budgets, a $\pm 2\%$ fixed-rate operating window, and a label-neutral-pipeline comparison show what the certified separation tolerates and what cannot repair it (Propositions 3.6–3.5).
- (iv) The crossing time is a sum of independent exponential holding times whose undepleted part has an exact beta law. With a growing reserve, depletion is a deterministic delay and the limiting tradeoff is the undepleted one; with fixed headroom m , the last $m + 1$ births add an independent $\text{Gamma}(m + 1)$ delay in logarithmic time (Theorems 4.2–4.3).

For every integer $h \geq 10^6$: at $R = h$ every crossing-time classifier at 1% blank error misses more than 5.63%, an explicit identity measurement misses less than 4.13% with blank error below 0.085%, and at $R = 2h$ an explicit deadline has errors below .009373 and .048732 (Theorem 4.4; ordinary proof with directed scalar certificates).

- (v) At threshold three the sources $(b, g, R) = (2, 1, 6)$ and $(8/3, 2/3, 4)$ have identical blank crossing-time measures (Proposition 5.1; Lean-verified). With a known and calibrated states observed, paired waits and ordered block ratios cancel an unknown constant clock and recover capacity; the block ordering probability is strictly decreasing and uniquely invertible on its physical range (Theorems 5.2–5.3), the pathwise ordering survives non-exponential, correlated, factorized holding times (Proposition 5.4), and plateau-relative thresholds converge to a capacity-free limit (Proposition 5.5).
- (vi) A signal-to-state interval construction, a censoring-safe confidence theorem that needs no independent-censoring assumption (Theorem 6.1), and a three-way decision procedure with Lean-verified band logic turn the theory into an executable analysis (Section 6).

Three kinds of evidence. Every statement below carries one of the following labels. *Lean-verified* statements are theorems of a frozen Lean 4/Mathlib development [7, 29], compiled with warnings treated as errors and depending only on the standard axioms; Section 8 and Appendix E list them. *Ordinary proof* means a complete mathematical argument printed here; it has not thereby become a Lean declaration. *Exact rational computation* and *directed computation* mean that integer or rational arithmetic, or outward rounding with explicit remainder bounds, certifies the stated finite inequalities; the script `check_paper.py` replays all of them. *Numerical illustration* means a floating-point evaluation, root or seeded simulation; such values are never premises of a theorem. *Empirical* conclusions require physical observations, and none is established here.

Relation to prior work. The tradeoff between speed, sensitivity and background is established experimental territory [11, 12, 25, 27], and kinetic and identity information have already been combined in real-time digital LAMP with high-resolution melting [26]; adding identity information is therefore not itself the contribution. On the statistical side, optimality of likelihood-ratio tests [21], the monotone-likelihood-ratio theory [16, 18], comparison of experiments [4], the beta–gamma algebra [19], the logarithmic limit of the Yule process [17] and the concentration inequalities we use [5, 14] are classical, and we claim no priority for them. The contribution is their source-connected combination: a concrete, formally checked separation between a timing experiment and an identity experiment; the general deadline-optimality theorem that makes the separation an information statement; an explicit large-threshold theorem showing that neither concentration of large counts nor the vanishing of initial hits removes it; and observation-design results for effective capacity that distinguish exact nonidentification, loss of conditioning, selection bias and normalization artefacts. The certified capacity classification at threshold five that this paper builds on is in a companion manuscript [2].

What is not claimed. The units of the source are effective amplification units, not molecules or fluorescence; rate constants are illustrative; the capacity is an effective inventory whose laboratory realization is an open question; the assumed 99%/1% channel capabilities are stated premises, not measured performance; and no biological dataset is used to fit or validate any constant. Neither clinical performance nor priority over the experimental literature is asserted. The results do not establish universal necessity of an identity readout: the reserve regime of Section 4 is a positive control in which timing alone works.

Organization. Section 2 defines the source, the observations and the decision unit. Section 3 contains the certified separation, the deadline-optimality theorem and the robustness propositions. Section 4 contains the large-threshold theory. Section 5 treats effective capacity. Section 6 treats signal calibration, incomplete observations and the decision procedure. Section 7 states the scope and the empirical frontier, and Section 8 the verification and reproduction record. Appendices A–E contain the source calculation, the sign certificate, the proof of Theorem 4.4 with its directed arithmetic, and the claims-to-declarations map.

2 The source, the observations, and the decision unit

2.1 Counts, capacity, loading and classifiers

Let X_t be a continuous-time pure-birth chain on $\{0, \dots, h\}$, latched (absorbed) at h , with pre-hit rates

$$q_z = (b + gz)\left(1 - \frac{z}{R}\right), \quad 0 \leq z < h, \quad b, g > 0, \quad R > h - 1. \quad (1)$$

Here b is background initiation, gz is growth proportional to the active count and $1 - z/R$ is the resource factor; $R - z$ is the remaining effective inventory. Target-origin and background-origin units grow with the same coefficient, so the total count is a closed observable. The crossing time is $T = \inf\{t : X_t = h\}$. A *blank* starts at zero. A *loaded* preparation starts at $\min(N, h)$, where $N \sim \text{Poisson}(\lambda)$ is independent of the subsequent source innovations; loads $N \geq h$ give an atom of T at zero, which is retained. Write

$$F(t) = \mathbb{P}_0(T \leq t), \quad M(t) = \mathbb{P}_1(T > t),$$

for the blank crossing probability and the loaded miss probability of the deadline t . Time is model time; a physical rate calibration is needed before a numerical deadline is an assay turnaround time. Writing $a = b/g$ and $\tau = gT$ removes g from every statement.

The decision unit is one reaction, and empty loaded reactions ($N = 0$) remain in the denominator. A *measurable randomized complete-time classifier* is a measurable $\phi : [0, \infty) \rightarrow [0, 1]$, interpreted as the probability of calling a reaction positive when its crossing time is T ; its errors are

$$\text{FP}(\phi) = \mathbb{E}_0 \phi(T), \quad \text{Miss}(\phi) = 1 - \mathbb{E}_1 \phi(T). \quad (2)$$

A deadline is $\phi(t) = \mathbf{1}\{t \leq t_0\}$. Complete timing reveals the eventual crossing time even beyond any finite acquisition, so an obstruction for complete-time classifiers is an obstruction for every finite acquisition of the same datum. The nominal instance is

$$b = \frac{1}{100}, \quad g = 1, \quad R = h = 5, \quad \lambda = 4, \quad \alpha = \frac{1}{100}, \quad \beta = \frac{1}{20}. \quad (3)$$

Both classes share the post-loading rates (1). An assay in which target and background products have different subsequent kinetics is a different source: the decision-theoretic method applies, but the constants do not transfer.

Proposition 2.1 (Loading floor; ordinary proof). *If an empty loaded reaction has the blank observation law, then every rule with blank error at most α has $\text{Miss} \geq e^{-\lambda}(1 - \alpha)$. Miss at most β therefore requires $\lambda \geq \log((1 - \alpha)/\beta)$, which is 2.9857 at the nominal limits; at $\lambda = 4$ the floor is .018133.*

Proof. The loaded law is $P_1 = e^{-\lambda}P_0 + (1 - e^{-\lambda})P_{\text{occ}}$. Integrate $1 - \phi \geq 0$ against it and rearrange. $\square$

The floor concerns loaded preparations, not occupancy-conditioned sensitivity. Splitting a fixed specimen among several wells changes the loading, so independence alone does not justify raising a per-reaction miss to a power.

2.2 The probability law and the holding-time representation

The formal development defines every probability through uniformization [15, 22]: with $\nu \geq \max_z q_z$ and $P = I + Q/\nu$ for the generator Q of the chain, the transition matrix is $K_t = e^{-\nu t} \sum_k (\nu t)^k P^k / k! = e^{tQ}$, and F, M are the corresponding Poisson mixtures (Appendix A). The conventional arguments use instead the holding-time construction.

Lemma 2.2 (Holding times; ordinary proof). *Let $E_n, \dots, E_{h-1}$ be independent unit exponential variables. The crossing time from state $n < h$ satisfies*

$$T_n \stackrel{d}{=} \sum_{z=n}^{h-1} \frac{E_z}{q_z} = \frac{1}{g} \frac{R}{R+a} \left[\sum_{z=n}^{h-1} \frac{E_z}{z+a} + \sum_{z=n}^{h-1} \frac{E_z}{R-z} \right]. \quad (4)$$

Proof. A pure-birth chain holds in state z for an $\text{Exp}(q_z)$ time and moves to $z+1$; successive holding times are independent by the strong Markov property [9, 22]. The transition matrix of this construction solves the finite forward equation $K'_t = K_t Q$, $K_0 = I$, whose unique solution is e^{tQ} ; so it coincides with the uniformized law (Appendix A). The second form is partial fractions. $\square$

The first sum in (4) is the undepleted crossing time and the second is the resource delay; they share the same E_z and are not independent at finite h . The identification of the uniformized law with the holding-time construction is an ordinary finite-state proof; it is not formalized, and no coupling argument below is labeled Lean-verified.

Question	Diagnostic supplied	Evidence
Does any readout time work?	Separating-time certificates; usable/excluded/unresolved bands (Section 6.5)	Lean at $h = 5$, $R = 5$ and 10; exact computation at $h = 20$; ordinary proof for $h \geq 10^6$
Can a better crossing-time rule help?	All-classifier bound (Theorem 3.1); deadline optimality (Theorem 3.2)	Lean; ordinary proof
What must an added measurement achieve?	Hit-conditioned channel bounds and a realizing joint law (Proposition 3.4)	Lean
Does the count scale matter?	Startup and terminal limit laws; uniform theorem at $h \geq 10^6$ (Section 4)	Ordinary proof, directed computation
Can endpoint data identify capacity?	Exact endpoint-law equality (Proposition 5.1)	Lean
What does intermediate data identify?	Paired waits, ordered blocks, strict inversion, conditioning, normalization limit (Section 5)	Lean algebra; ordinary proof; exact rational solves
How to treat incomplete observations?	Interval categories, censoring-safe confidence sets, capacity sets (Section 6)	Ordinary proof; seeded illustration

Table 1: The questions of the worked decision, the diagnostic each receives, and its evidence type.

3 Timing limits at small thresholds and an explicit repair

3.1 A certified separation

Theorem 3.1 (Certified decision separation; Lean-verified). *For the nominal instance (3), every measurable randomized complete-time classifier with $\text{FP} \leq .01$ has $\text{Miss} > .0555$. At*

model time seven there is a source-preserving three-category hit-and-identity experiment for which the rule “hit and identity positive” has

$$\text{FP}_J < .0003, \quad \text{Miss}_J < .0422443, \quad (5)$$

so its miss advantage over every admissible timing classifier exceeds .0132557.

The timing part is the declaration `full_timing_obstruction`; the joint part is `concrete_joint_rescue`; the gap is `decision_gap`. All three rest on source identities that are themselves verified, not assumed (Appendix E).

The certificate. Let f_0, f_1 be the densities on $(0, \infty)$ of the two crossing laws; the loaded law also has the atom $a_1 = \mathbb{P}(N \geq 5) = 1 - M(0)$ at zero. Appendix B proves, from five-term exponential polynomials with outward-rounded rational coefficients, the source-connected inequalities

$$f_1(t) \geq 5f_0(t) \quad (0 < t \leq 4), \quad f_1(t) \leq 5f_0(t) \quad (t \geq 5), \quad F(4) > .0073, \quad M(5) > .069. \quad (6)$$

No sign is needed on $(4, 5)$. For $0 \leq \phi \leq 1$, on the three regions $(0, 4]$, $(4, 5]$, $(5, \infty)$ respectively,

$$\phi f_1 \leq 5\phi f_0 + \mathbf{1}_{(0,5]} f_1 - 5 \mathbf{1}_{(0,4]} f_0.$$

Integrating, using $a_1 \phi(0) \leq a_1$ and $a_1 + \int_0^5 f_1 = 1 - M(5)$,

$$\text{Miss}(\phi) \geq M(5) + 5F(4) - 5\text{FP}(\phi) > .1055 - 5\text{FP}(\phi). \quad (7)$$

This is the declaration `timing_dual` combined with `dual_source_margin`. The argument credits the competitor with the maximal acceptance of the loaded atom, so immediately positive reactions are treated in its favour. The constant .0555 is a conservative certified lower bound, not an optimum.

3.2 Deadlines are optimal crossing-time rules

The certificate above is specific to one instance. The following theorem explains why the phenomenon it certifies is an information obstruction for the whole declared source family.

Theorem 3.2 (Deadline optimality; ordinary proof). *Fix positive pre-hit rates $q_0, \dots, q_{h-1}$, identical under both classes, not necessarily monotone or distinct. Under P_0 the chain starts at zero. Under P_1 it starts at a random state in $\{0, \dots, h\}$ with arbitrary fixed weights $\pi_0, \dots, \pi_h$, state h representing an atom of T at zero. Then for every $0 < \alpha < 1$, the deadline $\phi_* = \mathbf{1}\{T \leq t_\alpha\}$ at the blank α -quantile t_α maximizes $\mathbb{E}_1 \phi(T)$ among all measurable $\phi : [0, \infty) \rightarrow [0, 1]$ with $\mathbb{E}_0 \phi(T) \leq \alpha$. The loaded atom at zero is accepted at no blank-error cost.*

Proof. Let f_n be the density of $T_n = \sum_{z=n}^{h-1} E_z / q_z$ on $(0, \infty)$.

Step 1: each f_n is log-concave on $(0, \infty)$. The density f_{h-1} is exponential. Inductively $f_n(t) = q_n e^{-q_n t} \int_0^t e^{q_n s} f_{n+1}(s) ds$. Log-concavity is preserved by convolution [3, 24], which already proves the step; the following elementary argument avoids the general theorem. For a positive differentiable log-concave k on $(0, \infty)$ with a finite limit $k(0)$ at zero, put $H(t) = \int_0^t k$ and $m = k'/k$, which is nonincreasing. Then $k(t) - k(0) = \int_0^t m(s)k(s) ds \geq m(t)H(t)$, hence $H''H - (H')^2 = k'H - k^2 \leq -k(0)k(t) \leq 0$, so H is log-concave. Apply this to $k(s) = e^{q_n s} f_{n+1}(s)$, a product of log-concave functions (when $f_{n+1}(0) = 0$, integrate from a positive lower endpoint and let it tend to zero; the finite exponential sums here have the required endpoint regularity). Multiplying H by the log-concave factor $q_n e^{-q_n t}$ gives f_n .

Step 2: f_n/f_0 is nonincreasing for every n . For $n > 0$ write $T_0 = U_n + T_n$ with independent prefix $U_n = \sum_{z < n} E_z/q_z$ and suffix T_n ; then $f_0 = f_n * P_{U_n}$ and

$$\frac{f_0(t)}{f_n(t)} = \int_{[0,t]} \frac{f_n(t-u)}{f_n(t)} dP_{U_n}(u).$$

For fixed $u \geq 0$ the integrand, extended by zero for $t < u$, is nondecreasing in t : where positive its logarithmic derivative is $(\log f_n)'(t-u) - (\log f_n)'(t) \geq 0$ by Step 1. Hence f_0/f_n is nondecreasing, and f_n/f_0 nonincreasing.

Step 3: the likelihood ratio is nonincreasing. On $(0, \infty)$ the loaded density is $f_1 = \sum_{n < h} \pi_n f_n$, so $f_1/f_0 = \sum_{n < h} \pi_n (f_n/f_0)$ is a nonnegative combination of nonincreasing functions (the $n = 0$ term is constant).

Step 4: Neyman–Pearson. The blank distribution function is continuous and strictly increasing from zero to one, so t_α exists. Put $c = f_1(t_\alpha)/f_0(t_\alpha)$. For any admissible ϕ , $(\phi - \phi_*)(f_1 - cf_0) \leq 0$ pointwise on $(0, \infty)$: for $t < t_\alpha$ the first factor is ≤ 0 and the second ≥ 0 , and conversely for $t > t_\alpha$. Integrating, $\mathbb{E}_1^c \phi - \mathbb{E}_1^c \phi_* \leq c(\mathbb{E}_0 \phi - \mathbb{E}_0 \phi_*) = c(\mathbb{E}_0 \phi - \alpha) \leq 0$, where $\mathbb{E}_1^c$ is expectation against the continuous part. On the atom, $\pi_h \phi(0) \leq \pi_h = \pi_h \phi_*(0)$. Adding the two parts gives $\mathbb{E}_1 \phi \leq \mathbb{E}_1 \phi_*$ [16, 21]. $\square$

Corollary 3.3 (All-deadline exclusion is all-classifier exclusion). *For a source as in Theorem 3.2, the set of deadlines meeting both limits is empty if and only if no measurable randomized crossing-time classifier meets them. In particular, if some time s has $F(s) > \alpha$ and $M(s) > \beta$ then every classifier with blank error at most α has miss at least $M(s)$. Consequently:*

- (a) *at the nominal instance, since $F(5) > .01$ and $M(5) > .069$ are Lean-verified (`capacity5_separating_errors`, `dual_miss_lower`), every classifier with $\text{FP} \leq .01$ has $\text{Miss} > .069$ (ordinary proof plus Lean certificates), sharpening the fully formal .0555; the numerical optimum, the deadline at $t_\alpha = 4.4594$, has miss 8.9566%;*
- (b) *at $R = h = 20$ with the same a and λ , the exact enclosures $F(8) > .013$ and $M(8) > .087$ (Table 2) exclude every classifier with $\text{FP} \leq .01$ below miss .087;*
- (c) *every finite separating time produced in Section 4 excludes every classifier in the same way.*

Proof. Since F is nondecreasing, $F(s) > \alpha$ gives $t_\alpha < s$, and since M is nonincreasing, $\text{Miss}(\phi_*) = M(t_\alpha) \geq M(s)$. Apply Theorem 3.2. $\square$

The theorem covers serial exponential intermediate steps (a chain with additional states is still a sequential chain with ordered initial positions) but not branching intermediates, backward transitions, class-dependent kinetics or random rate mixtures; Proposition 3.5 covers some of the latter in a different way.

3.3 The required additional information

Define the occupied-hit mass

$$C(t) = \mathbb{P}_1(N > 0, T \leq t) = 1 - M(t) - e^{-\lambda} F(t). \quad (8)$$

Proposition 3.4 (Channel requirements; Lean-verified at the instance). *Suppose a marker is available on hits, with positive probability at most f on blank hits and at least s on occupied hits, as conditional bounds on those two populations; no independence between marker and crossing time is assumed. The rule “hit by t and marker positive” satisfies*

$$\text{FP}_J \leq fF(t), \quad \text{Miss}_J \leq 1 - sC(t). \quad (9)$$

At the nominal instance, exact enclosures give $F(7) < .03$, $M(7) < .032$, $e^{-4} < .019$, hence $C(7) > .96743$; with $f = .01$ and $s = .99$ this yields (5), since $1 - .99 \times .96743 = .0422443$. An explicit joint law realizes the premises: with $F = F(7)$, $M = M(7)$, $C = C(7)$, $e = e^{-4}$, the blank and loaded probabilities of (hit, positive), (hit, negative), (no hit) are

$$(.01F, .99F, 1 - F), \quad (.99C + .01eF, .01C + .99eF, M). \quad (10)$$

They are nonnegative, sum to one, and retain the hit marginals F and $1 - M$ (*joint_fixture_valid*).

The fixture assigns positive markers with probabilities .99 and .01 to occupied and empty hits, so it changes nothing about the count kinetics. It establishes that the required observation law is *compatible* with the same total-count source; it does not supply a biochemical implementation. A physical identity channel must respond to something retained by the products, not to the preparation label; if a reporter changes kinetics, both source and observation law must be re-evaluated, and an acquisition duration τ adds to turnaround. The fixture also credits accidental positive calls on empty loaded wells, a credit that (9) conservatively ignores. Figure 1 shows the forbidden region and the channel requirements as functions of the readout time; the miss advantage does not imply Blackwell dominance [4] of the identity experiment at every operating point.




Figure 1: (a) The numerical deadline tradeoff at the nominal instance, the region forbidden to every crossing-time classifier (Lean-verified 5.55% and ordinary-proof 6.9%), the numerical optimum, and the certified upper bounds of the hit-and-identity rule. (b) The hit-conditioned sensitivity a marker must reach for a 5% miss, $s \geq .95/C(t)$, near its feasible boundary. (c) The blank-hit false-positive rate it may have for a 1% blank error, $f \leq .01/F(t)$. Curves are numerical illustrations; the shaded region and the marked point are theorems.

Proposition 3.5 (Label-neutral pipelines cannot repair the obstruction; ordinary proof). *Suppose an observed datum Y is generated from the complete crossing time by a Markov kernel $K(t, dy)$ that is the same under both labels. Then every classifier $\psi(Y)$ has exactly the errors of the crossing-time classifier $\phi(t) = \int \psi(y)K(t, dy)$. Hence every all-classifier lower bound of this section and of Section 4 also holds for rules using Y .*

Proof. ϕ is measurable with values in $[0, 1]$ and $\mathbb{E}_i \psi(Y) = \mathbb{E}_i \phi(T)$ for $i = 0, 1$ by the tower property. $\square$

This covers censoring at a finite horizon, timing noise, frame quantization, and an exogenous random global time change whose law does not depend on the label, including one with a “never observed” outcome. It does not require the transformed law to keep a monotone likelihood ratio, so the obstruction survives some random-clock mixtures for which deadline optimality

itself needs separate analysis. It does not cover richer trajectory information or label-dependent clocks. An identity measurement escapes because it is not a label-neutral transformation of timing.

3.4 Robustness

Proposition 3.6 (Classwise discrepancy budgets; ordinary proof, scalar consequence Lean-verified). *With $d_{TV}(P, Q) = \sup_A |P(A) - Q(A)|$, suppose the actual blank and loaded crossing-time laws differ from the nominal ones by at most δ_0 and δ_1 in total variation. Then any rule with actual blank error at most .01 has actual miss greater than $.0555 - 5\delta_0 - \delta_1$; exclusion at 5% miss persists whenever $5\delta_0 + \delta_1 \leq .0055$. For the fixed joint rule, classwise discrepancies η_0, η_1 give $FP_J < .0003 + \eta_0$ and $Miss_J < .0422443 + \eta_1$, so the limits survive if $\eta_0 \leq .0097$ and $\eta_1 \leq .0077557$.*

Proof. For $0 \leq u \leq 1$, $\mathbb{E}_P u = \int_0^1 P(u > x) dx$, so $|\mathbb{E}_P u - \mathbb{E}_Q u| \leq d_{TV}(P, Q)$. The nominal blank error is at most .01 + δ_0 ; apply (7) and transfer the miss back with error δ_1 (**discrepancy_budget**). Apply the same inequality directly to the joint rule. $\square$

These are sufficient budgets, not measured discrepancies.

Proposition 3.7 (A fixed-rate operating window; ordinary proof, exact computation, scalar consequence Lean-verified). *Keep $R = h = 5$ and let each fixed rate satisfy $q'_z \in [.98q_z, 1.02q_z]$, the same vector in both classes. Let $\lambda \geq 3.99$, let the actual readout time lie in $[6.99, 7.01]$, and assume $f \leq .01$, $s \geq .99$ uniformly. Then $FP_J < .00032$ and $Miss_J < .04525192$.*

Proof. Couple holding times through common unit exponentials E_z : crossing from $n < h$ takes $\sum_{z \geq n} E_z/q'_z$, which decreases when any rate increases and when the initial state increases; couple Poisson loads by adding independent increments. By the source equivalence of Lemma 2.2, $F_{q'}(t) \leq F_q(1.02 \times 7.01)$ and $M_{q', \lambda}(t) \leq M_{q, 3.99}(.98 \times 6.99)$. Exact rational uniformization (Appendix A) gives $F_q(7.1502) < .032$, $M_{q, 3.99}(6.8502) < .035$ and $e^{-3.99} < .019$, hence $C > 1 - .035 - .019 \times .032 = .964392$ and, by (9), $Miss_J \leq 1 - .99 \times .964392 = 565649/12500000 = .04525192$ (**robust_channel_arithmetic**). $\square$

The inherited window of the companion manuscript is the same kind of statement for a usable timing rule: at $R = 10$, $h = 5$, the deadline 3.2 meets both limits for every fixed rate vector within 2%, every loading at least 3.99 and every actual time in $[3.19, 3.21]$ (**robust_guarantee**). Arbitrarily time-varying rates are outside both propositions.

3.5 A larger threshold with a negligible initial atom

At threshold five, $\mathbb{P}(N \geq 5) = .371163$ of loaded reactions are positive at time zero. To remove this feature, take $R = h = 20$ with $b = .01$, $g = 1$, $\lambda = 4$; then $\mathbb{P}(N \geq 20) = 1.0200522 \times 10^{-8}$. Table 2 gives exact enclosures at times eight and ten. Every deadline is excluded, and by Corollary 3.3 so is every crossing-time classifier, with miss above .087; at time ten the same channel premises give exactly the bounds (5).

4 Large thresholds: startup and terminal randomness

Throughout this section $g = 1$, $a = b/g > 0$, λ is fixed and $h \rightarrow \infty$; everything is ordinary proof or directed computation. Write $T_{R, h, n}$ for the crossing time from state $n < h$ at threshold h and capacity R .

Quantity at $R = h = 20$	Numerical value	Exact enclosure
$F(8)$	.0134183123	$> .013$
$M(8)$	.0876466114	$> .087$
$F(10)$	.0297115009	$< .030$
$M(10)$	.0318119401	$< .032$
Joint blank error at 10	.0002971150	$< .0003$
Joint miss at 10	.0420271221	$< .0422443$

Table 2: Threshold twenty. Numerical joint values use the fixture (10), which credits positive calls on empty loaded wells; the enclosures use (9) and rational uniformization with $L = 300$ terms. Neither threshold is a calibrated optical detection level.

4.1 The startup law

Lemma 4.1 (Beta law and gamma coupling; ordinary proof). *Let $A_{h,n} = \sum_{z=n}^{h-1} E_z/(z+a)$. Then $e^{-A_{h,n}} \sim \text{Beta}(n+a, h-n)$, and on an extension of the probability space there are $Z \sim \text{Gamma}(h+a, 1)$ and $W_n \sim \text{Gamma}(n+a, 1)$ with $A_{h,n} = \log Z - \log W_n$ exactly. Consequently $A_{h,n} - \log h \Rightarrow -\log W_n$.*

Proof. $\mathbb{E}e^{-sA_{h,n}} = \prod_{z=n}^{h-1} \frac{z+a}{z+a+s} = \frac{\Gamma(n+a+s)\Gamma(h+a)}{\Gamma(n+a)\Gamma(h+a+s)}$ is the s th moment of $\text{Beta}(n+a, h-n)$, and moments determine a law on $[0, 1]$. Draw $Z \sim \text{Gamma}(h+a, 1)$ independent of $e^{-A_{h,n}}$ and set $W_n = Ze^{-A_{h,n}}$; by the beta-gamma algebra [19], $W_n \sim \text{Gamma}(n+a, 1)$. Since $Z/(h+a) \rightarrow 1$ in probability, the limit follows. Z and W_n are not independent, and nothing below uses independence between them. $\square$

This is the logarithmic limit of the Yule process with immigration [17]: after $\log h$ units of time a reaction started with n units behaves like a $\text{Gamma}(n+a)$ effective initial mass. The blank member $\text{Gamma}(a)$ with $a = .01$ is extremely spread, and this spread is the whole difficulty of separating blanks from loaded reactions by timing. A large final count does not make a finite-founder startup deterministic. The limiting shift is minus the logarithm of a gamma variable; it is Gumbel only for shape one.

4.2 Two resource regimes

Under the coupling of Lemma 2.2, $T_{R,h,n} = \frac{R}{R+a}[A_{h,n} + D_{R,h,n}]$ with $D_{R,h,n} = \sum_{z=n}^{h-1} E_z/(R-z)$.

Theorem 4.2 (Growing reserve; ordinary proof). *Let $R = R(h) \geq h+1$ with $R-h \rightarrow \infty$, and let $d_h = \sum_{z=0}^{h-1} \frac{z}{(z+a)(R-z)}$. Then for every fixed n , $T_{R,h,n} - \log h - d_h \Rightarrow -\log W_n$ with $W_n \sim \text{Gamma}(n+a, 1)$. If $R/h \rightarrow \rho > 1$ then $d_h \rightarrow \log \frac{\rho}{\rho-1}$. For deadlines $t = \log h + d_h + s$,*

$$F(t) \rightarrow \mathbb{P}(W_0 \geq e^{-s}), \quad M(t) \rightarrow \sum_{n \geq 0} \mathbb{P}(N = n) \mathbb{P}(W_n \leq e^{-s}),$$

the tradeoff of the undepleted source after a common delay.

Proof. Write $T = A_{h,n} + D'_n$ with $D'_n = \sum_{z \geq n} \frac{zE_z}{(z+a)(R-z)}$ (partial fractions). $\text{Var } D'_n \leq \sum_{z < h} (R-z)^{-2} \leq (R-h)^{-1} \rightarrow 0$ and $|\mathbb{E}D'_n - d_h| \leq \sum_{z < n} (R-z)^{-1} \rightarrow 0$, so $D'_n - d_h \rightarrow 0$ in probability by Chebyshev; combine with Lemma 4.1. For $\rho > 1$, $d_h = \sum_{z < h} \frac{1}{R-z} - \sum_{z < h} \frac{a}{(z+a)(R-z)}$, whose first sum tends to $\log \frac{\rho}{\rho-1}$ and whose second is $O(\log h/(R-h))$. The error limits follow by continuity of the limit laws and dominated convergence over n . $\square$

Theorem 4.3 (Fixed headroom; ordinary proof). *Let $R = h + m$ with a fixed integer $m \geq 0$. Then for every fixed n ,*

$$T_{h+m,h,n} - 2 \log h \Rightarrow -\log W_n - \log Y_{m+1},$$

with $W_n \sim \text{Gamma}(n + a, 1)$ and $Y_{m+1} \sim \text{Gamma}(m + 1, 1)$ independent. For $m = 0$, $-\log Y_1$ is standard Gumbel. For deadlines $t = 2 \log h + s$, $F(t) \rightarrow \mathbb{P}(W_0 Y_{m+1} \geq e^{-s})$ and $M(t) \rightarrow \sum_n \mathbb{P}(N = n) \mathbb{P}(W_n Y_{m+1} \leq e^{-s})$.

Proof. Put $c = R/(R + a)$ and $k = \lfloor h/2 \rfloor$. By (4), $T/c = U_n + V + C_n$ with

$$U_n = \sum_{z=n}^{k-1} \frac{E_z}{z+a}, \quad V = \sum_{z=k}^{h-1} \frac{E_z}{R-z} = \sum_{j=m+1}^{R-k} \frac{E'_j}{j}, \quad C_n = \sum_{z=n}^{k-1} \frac{E_z}{R-z} + \sum_{z=k}^{h-1} \frac{E_z}{z+a},$$

where $E'_j = E_{R-j}$. The blocks U_n and V use disjoint exponentials and are independent. By Lemma 4.1 with h replaced by k , $U_n - \log k \Rightarrow -\log W_n$; the transform $\mathbb{E}e^{-sV} = \prod_{j=m+1}^{R-k} \frac{j}{j+s}$ is that of $\text{Beta}(m+1, R-k-m)$, so the same coupling gives $V - \log(R-k-m) \Rightarrow -\log Y_{m+1}$ with $Y_{m+1} \sim \text{Gamma}(m+1, 1)$, and the two limits are independent. The remainder has $\mathbb{E}C_n \rightarrow 2 \log 2$ (it equals $\log \frac{R}{R-k} + \log \frac{h}{k} + o(1)$) and $\text{Var } C_n \leq k/(R-k)^2 + (h-k)/k^2 = O(1/h)$, so $C_n - 2 \log 2 \rightarrow 0$ in probability. Since $\log k + \log(R-k-m) + 2 \log 2 = 2 \log h + o(1)$ and $(1-c)O(\log h) \rightarrow 0$, the claim follows. For $m = 0$, $Y_1 \sim \text{Exp}(1)$ and $-\log Y_1$ is Gumbel [8]. $\square$

The mechanism is concrete. At $R = h$ the last transition has rate $q_{h-1} = g(a+h-1)/h \rightarrow g$, so its mean waiting time stays of order $1/g$ however large the threshold; at $R = 2h$ the same rate is asymptotic to $gh/2$ and the waiting time vanishes. The last $m+1$ births before exhaustion contribute an independent $\text{Gamma}(m+1)$ variable in logarithmic time in addition to the startup shift. Random slowdown near an absorbing boundary is known in absorption-time asymptotics [13]; the point here is its exact form and its consequence for decisions. For $m = 0$ the product law has the closed form

$$\mathbb{P}(W_n Y_1 \geq y) = \frac{1}{\Gamma(n+a)} \int_0^\infty w^{n+a-1} e^{-w-y/w} dw = \frac{2y^{(n+a)/2}}{\Gamma(n+a)} K_{n+a}(2\sqrt{y}). \quad (11)$$

Regime	Centered limiting time	Blank boundary y	Miss at 1% blank error
Growing reserve ($R-h \rightarrow \infty$, e.g. $R = 2h$)	$-\log W_N$	.26505255	3.96249%
Zero headroom ($R = h$)	$-\log W_N - \log Y_1$	.19364809	10.91390%

Table 3: Limiting tradeoffs at $a = .01$, $\lambda = 4$ (numerical quadrature of the limit laws; the Bessel form (11) and direct quadrature agree). Both regimes have $R/h = O(1)$; what decides is the remaining headroom at detection. Between them, headroom $m = 1, 2$ gives limiting miss 6.108% and 5.027%, and $m \geq 3$ gives a limiting window.

4.3 A uniform theorem at every threshold of at least one million

The limits above are qualitative. The following statement is quantitative, holds at every large threshold with explicit deadlines, and includes the identity repair. Let $k = \lfloor h/2 \rfloor$, $\ell = h - k$, $c_h = h/(h+a)$, and

$$\mu_-(h) = \sum_{z=0}^{k-1} \frac{1}{h-z} + \sum_{z=k}^{h-1} \frac{1}{z+a}, \quad A_h = \log(k+a) + \log(\ell+1) + \mu_-(h), \quad (12)$$

$$\tau_-(h) = c_h[A_h - \log(.15)], \quad \tau_J(h) = c_h[A_h + \log 200], \quad (13)$$

$$\mu_+(h) = \sum_{z=0}^{h-1} \frac{z}{(z+a)(2h-z)}, \quad \tau_+(h) = \log(h+a) + \mu_+(h) - \log(.3), \quad (14)$$

all in units of $1/g$.

Theorem 4.4 (Uniform large-threshold decision separation; ordinary proof with directed scalar certificates). *Let $a = .01$, $\lambda = 4$, $\alpha = .01$, and let $h \geq 10^6$ be an integer.*

- (a) *At $R = h$, $F(\tau_-) > .01086$ and $M(\tau_-) > .0563$; hence, by Corollary 3.3, every measurable randomized crossing-time classifier with blank error at most .01 has miss greater than .0563.*
- (b) *At $R = h$, a source-preserving identity measurement available at τ_J with blank-hit positive probability at most .01 and occupied-hit positive probability at least .99 gives a hit-and-positive rule with blank error below .00085 and miss below .0413. Its miss advantage over every admissible crossing-time classifier exceeds 1.5 percentage points.*
- (c) *At $R = 2h$, the deadline τ_+ has blank error below .009373 and miss below .048732, so timing alone meets the limits.*

The proof is in Appendix C. Its ingredients are the couplings of Lemma 4.1 applied to the two blocks of Theorem 4.3, exponential concentration for the auxiliary gamma variables and for the weighted remainder, and a handful of scalar enclosures evaluated with directed rounding (Appendix D). The earlier version of this argument in the companion manuscript used Chebyshev bounds whose $O(1/h)$ allowances forced $h \geq 10^8$; replacing them by gamma Chernoff and weighted Bernstein bounds makes the allowances e^{-ch} and lowers the sufficient threshold to 10^6 . That constant remains a convenient sufficient value; Theorem 3.1 covers $h = 5$, Table 2 covers $h = 20$, and Figure 2(b) covers the range between numerically. Table 4 displays the explicit times.

Threshold h	τ_- (exhausted; separating)	τ_J (exhausted; joint readout)	τ_+ (reserve; usable deadline)
10^6	29.52814	32.92934	15.71263
10^8	38.73848	42.13968	20.31780

Table 4: Numerical evaluations of the exact formulas (13)–(14) in units of $1/g$. The times grow like $2 \log h$ and $\log h$; no experimentally achieved turnaround follows without calibrating g and the acquisition time. At $h = 10^6$ the limiting conservative joint rule at product threshold $y = .005$ has miss 3.18421% and blank error .04089% (numerical), which explains the margin the coarser proved bounds of part (b) leave.




Figure 2: (a) The limiting tradeoffs of Table 3 (curves and triangles, numerical) with the proved finite-threshold bounds of Theorem 4.4 for every $h \geq 10^6$ (square and circle). (b) Miss at the 1% blank boundary of the optimal timing rule as the threshold grows, for fixed headroom $R = h, \dots, h + 3$ and for $R = 2h$, from Laplace-transform inversion [1, 28] in the companion manuscript; dashed lines are the limits. Numerical.

4.4 What the large-threshold behaviour resolves

The scale objection is resolved within the declared family: an all-timing obstruction can persist at arbitrarily large thresholds, its cause being the order-one holding time of the last births before exhaustion together with the startup spread of the blank. It is not resolved by asserting that all physical assays operate in the exhausted regime. At $h = 10^6$, fixed headroom of a few units means a threshold extremely close to the modeled capacity, and whether a detector, a signal map and a kinetic law place an assay there is an empirical question (Section 7). The reserve regime is as important to report as the obstruction: the same theory identifies conditions in which timing works, so it should not be read as universal necessity of an identity measurement. The displayed limits keep a and λ fixed while the threshold grows; they are not statements about every joint scaling of loading, background, resource and noise.

5 Effective capacity: what intermediate observations identify

5.1 Exact endpoint nonidentification

Proposition 5.1 (Endpoint nonidentification; Lean-verified). *For blanks at threshold three, $(b, g, R) = (2, 1, 6)$ and $(8/3, 2/3, 4)$ have identical complete crossing-time probability measures.*

Proof. The ordered pre-hit rates are $(2, 5/2, 8/3)$ and $(8/3, 5/2, 2)$. The crossing time is a sum of independent exponential holds with these rates, and convolution is commutative (endpoint_law_nonidentification, an equality of measures, not of transforms or sampled values). $\square$

Unlimited endpoint data cannot identify capacity in this family without restrictions; this is structural, not a lack of power. The example changes $a = b/g$ and does not assert nonidentification when a is known, for loaded preparations, or for the intermediate trajectory. Figure 3(a) renders the equal laws.

5.2 Paired calibrated waits

Suppose a is known, calibrated states $0 \leq i < j < R$ are observed, and each reaction has an unobserved constant positive clock multiplier C_0 so that its rates are $C_0 q_z$. With independent unit exponentials, $D_z = E_z / (C_0 g(a + z)(1 - z/R))$, and the clock cancels pathwise in D_j/D_i (clock_cancel). Writing q_i, q_j for the rates without C_0 ,

$$p = \mathbb{P}(D_j > D_i) = \frac{q_i}{q_i + q_j}, \quad \kappa = \frac{1 - p}{p} \frac{a + i}{a + j} = \frac{R - j}{R - i}, \quad R = \frac{j - i\kappa}{1 - \kappa}, \quad 0 < \kappa < 1 \quad (15)$$

(paired_exceedance_formula, reconstruct_from_paired_probability). For $a = .01$, $i = 2$, $j = 4$ the exact probabilities are 603/1004, 335/736 and 268/669 at $R = 5, 7, 10$. The innovations must have the declared independent exponential law; a random constant clock does not affect the ratio, but completion selection does (Section 6.2).

5.3 Calibrated blocks instead of resolved jumps

Fluorescence does not resolve single count jumps. Let integers $k_0 < k_1 < k_2$ be calibrated states, $R > k_2 - 1$, and define the elapsed times across two consecutive blocks,

$$U_R = \sum_{z=k_0}^{k_1-1} \frac{E_z}{C_0 g(a + z)(1 - z/R)}, \quad V_R = \sum_{z=k_1}^{k_2-1} \frac{E_z}{C_0 g(a + z)(1 - z/R)}, \quad Q_R = \frac{V_R}{U_R}. \quad (16)$$

Theorem 5.2 (Ordered block ratios; ordinary proof, algebra Lean-verified). Q_R is invariant to the common positive clock. Under common positive innovations, $R_2 > R_1 > k_2 - 1$ implies $Q_{R_2} < Q_{R_1}$. Hence $p_r(R) = \mathbb{P}(Q_R > r)$ is nonincreasing in R for every $r > 0$.

Proof. Set $w_z(R) = R/((a+z)(R-z))$ and $m_z = w_z(R_2)/w_z(R_1)$. For $z < y$,

$$m_z - m_y = \frac{R_2(R_2 - R_1)(y - z)}{R_1(R_2 - z)(R_2 - y)} > 0.$$

Expanding,

$$V_{R_1}U_{R_2} - V_{R_2}U_{R_1} = \frac{1}{(C_0g)^2} \sum_{z=k_0}^{k_1-1} \sum_{y=k_1}^{k_2-1} E_z E_y w_z(R_1) w_y(R_1) (m_z - m_y) > 0,$$

a sum of positive terms (`ordered_multiplier_term`, `block_ratio_bound`). Dividing by the positive denominators orders the ratios; event inclusion orders the probabilities. $\square$

For the exponential source the ordering probability is computable without simulation. With transient block generators Q_U, Q_V , initial rows α_U, α_V and exit vector d_U , integrating the density of U against the survival of V at rU gives the phase-type expression [20]

$$p_r(R) = (\alpha_U \otimes \alpha_V) [-(Q_U \oplus rQ_V)]^{-1} (d_U \otimes \mathbf{1}). \quad (17)$$

The triangular generators have negative eigenvalues, and rational rates give exact rational linear systems: for blocks $[1, 3]$ and $[3, 5]$, $a = .01$, $r = 1$, $p_1(7) = 111503293/331058412$, and the values at $R = 6, 7, 10, 20, 100$ are .387971, .336809, .275672, .230952, .206321 (Figure 3(b)).

Theorem 5.3 (Strict probability ordering and unique inversion; ordinary proof). *In the independent-exponential block experiment, for every $r > 0$ the function p_r is continuous and strictly decreasing on $(k_2 - 1, \infty)$, with $p_r(R) \rightarrow 1$ as $R \downarrow k_2 - 1$ and a limit $p_r(\infty) \in (0, 1)$ as $R \rightarrow \infty$. Hence every probability in $(p_r(\infty), 1)$ determines a unique capacity in this family.*

Proof. Fix $R_2 > R_1$ and an innovation vector e in the positive orthant; by Theorem 5.2, $Q_{R_2}(e) < Q_{R_1}(e)$. Multiplying only the late-block coordinates of e by $s > 0$ multiplies both ratios by s , so for s in the nonempty interval $(r/Q_{R_1}(e), r/Q_{R_2}(e))$ the scaled vector satisfies $Q_{R_2} < r < Q_{R_1}$. Both ratios are continuous in e , so this holds on an open set, which has positive probability because the innovation vector has a positive density on the orthant. The event $\{Q_{R_1} > r\}$ therefore exceeds $\{Q_{R_2} > r\}$ by an event of positive probability: $p_r(R_2) < p_r(R_1)$. For continuity in R , $Q_R(e)$ is continuous in R and $\{Q_R = r\} = \{V_R - rU_R = 0\}$ is a hyperplane in e with a nonzero normal, hence null; dominated convergence gives continuity of p_r . As $R \downarrow k_2 - 1$ the weight $w_{k_2-1}(R)$ diverges while U_R stays finite, so $Q_R \rightarrow \infty$ almost surely and $p_r \rightarrow 1$. As $R \rightarrow \infty$ every weight tends to $1/(a+z)$; the limiting ratio of two independent positive continuous variables exceeds r with a probability strictly between zero and one. The intermediate value theorem gives unique inversion. $\square$

The theorem does not make a finite-sample estimate precise, identify an unknown a , or recover the mean of a capacity mixture; if a larger source restricts capacities to a smaller interval, inversion is on that interval.

Proposition 5.4 (Block ordering under general positive holding innovations; ordinary proof). *Suppose successive residence times admit a coupling $D_z(R, C_0) = A_z/(C_0 v_z d_z(R))$ with $A_z > 0$, $v_z > 0$, $C_0 > 0$, where the joint law of (A_z) is the same for the compared capacities; neither exponential marginals nor independence among the A_z is required. If for $R_2 > R_1$ the multipliers $m_z = d_z(R_1)/d_z(R_2)$ are strictly smaller at every late-block state than at every early-block state, then the block ratio cancels C_0 and strictly decreases from R_1 to R_2 pathwise. If moreover (A_z) has a positive density on the positive orthant, the strict probability ordering of Theorem 5.3 holds as well.*

Proof. Cancel C_0 , write each weight at R_2 as its weight at R_1 times m_z , and expand the cross product as in Theorem 5.2; every summand is positive. The support argument of Theorem 5.3 uses only continuity and positive density. $\square$

Concrete cases: serial multistep delays whose substeps share the factor $C_0 d_z(R)$ (Erlang or other non-exponential residence times); correlated innovations with a capacity-independent joint law; arbitrary fixed state-specific baseline kinetics v_z ; and several proportionally scaled resource pools, $d_z(R) = \prod_s (1 - z/(c_s R))^{\gamma_s}$ with fixed $c_s, \gamma_s > 0$, for which the inferred R is a common scale rather than an identification of each pool. What these extensions preserve is the ordering; they do not preserve the closed form (15), the phase-type formula (17) or the numerical timing certificates. Cases outside the proposition are direct diagnostics rather than vague realism caveats: an additive delay L_z that does not scale with the clock, $D_z = A_z/(C_0 q_z) + L_z$, generally does not cancel C_0 in a ratio; capacity-dependent innovation laws, state-specific drift and backward transitions are likewise excluded. A known time-dependent activity $q_z(t) = C_0 c(t) q_z$ with $c > 0$ is handled by the operational time $\Lambda(t) = \int_0^t c$: increments of Λ replace clock-time increments and recover holds with rates $C_0 q_z$; trajectories that do not complete within the finite total operational time must be retained as right-censored, and an unobserved completion must not be mapped to the finite value $\Lambda(\infty)$.

5.4 Conditioning and depletion strength

Put $A = a + i$, $B = a + j$, $d = j - i$, $D(R) = (A + B)R - (Ai + Bj)$. From (15),

$$p(R) = \frac{A(R - i)}{D(R)}, \quad p_\infty = \frac{A}{A + B}, \quad p(R) - p_\infty = \frac{ABd}{(A + B)D(R)}, \quad p'(R) = -\frac{ABd}{D(R)^2}. \quad (18)$$

Departure from unlimited capacity is $O(1/R)$ and sensitivity to an absolute change is $O(1/R^2)$: at large capacity the parameter is identified but poorly conditioned, which is a different phenomenon from Proposition 5.1. The depletion strength $\rho = 1/R$ gives $p(\rho) = A(1 - i\rho)/(A + B - (Ai + Bj)\rho)$ with $p'(0) = ABd/(A + B)^2 > 0$, a regular coordinate near unlimited capacity; an upper bound on ρ stays informative when the capacity set is unbounded above (Figure 3(b)). For a target bracket $[R_-, R_+]$ let $\varepsilon_p = \min\{p(R_-) - p(R), p(R) - p(R_+)\}$. By Hoeffding's inequality [14], $n \geq \log(2/\delta)/(2\varepsilon_p^2)$ independent complete paired observations place the inverted point estimate in the bracket with probability at least $1 - \delta$; to place an entire confidence interval centred at the estimate inside the bracket on its coverage event, use half the tolerance and four times the bound. With $a = .01$, $i = 2$, $j = 4$, 10% relative precision and $\delta = .05$, the point-estimate bounds are $n = 4,883$, 225,627 and 9,911,503 at $R = 7$, 20, 100: sufficient scalings $O(R^2/\varepsilon^2)$ for relative and $O(R^4/\varepsilon^2)$ for absolute precision, for this fixed-state design and without any minimax claim.

5.5 Plateau-normalized thresholds lose the leading capacity information

A tempting preprocessing step chooses the count boundaries as fractions of each reaction's own plateau. In the source this means boundaries proportional to R , and it removes what the block method is meant to retain.

Proposition 5.5 (Plateau normalization; ordinary proof). *Let $0 < u_0 < u_1 < u_2 < 1$ and $k_r(R) = \lfloor u_r R \rfloor$ in the exponential source with fixed a . As $R \rightarrow \infty$,*

$$C_0 g U_R \xrightarrow{P} L(u_0, u_1), \quad C_0 g V_R \xrightarrow{P} L(u_1, u_2), \\ L(u, v) = \int_u^v \frac{dx}{x(1-x)} = \log \frac{v}{1-v} - \log \frac{u}{1-u},$$

so $Q_R \rightarrow L(u_1, u_2)/L(u_0, u_1)$ in probability, a limit independent of capacity.

Proof. $\mathbb{E}(C_0 g U_R) = \sum_{z=k_0}^{k_1-1} \frac{1}{(a+z)(1-z/R)}$ is a Riemann sum for the displayed integral; each coefficient is $O(1/R)$ and there are $O(R)$ independent unit-exponential terms, so the variance is $O(1/R)$. Chebyshev and the positive limit of the early block give the ratio limit; the late block is identical. $\square$

For fractions .2, .4, .6 the limit is .82678021, and the ratios of expected block durations at $R = 10^2, 10^3, 10^4, 10^5$ are .81814128, .82591487, .82669366, .82677156 (Figure 3(c)). The proposition does not prove complete nonidentification, since finite-size corrections and fluctuation amplitudes may retain information, but it shows that plateau normalization cannot be substituted into Theorem 5.2: it moves the boundaries with R and changes the experiment. The method should use externally calibrated fixed bands and propagate their uncertainty (Section 6.1).




Figure 3: (a) The exactly equal blank endpoint laws of Proposition 5.1, rendered numerically. (b) Calibrated single-wait and block ordering probabilities retain capacity information under their known-shape assumptions; the dotted line is p_∞ . (c) Plateau-relative bands at 20%, 40%, 60% of capacity: the ratio of expected block durations converges to the capacity-free limit of Proposition 5.5.

6 From signal to inference with incomplete observations

6.1 Signal-to-state intervals

Begin with a declared measurement relation on a validated range, for instance $Y(t_m) = B + Kz(t_m) + \epsilon_m$ with baseline B , gain $K > 0$ per calibrated state unit and readout error ϵ_m ; a known monotone nonlinear response can replace Kz with interval inversion. If calibration provides $B \in [B_L, B_U]$, $K \in [K_L, K_U]$ with $K_L > 0$, and a simultaneous error allowance $|\epsilon_m| \leq \nu$ on the claimed event, then a recorded value y implies

$$z_L = \max\left(0, \frac{y - B_U - \nu}{K_U}\right), \quad z_U = \frac{y - B_L + \nu}{K_L}, \quad (19)$$

intersected with the state space and rounded outward for integer states; an empty intersection is a measurement/model incompatibility and must be reported, not clamped. For illustration only, a baseline in $[95, 105]$ RFU, a gain in $[.9, 1.1] \times 10^{-3}$ RFU per unit, an allowance of 5 RFU and a reading of 1110 RFU give $z \in [909, 091, 1, 133, 333]$: a nominal million-unit crossing is not a precisely calibrated million-state crossing, and if the scientific distinction concerns a few units of headroom this calibration is insufficient. For each fixed calibrated threshold k , a frame is certainly below when $z_U < k$ and certainly at or above when $z_L \geq k$; under the monotone latched count model the last certainly-below and first certainly-above frames bracket the crossing time, an infinite upper endpoint is retained when no above-threshold frame is

certified, and initially-above cases form their own category. If $\tau_{k_r} \in [L_r, U_r]$ then consecutive block durations satisfy

$$\Delta_r \in [\max(0, L_{r+1} - U_r), U_{r+1} - L_r]. \quad (20)$$

Shared crossing times make neighbouring block bounds dependent; the inference below does not require independence between them. Crossings bracketed by $[2, 2.1]$, $[3.5, 3.6]$ and $[4.9, 5.0]$ give $U \in [1.4, 1.6]$ and $V \in [1.3, 1.5]$, and the comparison $V > U$ is unresolved; replacing the intervals by midpoints would discard exactly the uncertainty that matters. A pointwise calibration band is not a simultaneous band over thousands of frames; a shared calibration failure probability δ_{cal} is combined with the inference failure allowance by a union bound, and per-reaction bad-bracket probabilities are handled in Theorem 6.1.

6.2 The selection failure

For $a = .01$, $i = 2$, $j = 4$, $R = 7$, the full-cohort ordering probability is $p = 335/736 = .455163$ for every constant clock. Conditioning on the sum-of-waits completion event $D_i + D_j \leq 1$ changes it to .491088, .479803 and .459864 at clock multipliers .4, 1 and 3: with rates including the clock, the selected numerator is $\int_0^{1/2} q_i e^{-q_i u} (e^{-q_j u} - e^{-q_j(1-u)}) du$ and the completion probability $\int_0^1 q_i e^{-q_i u} (1 - e^{-q_j(1-u)}) du$ (Figure 4(a)). Pathwise cancellation before selection does not imply invariance in a selected subpopulation, so deleting slow observations defeats the method; the remedy is to keep them.

6.3 Censoring-safe partial information

Theorem 6.1 (Censoring-safe confidence sets; ordinary proof). *Suppose genuine recorded bounds give $U \in [L_U, U_U]$ and $V \in [L_V, U_V]$, with infinite upper endpoints allowed. For $r > 0$ call an observation certain positive if $L_V > rU_U$, certain negative if $U_V \leq rL_U$, and unknown otherwise. If the first two categories have probabilities π_+, π_- then*

$$\pi_+ \leq p = \mathbb{P}(V > rU) \leq 1 - \pi_-. \quad (21)$$

For n independent observations with the same latent p and possibly nonidentical censoring, with counts n_+, n_- , with probability at least $1 - \delta$,

$$p \in \left[\max\left(0, \frac{n_+}{n} - e_n\right), \min\left(1, 1 - \frac{n_-}{n} + e_n\right) \right], \quad e_n = \sqrt{\frac{\log(2/\delta)}{2n}}. \quad (22)$$

Independence of the censoring from the latent waits is not needed. If the bracket of observation i is invalid with probability at most ϵ_i , the statement holds with e_n replaced by $e_n + n^{-1} \sum_i \epsilon_i$.

Proof. The certain-positive event lies inside $\{V > rU\}$ and the certain-negative event inside its complement, which proves (21); the width of the interval equals the unknown mass, and it is sharp given only the category probabilities, since assigning every unknown outcome to either latent class attains an endpoint. For (22), apply one-sided Hoeffding bounds [14] to the two independent indicator sums around their average expectations, each with tail $\delta/2$, and use a union bound; each average expectation obeys (21). With invalid brackets, a certain-positive call can exceed the true positive event only on the bad-bracket event, so the category probability can exceed p by at most ϵ_i , and similarly for certain negatives; average over i . $\square$

Observation vectors must be independent across reactions; shared run effects need a justified independent unit or a conditional model. A shared nuisance is harmless when, conditionally on it, the observations are independent with the same latent p , as for a shared constant clock under the exact invariance model; arbitrary run effects are not. Continuous latent waits permit the displayed strict and nonstrict rules; recording atoms need conservative boundary handling.

6.4 Capacity sets

For single waits, transform the confidence set by (15), intersecting with $0 < \kappa < 1$; since κ decreases in p and increases in a , an uncertain $a \in [a_L, a_U]$ has conservative extreme values at (p_U, a_L) and (p_L, a_U) . An empty physical intersection is a model incompatibility at the stated confidence level; an interval reaching $\kappa = 1$ has unbounded upper capacity; $\kappa = 0$ corresponds to $R = j$, excluded for a positive dwell at j . For blocks, take $\{R : p_r(R) \in [p_L, p_U]\}$ with validated enclosures of (17); Theorem 5.3 makes this an interval, and Theorem 5.2 alone still supports brackets. The output may be a finite interval, a one-sided bound, the whole admissible domain, or an empty set, and none of these is to be replaced by a forced point estimate.

Simulation	R	n_+	n_-	Unknown	95% set for p	95% capacity set
Resolved waits	7	1752	2229	19	[.416527, .464223]	[6.745, 8.715]
Short follow-up	7	478	639	2883	[.098027, .861723]	[4.175, ∞)
Large capacity	100	1367	2600	33	[.320277, .371473]	[15.167, ∞)

Table 5: Prespecified seeded simulations (seed 20260916), $n = 4000$ each, $a = .01$, $i = 2$, $j = 4$, clocks drawn independently from $\{.4, 1, 3\}$; every observation is retained. Each designated dwell has its own horizon (8 with frame width .002 in rows one and three; .12 with frame width .05 in row two), unlike the sum-of-waits selection of Section 6.2. The last row gives $\rho \in [0, .06594]$ although the capacity set is unbounded. Simulations illustrate the theorem; they neither prove coverage nor validate anything empirically.




Figure 4: (a) Completion selection restores clock dependence (Section 6.2). (b) Probability confidence sets of Table 5 with the generating probabilities and the returned status.

6.5 A usable decision procedure

Suppose $[F_L, F_U]$ and $[M_L, M_U]$ simultaneously contain both errors over the candidate readouts. A readout is *usable* if $F_U \leq \alpha$ and $M_U \leq \beta$; *excluded* if $F_L > \alpha$ or $M_L > \beta$; and *unresolved* otherwise. Containment makes the first two verdicts sound, selection among simultaneously covered candidates preserves the guarantee, and for an actual time in $[t_-, t_+]$ the blank error is bounded at t_+ and the miss at t_- (**usable_sound**, **excluded_sound**, **selected_time_sound**, **jitter_endpoints**). The inherited capacity-ten instance at time 3.2 is usable; the coarse bands $F \in [0, .02]$, $M \in [0, .1]$ leave that same instance unresolved; and at capacity five the miss lower bound at time five excludes that deadline (**three_decision_examples**). Wide evidence and physical impossibility are different verdicts, and collecting better evidence can resolve the former but not the latter.

A researcher-facing version of the analysis, implementable without reading any formal proof, is the following.

- (1) Declare the blank and loaded populations, the loading model including empties, the independent unit, the error limits, and the available readout times and measurements.
- (2) Retain noncrossers and incomplete observations. Obtain model or empirical error enclosures for each candidate time, including the source uncertainty and the timing allowance, and return usable, excluded or unresolved.
- (3) If every deadline is excluded, the exclusion applies to every crossing-time classifier (Theorem 3.2). Evaluate a candidate identity channel through $fF(t) \leq \alpha$ and $sC(t) \geq 1 - \beta$ with hit-conditioned f and s , and its acquisition time; if a change of reserve is within scope, compare the resource regimes of Section 4.
- (4) For mechanism questions, convert calibrated frames into state and crossing intervals (19)–(20), keep the unknown category, and return a probability confidence set and its capacity or depletion-strength image.
- (5) Report one finite interval and one one-sided or unresolved case explicitly, and state which additional information would narrow the latter.

For a prespecified rule and independent Bernoulli trials, zero errors in n trials give the one-sided upper bound $1 - \gamma^{1/n}$ at tail probability γ [6, 10]. Allocating $\gamma = .025$ to each error, 368 blank and 72 loaded zero-error trials give simultaneous 1% and 5% upper limits, and a .03% bound needs 12,295 eligible trials. These are planning calculations. Conditional marker calibration uses the hit-conditioned denominator, and selecting a time or a rule from the same data needs simultaneous coverage or an independent evaluation set.

7 Scope and the empirical frontier

Three consequences change an analysis. An all-rule timing certificate ends futile cutoff tuning, and the deadline-optimality theorem says that no cleverer function of the same datum will do better. Equal endpoint laws defeat unlimited endpoint replication, while labeled intermediate observations retain rate order, and the strict inversion theorem tells what a block probability determines. Incomplete observations support an honest one-sided capacity bound instead of a biased complete-case estimate.

The results are conditional on the declared source, and Table 6 lists what connects each model object to a measurement. The primary empirical question is whether a realizable identity measurement meets the hit-conditioned bounds while preserving the total-count source across its operating range; the assumed 99%/1% capabilities should not be called standard performance without physical evidence. Calibrated state intervals with propagated band uncertainty, a constant within-reaction clock, and the identity of the consumed resource behind the effective capacity are the other substantive assumptions; an effective capacity fit is not a mechanism, and inversion of a capacity mixture does not recover its mean. The headroom prediction of Section 4 is testable: under a controlled resource intervention with matched blanks, known loading, all trajectories retained and independent run labels, the sensitivity–blank tradeoff should change with the remaining headroom at detection and should disappear when the reserve is large, whereas a common clock change only shifts all crossing times.

Open problems. (1) A measured interpretation of effective capacity, background and growth, and the headroom intervention. (2) Sharper constants: the certified 5.55% and 6.9% against the numerical 8.96%, and the sufficient threshold 10^6 . (3) Formalization of Theorem 3.2, of the holding-time and beta–gamma arguments, and of the full probabilistic block theorem, which would raise Sections 4–5 to the evidence level of Theorem 3.1. (4) Sources with unequal target and background kinetics, activation delays, extinction, time-dependent rates of unknown form, and adaptive stopping, each of which changes the source and needs its own separation analysis. (5) Specimen-level rules that aggregate several reactions.

Model object	Candidate measurement	Required bridge
State z , threshold h	Effective product amount	Calibrated signal/state map with uncertainty (19); small thresholds are illustrative.
Capacity R , headroom	Growth suppression near detection	Rate-law fit; separate mechanistic evidence for a consumed resource; the headroom intervention above.
Scale g^{-1}	Physical elapsed time	Rate calibration and acquisition duration.
Poisson loading	Target allocation	Loading distribution and a denominator including empties.
Identity channel	Product-sensitive readout	Hit-conditioned errors against an independent product-identity reference; availability at the readout time; no kinetic effect.
Block intervals	Calibrated signal crossings	State bounds, latching convention, frame resolution, retained missingness.
Independent unit	Reaction, run, batch	Independent identifiers or a justified conditional model.

Table 6: Source-to-observation requirements. No row is empirically established by the theoretical results.

8 Verification and reproducibility

Lean-verified. The formal development lives in the namespaces `DiagnosticWindows` and `AssayInformation` of a frozen Lean 4 project (toolchain `leanprover/lean4:v4.30.0`, pinned Mathlib manifest with SHA-256 prefix `a8df0b60`), compiled with warnings treated as errors; every strict receipt lists only the axioms `propext`, `Classical.choice` and `Quot.sound`, and no `sorry`, asserted source bridge or unchecked numerical solver closes any formal root. The combined declaration `AssayInformation.main` (source SHA-256 prefix `6fce3549`, 34 imported modules, 67 s) asserts: the all-classifier obstruction $\text{Miss} > 111/2000$ for every measurable ϕ with values in $[0, 1]$ and $\text{FP} \leq 1/100$; the joint fixture bounds $\text{FP}_J < 3/10000$ and $\text{Miss}_J < 422443/10^7$; the three decision-band examples; the endpoint-law equality; and the paired-observation reconstruction identity. The extension modules `DecisionGap` (SHA-256 prefix `d2c3b4ac`) and `BlockRatio` (`c777cd3d`) supply the strict gap, the two scalar robustness consequences and the block algebra. Appendix E maps claims to declarations.

Ordinary proof. Lemma 2.2 and the source equivalence of Appendix A; Theorem 3.2 and Corollary 3.3; Propositions 2.1, 3.5, 3.6 and 3.7 (beyond their compiled scalar parts); all of Section 4, including Theorem 4.4 with Appendices C–D; Theorems 5.2 (beyond its compiled algebra), 5.3 and 6.1; Propositions 5.4 and 5.5. None of these is labeled Lean-verified.

Exact and directed computation. Every printed enclosure of Sections 3 and 4 and of Appendices B–D, the rational block probabilities, the capacity sets and the design counts are replayed by `check_paper.py` (102 named checks, exact fractions or Python `decimal` with explicit rounding directions, cross-checked with `mpmath`); it also regenerates the coefficient table of Appendix B.

Numerical illustration. The deadline curves, the optimum $t_\alpha = 4.4594$ with miss 8.9566%, the limit-law values of Tables 3 and 4, the convergence panel of Figure 2(b), the numerical column of Table 2, the selection regression and the seeded simulations of Table 5. The figures are produced by `make_figures.py`. The repository location will be supplied with the author information.

A The uniformized source and the holding-time equivalence

Let Q be the $(h+1)$ -state generator with $Q_{zz} = -q_z$, $Q_{z,z+1} = q_z$ for $z < h$ and a zero absorbing row at h . For $\nu \geq \max_z q_z$ put $P = I + Q/\nu$, a stochastic matrix; the uniformized transition matrix is $K_t = e^{-\nu t} \sum_{k \geq 0} \frac{(\nu t)^k}{k!} P^k = e^{tQ}$ [15]. The independent-holding-time construction has the same generator, since exponential memorylessness gives the first-step transition and holding rate at every state; its transition matrix solves $K'_t = K_t Q$, $K_0 = I$, and uniqueness of this finite linear system identifies it with e^{tQ} . Writing $S_n(t) = \mathbb{P}(T > t \mid X_0 = n)$, one has $S_h = 0$, $S_n(0) = 1$ for $n < h$, $S'_n = -q_n S_n + q_n S_{n+1}$, and for distinct rates

$$S_n(t) = \sum_{i=n}^{h-1} c_{ni} e^{-q_i t}, \quad c_{ni} = \frac{\prod_{j=n}^{h-1} q_j}{q_i \prod_{j=n, j \neq i}^{h-1} (q_j - q_i)}, \quad (23)$$

with $F(t) = 1 - S_0(t)$ and $M(t) = e^{-\lambda} \sum_{n < h} \frac{\lambda^n}{n!} S_n(t)$; the omitted terms $N \geq h$ are the initial atom. The concrete Lean source uses $\nu = 5/2$ and rational transition probabilities at threshold five, and verifies the spectral identities algebraically inside the uniformized law. For rigorous enclosures at arbitrary rational times, compute $P^k \mathbf{1}_{z < h}$ recursively with exact fractions, truncate the Poisson series at L terms, bound the omitted tail through the enclosure $1/(A_L(x) + R_L(x)) \leq e^{-x} \leq 1/A_L(x)$ with $A_L(x) = \sum_{k \leq L} x^k/k!$ and $R_L(x) = \frac{x^{L+1}}{(L+1)!} \frac{1}{1-x/(L+2)}$, and combine with the finite loading sum. The fixed-rate window uses $L = 200$ at the effective times 1.02×7.01 and $.98 \times 6.99$ with loading 3.99; the threshold-twenty example uses $L = 300$ at times eight and ten.

B The sign certificate for the timing obstruction

For (3) the rates in state order are $(5, 404, 603, 602, 401)/500$. Differentiating (23),

$$f_0(t) = \sum_i q_i c_{0i} e^{-q_i t}, \quad f_1(t) = e^{-4} \sum_i q_i B_i e^{-q_i t}, \quad B_i = \sum_{n=0}^4 \frac{4^n}{n!} c_{ni}.$$

With $x = e^{-t/500}$, $f_1(t) - 5f_0(t) = x^5 P(x)$ where

$$P(x) = -A + Bx^{396} - Dx^{399} + Ex^{597} - Gx^{598} = x^{396} H(x) - A, \\ H(x) = B - Dx^3 + x^{201}(E - Gx).$$

Table 7 gives outward terminating-decimal enclosures of the five signed coefficients, reconstructed from the rational c_{ni} and an enclosure of e^{-4} ; substituting lower endpoints gives P_- and upper endpoints P_+ , with $P_- \leq P \leq P_+$ for $x \geq 0$. *Late region*. Put $v = 9901/10000$.

Exponent	Lower coefficient	Upper coefficient
0	-0.051936436911	-0.051936436910
396	133.156300607087	133.156300607088
399	-135.013385312608	-135.013385312607
597	127.123028975878	127.123028975879
598	-125.057323647968	-125.057323647967

Table 7: Outward coefficient enclosures (regenerated by `check_paper.py`); the formal proof uses the same source identities with a slightly different exact enclosure and the same four-cell argument.

The upper coefficients satisfy $D, G \geq 0$, $396B - 399Dv^3 \geq 0$, $597E - 598Gv \geq 0$ and

$P_+(v) < -.0021371$; since $P'_+(x) = x^{395}\{396B - 399Dx^3 + x^{201}(597E - 598Gx)\} \geq 0$ on $[0, v]$, $P_+ < 0$ there. For $t \geq 5$, $e^{t/500} \geq 1.01$ gives $x \leq 100/101 < v$. *Early region.* Divide $[.992, 1]$ at $.994, .996, .998$. On a cell $[l, u]$ the lower coefficients satisfy $D, G \geq 0$, $E - Gu \geq 0$ and $H'(x) \leq -3Dl^2 + 201u^{200}(E - Gl) - Gl^{201} < 0$; also $H(u) > 0$, so $P_-(x) \geq l^{396}H(u) - A$, which exact checks bound below by $.0061843, .0313825, .0503130$ and $.0424808$ on the four cells (with H' bounded above by $-238.51, -183.57, -112.10, -21.56$). For $0 \leq t \leq 4$, $x \geq 1 - t/500 \geq .992$. This proves (6); the enclosures $F(4) > .0073$ and $M(5) > .069$ come from Appendix A. The displayed decimals are conservatively rounded reports of exact rational checks, and their numerical appearance is not the evidence.

C Proof of Theorem 4.4

Fix $a = .01$, $\varepsilon = .01$, $n_* = 25$, and let $h \geq 10^6$. All times are in units of $1/g$.

Concentration inequalities. For $Z \sim \text{Gamma}(s, 1)$ with real $s > 0$ and $0 < \varepsilon \leq 1$, exponential Markov bounds with the optimal parameters $t = 1 - e^{-\varepsilon}$ and $t = e^\varepsilon - 1$ give

$$\mathbb{P}(|\log(Z/s)| > \varepsilon) \leq e^{-s(e^\varepsilon - 1 - \varepsilon)} + e^{-s(e^{-\varepsilon} - 1 + \varepsilon)} \leq 2e^{-s\varepsilon^2/3}, \quad (24)$$

using $e^\varepsilon - 1 - \varepsilon \geq \varepsilon^2/2$ and $e^{-\varepsilon} - 1 + \varepsilon \geq \varepsilon^2/2 - \varepsilon^3/6 \geq \varepsilon^2/3$. For independent unit exponentials and $S = \sum_r a_r(E_r - 1)$ with $0 \leq a_r \leq b_*$ and $V = \sum_r a_r^2$,

$$\mathbb{P}(|S| \geq \varepsilon) \leq 2 \exp\left[-\frac{\varepsilon^2}{2(V + b_*\varepsilon)}\right] : \quad (25)$$

for $0 < t < 1/b_*$, $\log \mathbb{E}e^{tS} = \sum_r [-\log(1 - ta_r) - ta_r] \leq \frac{t^2V}{2(1 - tb_*)}$, and $t = \varepsilon/(V + b_*\varepsilon)$ gives the upper tail; the lower tail is smaller since $-\log(1 + x) + x \leq x^2/2$ [5].

Elementary bounds on the limit laws. (i) $99 \leq \Gamma(a) \leq 100$: with $X \sim \text{Exp}(1)$, $\Gamma(a) = \mathbb{E}[X^a]/a$, concavity gives $\mathbb{E}X^a \leq 1$, and $X^a \geq 1 + a \log X$ with $\mathbb{E} \log X \geq -1$ gives $\mathbb{E}X^a \geq 1 - a$. (ii) From $w^a \leq 1 - a + aw$,

$$\mathbb{P}(\text{Gamma}(a, 1) \geq y) \leq .01E_1(y) + \frac{e^{-y}}{9900}, \quad E_1(y) = \int_y^\infty \frac{e^{-w}}{w} dw. \quad (26)$$

(iii) For $n \geq 1$, $\text{Gamma}(n + a)$ stochastically dominates $\text{Gamma}(n)$ and $\mathbb{P}(\text{Gamma}(n, 1) \leq y) \leq y^n/n!$; bounding the $n = 0$ term by one,

$$\sum_{n \geq 0} \mathbb{P}(N = n) \mathbb{P}(\text{Gamma}(n + a, 1) \leq y) \leq e^{-4} \sum_{n \geq 0} \frac{(4y)^n}{(n!)^2}. \quad (27)$$

(iv) For $Y \sim \text{Exp}(1)$ independent of $W \sim \text{Gamma}(a, 1)$, since $.97^{100} < .05$ gives $w^a \geq .97$ on $w \geq .05$,

$$\mathbb{P}(WY \geq y) = \frac{1}{\Gamma(a)} \int_0^\infty w^{a-1} e^{-w-y/w} dw \geq .0097 \int_{.05}^{10} \frac{e^{-w-y/w}}{w} dw. \quad (28)$$

(v) For $W_n \sim \text{Gamma}(n + a, 1)$, by $1 - e^{-u} \geq u/(1 + u)$ and Jensen,

$$\mathbb{P}(W_n Y \leq y) = \mathbb{E}[1 - e^{-y/W_n}] \geq \mathbb{E}\left[\frac{y}{W_n + y}\right] \geq \frac{y}{n + a + y}. \quad (29)$$

(vi) For $n \geq 2$, $\mathbb{P}(W_n Y \leq v) = \mathbb{E}[1 - e^{-v/W_n}] \leq v \mathbb{E}[W_n^{-1}] = \frac{v}{n+a-1} \leq \frac{v}{n-1}$; for $n = 1$, W_1 dominates $\text{Exp}(1)$, and a product of two independent unit exponentials is at most v only if one factor is at most $\sqrt{v}$, so $\mathbb{P}(W_1 Y \leq v) \leq 2\sqrt{v}$. (vii) $\mathbb{P}(N > 25) \leq e^{-4} \frac{4^{26}}{26!} (1 - \frac{4}{27})^{-1} < 3 \times 10^{-13}$.

The exhausted source. With $c_h = h/(h+a)$, $k = \lfloor h/2 \rfloor$, $\ell = h - k$, Lemma 2.2 gives $T_{h,h,n}/c_h = U_n + V_h + C_n$ with $U_n = \sum_{z=n}^{k-1} E_z/(z+a)$, $V_h = \sum_{z=k}^{h-1} E_z/(h-z)$ and $C_n = \sum_{z=n}^{k-1} E_z/(h-z) + \sum_{z=k}^{h-1} E_z/(z+a)$. For $n \leq 25$ the couplings of Lemma 4.1 give $U_n = \log Z_1 - \log W_n$ and $V_h = \log Z_2 - \log Y$, with $Z_1 \sim \text{Gamma}(k+a, 1)$, $Z_2 \sim \text{Gamma}(\ell+1, 1)$, $W_n \sim \text{Gamma}(n+a, 1)$, $Y \sim \text{Exp}(1)$, and W_n, Y independent; the Z 's need not be independent of anything, and only a union bound is used. Every coefficient of C_n is at most $1/k$ and there are at most h terms; with $k \geq (h-1)/2$ and $h \geq 10^6$, $V \leq h/k^2 \leq 4.01/h$, $b_* \leq 1/k \leq 2.01/h$, and $k+a, \ell+1 \geq h/2.01$; also $|\mathbb{E}C_n - \mu_-(h)| \leq 25/(h-24)$. Applying (24) to Z_1, Z_2 and (25) to $C_n - \mathbb{E}C_n$ with $\varepsilon = .01$, the failure probability is at most

$$\Delta_-(h) = 4e^{-h/60300} + 2e^{-h/80602} < 10^{-5}, \quad (30)$$

decreasing in h and equal to 8.43×10^{-6} at 10^6 . Outside the failure event, since $3\varepsilon + 25/(h-24) < .0301$,

$$|T_{h,h,n}/c_h - A_h + \log(W_n Y)| < .0301 \quad (n \leq 25). \quad (31)$$

(a) *Timing obstruction.* At τ_- , on the good event, $W_0 Y \geq .155 > .15e^{.0301}$ forces $T \leq \tau_-$, and $W_n Y \leq .145 < .15e^{-.0301}$ forces $T > \tau_-$. Hence, keeping only loads $n \leq 5$ in the second bound and using (29),

$$F(\tau_-) \geq \mathbb{P}(W_0 Y \geq .155) - \Delta_-, \quad M(\tau_-) \geq e^{-4} \sum_{n=0}^5 \frac{4^n}{n!} \frac{.145}{n + .155} - \Delta_-.$$

The directed computation of Appendix D gives $\mathbb{P}(W_0 Y \geq .155) > .01087$ and a finite sum exceeding .056383, so $F(\tau_-) > .01086 > .01$ and $M(\tau_-) > .056373 > .0563$. Thus τ_- is a separating time, every deadline is excluded, and Corollary 3.3 gives miss greater than .0563 for every crossing-time classifier with blank error at most .01.

(b) *Identity repair.* At τ_J , put $y = 1/200$. On the good event a missed occupied reaction with load $n \geq 1$ has $W_n Y < ye^{.0301} < v := .00516$. By (vi), (vii) and the exact bounds $e^{-4} < .019$, $4e^{-4} < .0733$, $\sum_{n \geq 2} \mathbb{P}(N = n)/(n-1) < .39$ and $\sqrt{.00516} < .07184$, the unconditional occupied-miss mass is less than

$$2(.0733)(.07184) + (.00516)(.39) + 10^{-5} + 3 \times 10^{-13} < .0126.$$

The occupied-hit mass therefore exceeds $1 - .019 - .0126 = .9684$, and with $s \geq .99$, $\text{Miss}_J < 1 - .99 \times .9684 = .041284 < .0413$. For the blank, a hit on the good event requires $W_0 Y > ye^{-.0301} > .004$, hence $W_0 \geq 1/2500$ or $Y \geq 10$, so $F(\tau_J) \leq \mathbb{P}(W_0 \geq 1/2500) + e^{-10} + \Delta_-$. By (26) with $E_1(1/2500) \leq \log 2501 < 8.4$, $F(\tau_J) < .01 \times 8.4 + 1/9900 + .00005 + 10^{-5} < .085$, and $f \leq .01$ gives $\text{FP}_J < .00085$. The advantage exceeds $.0563 - .0413 = .015$.

(c) *The reserve source.* At $R = 2h$, $T_{2h,h,n} = A_{h,n} + D_n$ with $D_n = \sum_{z=n}^{h-1} \frac{zE_z}{(z+a)(2h-z)}$, whose coefficients are at most $1/h$, so $V \leq 1/h$ and $b_* \leq 1/h$; and $A_{h,n} = \log Z - \log W_n$ with $Z \sim \text{Gamma}(h+a, 1)$. The failure probability is at most $\Delta_+(h) = 2e^{-h/30000} + 2e^{-h/20200} < 10^{-12}$, and outside it $|T_{2h,h,n} - \log(h+a) - \mu_+(h) + \log W_n| < .0201$ for $n \leq 25$. At τ_+ , a blank hit forces $W_0 \geq .3e^{-.0201} > .293$ and a loaded miss forces $W_n \leq .3e^{.0201} < .307$, so by (26), (27) and Appendix D,

$$F(\tau_+) \leq .01E_1(.293) + \frac{e^{-.293}}{9900} + \Delta_+ < .009372 + \Delta_+ < .009373,$$

$$M(\tau_+) \leq e^{-4} \sum_{n \geq 0} \frac{1.228^n}{(n!)^2} + \Delta_+ + 3 \times 10^{-13} < .048732. \quad \square$$

D Directed arithmetic for the scalar enclosures

The scalar bounds of Appendix C are evaluated in `check_paper.py` with Python’s `decimal` module at 50 digits and explicit rounding directions, with e^{-x} enclosed by correctly rounded exponentials moved one representable step outward after rounding the argument in the conservative direction, so that every printed bound is an outward bound of the exact quantity. The integral in (28) is bounded below by a lower Riemann sum on $[\cdot05, 10]$ with step $\cdot005$, using the convexity of $w + \cdot155/w$ on each cell and the right endpoint in the denominator. The exponential integral is bounded above by an upper Riemann sum on $[\cdot293, 20]$ with step $\cdot005$ plus the tail $e^{-20}/20$. The series $\sum 1.228^n/(n!)^2$ is summed to $n = 20$ with a geometric tail, and the Poisson tail uses its first omitted term with ratio $4/27$. The results are

Quantity	Directed bound
$\mathbb{P}(W_0 Y \geq \cdot155)$, lower	$\cdot0108763$
$e^{-4} \sum_{n \leq 5} \frac{4^n}{n!} \frac{\cdot145}{n + \cdot155}$, lower	$\cdot0563833$
$\cdot01 E_1(\cdot293) + e^{-\cdot293}/9900$, upper	$\cdot0093715$
$e^{-4} \sum_n 1.228^n/(n!)^2$, upper	$\cdot0487303$
$\Delta_-(10^6)$, $\Delta_+(10^6)$, upper	8.43×10^{-6} , 6.68×10^{-15}
occupied-miss mass at τ_J , upper	$\cdot0125541$
blank hit probability at τ_J , upper	$\cdot0841610$

The script asserts the coarser inequalities used in the theorem. `mpmath` at 40 digits agrees with each bound to the printed digits, and the exact limiting quantities are more favourable ($\mathbb{P}(\text{Gamma}(a) \geq \cdot293) = \cdot009247$, $\mathbb{P}(\text{Gamma}(a)Y \geq \cdot155) = \cdot011485$). The arithmetic is h -independent except through $\Delta_{\pm}(h)$, which decrease with h , so the theorem holds for every $h \geq 10^6$ without testing thresholds individually.

E Claims-to-declarations map

Claim	Lean declarations and formal boundary
Uniformized source, loading mixture, monotonicity, rational spectral formulas, exponential enclosures at threshold five	<code>DiagnosticWindows.birthKernel</code> , <code>loadedSurvival</code> , <code>spectral_survival</code> , <code>spectral_antitone</code> , <code>capacity5_survival</code> , <code>exp_enclosure16</code>
Separating certificate at $t = 5$ and window at $R = 10$	<code>capacity5_separating_errors</code> , <code>capacity10_errors</code> , <code>capacity5_no_deadline</code> , <code>DiagnosticWindows.main</code>
Fixed-rate window at $R = 10$	<code>robust_guarantee</code>
Source densities, integrability, tails	<code>AssayInformation.CrossingDensity</code> : <code>suffixDensity_source</code> , <code>blankDensity_total</code> , <code>loadedDensity_tail_source</code>
Sign polynomial and outward coefficients (6)	<code>DensitySigns</code> : <code>early_density_order</code> , <code>late_density_order</code> ; <code>SignCertificates</code> ; <code>GroupedPolynomial</code>
Dual inequality (7) and margins	<code>TimingDual.timing_dual</code> , <code>DualMargin.dual_source_margin</code> , <code>dual_blank_lower</code> , <code>dual_miss_lower</code>
All-classifier obstruction (Theorem 3.1)	<code>FullTimingObstruction.full_timing_obstruction</code>
Occupied-hit mass, channel rescue, fixture (10)	<code>IdentityRescue.occupiedHit_seven</code> , <code>identity_rescue</code> , <code>JointFixture.joint_fixture_valid</code> , <code>concrete_joint_rescue</code> , <code>RescueCertificates</code>
Strict advantage; scalar TV and fixed-rate consequences	<code>DecisionGap.decision_gap</code> , <code>discrepancy_budget</code> , <code>robust_channel_arithmetic</code> ; the measure-theoretic and coupling arguments are ordinary proof

The remaining claims concern the decision procedure, capacity, and the results that are not formalized.

Claim	Lean declarations and formal boundary
Three-way bands	<code>WindowBands.usable_sound</code> , <code>excluded_sound</code> , <code>selected_time_sound</code> , <code>jitter_endpoints</code> ; <code>WindowInstance.three_decision_examples</code>
Endpoint-law equality (Proposition 5.1)	<code>EndpointLaw.endpoint_law_nonidentification</code> ; <code>Identification.rate_reversal</code>
Paired waits (15)	<code>PairedObservation.paired_exceedance_formula</code> , <code>dwel_order_clock_cancel</code> , <code>reconstruct_from_paired_probability</code>
Block algebra (Theorem 5.2)	<code>BlockRatio.ordered_multiplier_term</code> , <code>block_ratio_bound</code> ; the probabilistic statement is ordinary proof
Theorem 3.2, Section 4, Theorem 5.3, Propositions 5.4, 5.5, Theorem 6.1	Not formalized; ordinary proof with directed scalar certificates where stated
Physical assay validation	Not established in this work

References

- [1] Joseph Abate and Ward Whitt. A unified framework for numerically inverting Laplace transforms. *INFORMS Journal on Computing*, 18(4):408–421, 2006. doi: 10.1287/ijoc.1050.0137.
- [2] Anonymous. Finite-resource amplification can eliminate diagnostic readout windows: certified small-threshold instances and large-threshold regimes. Companion manuscript, September 2026; supplies the certified capacity classification at threshold five and the Laplace-inversion numerics reused in Figure 3(b), 2026.
- [3] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 2005. doi: 10.1007/s00199-004-0514-4.
- [4] David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2): 265–272, 1953. doi: 10.1214/aoms/1177729032.
- [5] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, Oxford, 2013. doi: 10.1093/acprof:oso/9780199535255.001.0001.
- [6] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404.
- [7] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.
- [8] Paul Erdős and Alfréd Rényi. On a classical problem of probability theory. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, 6:215–220, 1961.
- [9] William Feller. *An Introduction to Probability Theory and Its Applications*, volume I. Wiley, New York, third edition, 1968.
- [10] James A. Hanley and Abby Lippman-Hand. If nothing goes wrong, is everything all right? Interpreting zero numerators. *JAMA*, 249(13):1743–1745, 1983. doi: 10.1001/jama.1983.03330370053031.
- [11] Patrick Hardinge and James A. H. Murray. Reduced false positives and improved reporting of loop-mediated isothermal amplification using quenched fluorescent primers. *Scientific Reports*, 9: 7400, 2019. doi: 10.1038/s41598-019-43817-z.

- [12] Patrick Hardinge and James A. H. Murray. Full dynamic range quantification using loop-mediated amplification (LAMP) by combining analysis of amplification timing and variance between replicates at low copy number. *Scientific Reports*, 10:916, 2020. doi: 10.1038/s41598-020-57473-1.
- [13] David Hathcock and Steven H. Strogatz. Asymptotic absorption-time distributions in extinction-prone Markov processes. *Physical Review Letters*, 128(21):218301, 2022. doi: 10.1103/PhysRevLett.128.218301.
- [14] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.
- [15] Arne Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 36:87–91, 1953. doi: 10.1080/03461238.1953.10419459.
- [16] Samuel Karlin and Herman Rubin. The theory of decision procedures for distributions with monotone likelihood ratio. *Annals of Mathematical Statistics*, 27(2):272–299, 1956. doi: 10.1214/aoms/1177728259.
- [17] David G. Kendall. On the generalized “birth-and-death” process. *Annals of Mathematical Statistics*, 19(1):1–15, 1948. doi: 10.1214/aoms/1177730285.
- [18] Erich L. Lehmann and Joseph P. Romano. *Testing Statistical Hypotheses*. Springer, New York, third edition, 2005. doi: 10.1007/0-387-27605-X.
- [19] Eugene Lukacs. A characterization of the gamma distribution. *Annals of Mathematical Statistics*, 26(2):319–324, 1955. doi: 10.1214/aoms/1177728549.
- [20] Marcel F. Neuts. *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Johns Hopkins University Press, Baltimore, 1981.
- [21] Jerzy Neyman and Egon S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A*, 231:289–337, 1933. doi: 10.1098/rsta.1933.0009.
- [22] James R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, 1997. doi: 10.1017/CBO9780511810633.
- [23] Tsugunori Notomi, Hiroto Okayama, Harumi Masubuchi, Toshihiro Yonekawa, Keiko Watanabe, Nobuyuki Amino, and Tetsu Hase. Loop-mediated isothermal amplification of DNA. *Nucleic Acids Research*, 28(12):e63, 2000. doi: 10.1093/nar/28.12.e63.
- [24] András Prékopa. On logarithmic concave measures and functions. *Acta Scientiarum Mathematicarum (Szeged)*, 34:335–343, 1973.
- [25] Justin C. Rolando, Erik Jue, Nathan G. Schoepp, and Rustem F. Ismagilov. Real-time, digital LAMP with commercial microfluidic chips reveals the interplay of efficiency, speed, and background amplification as a function of reaction temperature and time. *Analytical Chemistry*, 91(1):1034–1042, 2019. doi: 10.1021/acs.analchem.8b04324.
- [26] Justin C. Rolando, Erik Jue, Jacob T. Barlow, and Rustem F. Ismagilov. Real-time kinetics and high-resolution melt curves in single-molecule digital LAMP to differentiate and study specific and non-specific amplification. *Nucleic Acids Research*, 48(7):e42, 2020. doi: 10.1093/nar/gkaa099.
- [27] Lindsay Schneider, Hannah Blakely, and Anubhav Tripathi. Mathematical model to reduce loop mediated isothermal amplification (LAMP) false-positive diagnosis. *Electrophoresis*, 40(20): 2706–2717, 2019. doi: 10.1002/elps.201900167.
- [28] Alan Talbot. The accurate numerical inversion of Laplace transforms. *IMA Journal of Applied Mathematics*, 23(1):97–120, 1979. doi: 10.1093/imamat/23.1.97.
- [29] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi: 10.1145/3372885.3373824.

- [30] Bert Vogelstein and Kenneth W. Kinzler. Digital PCR. *Proceedings of the National Academy of Sciences of the USA*, 96(16):9236–9241, 1999. doi: 10.1073/pnas.96.16.9236.