

Finite-resource amplification can eliminate diagnostic readout windows: certified small-threshold instances and large-threshold regimes

16 September 2026

Abstract

An amplification reaction may become sensitive enough only after its background-positive probability has already exceeded the permitted limit. We study this timing conflict for an explicit finite-resource birth–immigration source: z active units grow at rate $(b + gz)(1 - z/R)$ until an aggregate detection threshold h is reached, a blank starts empty, and a loaded reaction starts with a Poisson number of units. A deadline is acceptable when the blank-positive probability is at most 1% and the missed-call probability at most 5%. At threshold five, Poisson mean four, $b = 0.01$ and $g = 1$ per minute, capacities five, six and seven admit no acceptable deadline at all, capacity eight admits one but none on a 0.1-minute observation grid, and capacity nine is the least grid-feasible capacity. Capacity ten is feasible at every deadline in $[3.15, 3.30]$ minutes, and remains feasible for every fixed rate vector within 2% of nominal, every loading mean at least 3.99 and every actual deadline in $[3.19, 3.21]$ minutes. These statements, including the source law, the monotonicity arguments and the rational exponential enclosures, are verified in Lean 4. A uniform rescaling of all rates cannot create a window; a state-dependent resource change can. We then analyse arbitrarily large thresholds through the independent holding times of the count chain and the beta–gamma algebra. If the capacity exceeds the threshold by a margin that grows without bound, depletion becomes a common additive delay and the undepleted design rules apply after re-centering. If the margin is a fixed number m of units, the last births near exhaustion add an independent $\text{Gamma}(m + 1)$ delay in logarithmic time, and the limiting error tradeoff depends only on b/g , the loading mean and m : with these parameters, at most two units of headroom leave no window, three or more do. For every threshold $h \geq 10^8$ we prove, with explicit constants and without any h -state computation, that capacity h has no acceptable deadline while capacity $2h$ has one at an explicit time; the certified instances at threshold five are therefore not an artefact of small counts or of reactions that are positive at time zero. Numerical Laplace inversion at $h = 10^8$ confirms the bounds. The large-threshold results are conventional proofs and are identified as such. The model is an effective one; no clinical calibration or specific molecular capacity intervention is established.

1 Introduction

Isothermal nucleic-acid amplification assays such as loop-mediated isothermal amplification (LAMP) [18] call a reaction positive when an aggregate signal crosses a threshold before a deadline. Two failure modes compete. Stopping too early misses reactions that contain target but have not yet amplified enough; stopping too late lets non-specific background amplification make blank reactions positive [9, 19]. In a digital format the target content of a reaction is a Poisson random variable [24], so both failure probabilities are properties of a stochastic amplification law, and the assay designer needs a deadline at which both are acceptable.

This paper treats that decision for one explicitly defined stochastic source and asks a sharper question than “which deadline is best”: *does any acceptable deadline exist at all*, and what determines the answer? Let $F_0(t)$ be the probability that a blank has called positive by time t and $M_\lambda(t)$ the probability that a reaction with Poisson mean loading λ is still negative at t . The acceptable set is

$$\mathcal{W}_{R,\lambda} = \{t \geq 0 : F_{0,R}(t) \leq \alpha, M_{R,\lambda}(t) \leq \beta\}, \quad \alpha = 0.01, \beta = 0.05, \quad (1)$$

where R is the capacity of the source defined in Section 2. The distinction between a poorly chosen deadline and an empty $\mathcal{W}_{R,\lambda}$ matters in practice: the first is corrected by re-timing, the second only by changing the reaction. A second, more mundane distinction arises when an instrument observes on a fixed time grid: a nonempty continuous window may contain no observation frame.

Context. The tradeoff itself is established experimental territory. Rolando and colleagues measured amplification efficiency, background and time to threshold across temperatures and polymerases in real-time digital LAMP [19], and later combined real-time kinetics with melt curves to separate specific from non-specific products [20]; the latter observation rule uses information that a timing-only decision does not have. Hardinge and Murray related low-copy quantification to timing and replicate variance [10] and reduced false positives with quenched primers [9]. Schneider, Blakely and Tripathi modelled product sizes for electrophoretic discrimination [21], a different observable. None of these is the subject of the present paper, which is a mathematical case study, and none of the numerical parameters used below is fitted to their data.

Main results in words. The source is a pure-birth process with immigration and a finite effective resource: from count $z < h$ the next unit appears at rate $(b + gz)(1 - z/R)$, and the reaction latches positive at count h . With threshold $h = 5$, mean loading $\lambda = 4$, $b = 0.01$ and $g = 1$ per minute we prove:

- (i) capacities $R = 5, 6, 7$ admit no acceptable deadline; capacity eight admits one (3.45 minutes) but none on the grid $\{0, 0.1, 0.2, \dots\}$ minutes; capacity nine is grid-feasible at 3.3 minutes; capacity ten is feasible at every deadline in $[3.15, 3.30]$ minutes (Theorem 3.1);
- (ii) capacity ten remains feasible for every fixed state-specific rate vector within 2% of nominal, every loading mean at least 3.99 and every actual deadline in $[3.19, 3.21]$ minutes (Theorem 3.2);
- (iii) multiplying all rates by a common factor rescales the window and cannot create one (Proposition 4.1); a fixed-rate-constant mass-action realization of the source inherits the continuous capacity classification but not the grid classification (Proposition 4.2).

Items (i)–(iii) are checked in Lean 4 with Mathlib [4, 23], from the probability law down to the rational enclosures of the exponentials involved. They concern a six-state chain, and a legitimate objection is that a real detector fires at counts many orders of magnitude larger, where 37% of loaded reactions are not already positive at time zero as they are here. Section 5 answers this with conventional proofs:

- (iv) the threshold time is a sum of independent exponential holding times whose undepleted part has an exact beta law; with a growing capacity margin $R - h \rightarrow \infty$, depletion is asymptotically a deterministic additive delay and the limiting tradeoff is the undepleted one (Theorem 5.3); with a fixed margin $R = h + m$, the last births contribute an independent Gamma($m + 1$) delay in logarithmic time (Theorem 5.4), and the limiting tradeoff is a function of $(b/g, \lambda, m)$ alone;

- (v) for every integer $h \geq 10^8$, capacity $R = h$ has no acceptable deadline and capacity $R = 2h$ has one at an explicit time, with explicit error bounds (Theorem 5.5). The proof uses four scalar enclosures and no computation of dimension h .

Numerical Laplace-transform inversion, which also costs $O(h)$ rather than $O(h^2)$, corroborates the bounds at $h = 10^8$ and shows the finite-threshold tradeoffs converging to their limits.

Three kinds of evidence. Every statement below carries one of three labels. *Lean-verified* statements are theorems in the frozen Lean 4/Mathlib development described in Section 7; they depend only on the standard axioms. *Conventional* statements are proved in the text with ordinary mathematics; their scalar enclosures are replayed with directed rounding by the script `check_paper.py`. *Numerical* statements are floating-point evaluations, labelled as such, and are never premises of a theorem. Section 7 lists which results fall in which class.

What is not claimed. The units of the source are effective amplification units, not nucleotides or fluorescence; the rate constants are illustrative; the capacity is an effective inventory whose laboratory realization is an open question discussed in Section 4; the public data replay in Section 6 validates a crossing rule, not the model. Neither clinical performance nor priority over the experimental literature is asserted.

Organization. Section 2 defines the source, the probability law and the deadline sets. Section 3 states the certified threshold-five theorems. Section 4 analyses the mechanism, clock changes, loading and a mass-action realization. Section 5 contains the large-threshold theory. Section 6 discusses data, a discriminating experiment and sampling confidence. Section 7 describes verification and reproducibility, and Section 8 concludes. Appendices A–E contain the exact source calculation, the comparison lemmas, the proof of Theorem 5.5, the directed arithmetic and the claims-to-declarations map.

2 The effective finite-resource source

2.1 Counts, capacity and latch

Fix an integer threshold $h \geq 1$, an integer capacity $R \geq h$, a background coefficient $b > 0$ and a growth coefficient $g > 0$. Let $z \in \{0, \dots, h\}$ be the number of active signal units; $R - z$ is the remaining effective precursor inventory. Before detection the only transition is $z \mapsto z + 1$, at rate

$$q_R(z) = (b + gz) \left(1 - \frac{z}{R}\right), \quad 0 \leq z < h. \quad (2)$$

The detector latches positive when the count reaches h , and the reaction is then stopped: state h is absorbing. The factor b is background initiation, gz is growth proportional to the active count, and $1 - z/R$ is the resource factor. Both target-origin and background-origin units grow with the same coefficient, so their sum is a closed observable and no independence of resource-sharing lineages is assumed. There is no activation delay, extinction, optical error or between-run heterogeneity. Except where stated, $h = 5$, $b = 1/100$ and $g = 1$ per minute; time is in minutes.

A *blank* starts at $z = 0$. A *loaded* reaction has $N \sim \text{Poisson}(\lambda)$ potential active units and initially activates $\min(N, R)$ of them, consuming inventory; its stopped initial state is $\min(N, h)$. A load at or above threshold is positive at time zero; this is part of the source and is neither discarded nor renormalized.

2.2 The probability law used in the proofs

The formal development defines all probabilities through uniformization [7, 13], avoiding any unverified simulation. For the six-state instances $5 \leq R \leq 10$ take the common rate $\nu = 5/2$ per minute, which bounds every $q_R(z)$ (the largest is $q_{10}(4) = 2.406$). The kernel P_R on $\{0, \dots, 5\}$ is

$$P_R(z, z+1) = q_R(z)/\nu, \quad P_R(z, z) = 1 - q_R(z)/\nu \quad (z < 5), \quad P_R(5, 5) = 1, \quad (3)$$

with all other entries zero; every entry is nonnegative and each row sums to one. Let $f(z) = \mathbf{1}_{\{z < 5\}}$ be the uncalled indicator. The survival (uncalled) probability from state z at time t is the Poisson mixture

$$S_R(t, z) = \sum_{n=0}^{\infty} e^{-\nu t} \frac{(\nu t)^n}{n!} (P_R^n f)(z), \quad (4)$$

which is the transition law of the continuous-time chain with rates (2) [17]. The operating probabilities are

$$F_{0,R}(t) = 1 - S_R(t, 0), \quad (5)$$

$$M_{R,\lambda}(t) = \sum_{n=0}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} S_R(t, \min(n, 5)) = e^{-\lambda} \sum_{n=0}^4 \frac{\lambda^n}{n!} S_R(t, n). \quad (6)$$

The second equality is exact: loads $n \geq 5$ are absorbed at time zero and contribute no negative calls. All probabilities and error limits are per reaction.

Since the count never decreases, a positive latch at deadline t is equivalent to threshold attainment by t . On the observation grid $G = \{k/10 : k \in \mathbb{N}\}$ minutes, anchored at zero, the acceptable frames are $\mathcal{W}_{R,\lambda} \cap G$. All statements use a fixed threshold and deterministic deadlines.

2.3 Monotonicity and the separating-time principle

The survival $S_R(t, z)$ is nonincreasing in t (Appendix A, Lemma A.1), hence $F_{0,R}$ is nondecreasing and $M_{R,\lambda}$ nonincreasing. Two consequences organise every negative and every positive result below.

Lemma 2.1 (Separating time). *If some $s \geq 0$ satisfies $F_{0,R}(s) > \alpha$ and $M_{R,\lambda}(s) > \beta$, then $\mathcal{W}_{R,\lambda} = \emptyset$.*

Proof. If $t \leq s$ then $M_{R,\lambda}(t) \geq M_{R,\lambda}(s) > \beta$; if $t \geq s$ then $F_{0,R}(t) \geq F_{0,R}(s) > \alpha$. Every $t \geq 0$ falls in one of the two cases. $\square$

Lemma 2.2 (Interval structure). *$\mathcal{W}_{R,\lambda}$ is an interval, closed in $[0, \infty)$. If $s \leq u$ are feasible then every $t \in [s, u]$ is feasible.*

Proof. For $s \leq t \leq u$, $F_{0,R}(t) \leq F_{0,R}(u) \leq \alpha$ and $M_{R,\lambda}(t) \leq M_{R,\lambda}(s) \leq \beta$. Closedness follows from the continuity of the finite spectral expression (15). $\square$

Thus an empty window is certified by one separating time, and a whole interval of feasible deadlines by its two endpoints. No interpolation between sampled times is involved.

3 Certified results at threshold five

3.1 Minimum capacity and observation grid

Theorem 3.1 (Minimum capacity and timing allowance; Lean-verified). *For the source (2)–(6) with $h = 5$ and $\lambda = 4$:*

- (a) $\mathcal{W}_{5,4} = \mathcal{W}_{6,4} = \mathcal{W}_{7,4} = \emptyset$, while $3.45 \in \mathcal{W}_{8,4}$;
- (b) $\mathcal{W}_{8,4} \cap G = \emptyset$, while $3.3 \in \mathcal{W}_{9,4} \cap G$;
- (c) $[3.15, 3.30] \subseteq \mathcal{W}_{10,4}$.

Consequently the least feasible integer capacity among $R \geq 5$ is eight with continuous deadlines and nine on the prescribed grid.

Proof. Appendix A derives the exact finite spectral form of (4) and certifies the inequalities in the last column of Table 1 by rational enclosures. For capacities five, six and seven the listed time is a separating time, so Lemma 2.1 applies. For capacity eight the evaluation at 3.45 meets both limits. Every grid point is either at most 3.4 or at least 3.5: the miss bound at 3.4 and monotonicity exclude the first set, the blank bound at 3.5 the second. The capacity-nine witness 3.3 lies on the grid. These statements exhaust the admissible smaller capacities; no monotonicity of feasibility in R is used. For (c), Lemma 2.2 with the two endpoint bounds for capacity ten. $\square$

R	t (min)	$F_{0,R}(t)$	$M_{R,4}(t)$	Certified inequality
5	5.00	.01355402	.06920700	both $> \alpha, \beta$ (separating)
6	4.00	.01076152	.05773912	both $> \alpha, \beta$ (separating)
7	3.67	.01026447	.05146115	both $> \alpha, \beta$ (separating)
8	3.45	.00976997	.04923440	both $\leq \alpha, \beta$ (witness)
8	3.40	—	.05070654	miss $> \beta$ (excludes frames ≤ 3.4)
8	3.50	.01012746	—	blank $> \alpha$ (excludes frames ≥ 3.5)
9	3.30	.00940708	.04807559	both $\leq \alpha, \beta$ (grid witness)
10	3.15	—	.04857072	miss $\leq \beta$ (interval endpoint)
10	3.30	.00991967	—	blank $\leq \alpha$ (interval endpoint)

Table 1: Source evaluations behind Theorem 3.1. The decimals are numerical displays of the exact spectral formula; Lean verifies the inequalities in the last column through rational enclosures (Appendix A).

Numerically, the full windows are about $[3.4237, 3.4822]$ minutes for capacity eight (about 3.5 seconds wide), $[3.2346, 3.3824]$ for capacity nine and $[3.1029, 3.3110]$ for capacity ten (Figure 2); capacities six and seven have their sensitivity endpoint after their blank endpoint ($4.268 > 3.889$ and $3.721 > 3.632$). These root locations are not formal claims. The grid statement in (b) concerns the grid anchored at zero; a phase-shifted six-second grid could intersect the capacity-eight window, so the observation schedule is part of the decision problem.

3.2 A uniform operating tolerance

The same source permits a guarantee over a continuum of fixed rate vectors, loading means and actual observation times; it is not a statement about time-varying or random rate fluctuations within a reaction.

Theorem 3.2 (Rate, loading and timing tolerance; Lean-verified). *Let $q_z = q_{10}(z)$ be the nominal capacity-ten rates. For every fixed vector with $0.98 q_z \leq \tilde{q}_z \leq 1.02 q_z$ for $0 \leq z < 5$,*




Figure 1: The error tradeoff as the deadline varies, for capacities five to ten; the lower-left rectangle is acceptable. Additional capacity advances both desired and background amplification. Curves are numerical evaluations of the exact spectral formula; the theorem rests on the certified points of Table 1, not on plotted intersections.




Figure 2: Magnified deadline windows. Bars use numerical endpoints; circles mark feasible six-second frames; the diamond is the certified continuous capacity-eight witness and the black segment the certified capacity-ten interval. Capacity eight has a window but no frame on this grid.

every loading mean $\lambda \geq 3.99$ and every actual deadline $t \in [3.19, 3.21]$ minutes, the source with rates $\tilde{q}$ satisfies $F_{0,\tilde{q}}(t) \leq 0.01$ and $M_{\tilde{q},\lambda}(t) \leq 0.05$.

Proof. Appendix B proves, for the law (4), that increasing state-specific rates advances threshold crossing (Lemma B.1), that increasing the loading mean reduces missed calls (Lemma B.2), and that scaling every rate by c replaces t by ct (Lemma B.3). The largest allowed rate, 1.02×2.406 , is below $\nu = 5/2$, so all comparisons live under the common uniformization clock. Hence

$$\begin{aligned} F_{0,\tilde{q}}(t) &\leq F_{0,1.02q}(t) = F_{0,q}(1.02t) \leq F_{0,q}(3.2742) \leq .01, \\ M_{\tilde{q},\lambda}(t) &\leq M_{0.98q,3.99}(t) = M_{q,3.99}(0.98t) \leq M_{q,3.99}(3.1262) \leq .05, \end{aligned}$$

where the last two inequalities are rationally certified evaluations of the exact source (numerically .00973128 and .04967252). $\square$

The theorem allows 2% variation in each pre-threshold rate and ± 0.6 seconds around a nominal 3.2-minute deadline (Figure 3). It does not assert that a laboratory assay meets these tolerances or that rate uncertainty is its only uncertainty. The two worst cases need not occur under one rate vector; each bounds its own error for every admitted vector, which is what a joint operating requirement needs.




Figure 3: Worst-case comparison bounds for capacity ten: the blank curve uses rates $1.02q$, the miss curve $0.98q$ with loading 3.99 . Shading marks the certified actual-time interval; dashed lines are the error limits. Curves are numerical; the guarantee follows from the comparison lemmas and the two certified endpoint bounds.

4 Mechanism, clocks, loading and a mass-action realization

4.1 Why capacity changes feasibility and a common clock cannot

The resource factor in (2) depends on the current count. A blank traverses the whole chain from zero; a loaded reaction starts at a random higher count and traverses a shorter final segment. Changing R changes the state-dependent delays of these two populations differently and thereby the error tradeoff. By contrast:

Proposition 4.1 (Common clock invariance; Lean-verified). *If every rate is multiplied by $c > 0$, then $F_0^{(c)}(t) = F_0(ct)$, $M_\lambda^{(c)}(t) = M_\lambda(ct)$ and $\mathcal{W}^{(c)} = c^{-1}\mathcal{W}$. In particular a continuous window exists for the scaled source if and only if it exists for the original.*

Proof. Lemma B.3; the bijection $t \mapsto ct$ preserves existence. $\square$

Common acceleration can move which grid frames a window contains but cannot create a window. With depletion removed altogether (rates $b + gz$), the otherwise identical mean-four source has numerical blank and miss probabilities .00949563 and .03721034 at 2.8 minutes, an acceptable deadline; the capacity-five source has miss .20452023 at that time and, by Theorem 3.1, no acceptable deadline anywhere. An undepleted approximation therefore predicts a usable deadline where the finite-resource source has none. The undepleted comparison is numerical; the all-deadline exclusion is certified.

4.2 Loading requirements

For each fixed deadline, $M_{R,\lambda}(t)$ is nonincreasing in λ (Lemma B.2). Empty-loaded reactions follow the blank law, so retaining the $N = 0$ term of (6) gives, whenever the blank constraint holds,

$$M_{R,\lambda}(t) \geq e^{-\lambda} S_R(t, 0) \geq e^{-\lambda} (1 - \alpha). \quad (7)$$

Feasibility therefore requires $\lambda \geq \log((1 - \alpha)/\beta) = \log(0.99/0.05) \approx 2.98568$. At $\lambda = 4$ the floor is about 0.0181, well below β ; capacity five fails for kinetic reasons, not for lack of loading. Defining $\lambda_{\min}(R) = \inf\{\lambda : \mathcal{W}_{R,\lambda} \neq \emptyset\}$, the boundary is located numerically where the miss at the latest blank-admissible time equals β :

R	5	6	7	8	9	10
$\lambda_{\min}(R)$	4.9490	4.2892	4.0666	3.9562	3.8899	3.8456

At the fixed capacity-ten deadline 3.2 the threshold is about 3.9253. These are not clinical limits of detection: relating specimen concentration to activated units needs an extraction, capture and activation model.

4.3 What the capacity comparison holds fixed, and a mass-action realization

The capacity comparison in Theorem 3.1 holds λ , b and the undepleted growth coefficient g fixed. It does not assert that adding reagent to a fixed-volume assay realises that comparison. Consider the elementary reactions $P \rightarrow X$ and $P + X \rightarrow 2X$ at fixed volume V with rate constants k_b and k_g , where R is the conserved total $P + X$. Their propensity at $X = z$ is

$$q_R^{\text{MA}}(z) = k_b(R - z) + \frac{k_g}{V} z(R - z) = \kappa R(a + z)\left(1 - \frac{z}{R}\right), \quad \kappa = \frac{k_g}{V}, \quad a = \frac{k_b}{\kappa}. \quad (8)$$

Matching (2) with $g = 1$ requires $b = a$, and the physical rates are the normalized rates multiplied by κR . Two consequences follow.

Proposition 4.2 (Fixed-rate-constant family; conventional). *For the family (8) with $a = 1/100$, $\mathcal{W}_{R,\lambda}^{\text{MA}} = (\kappa R)^{-1}\mathcal{W}_{R,\lambda}$. Hence at $\lambda = 4$ and threshold five, capacities five, six and seven admit no deadline and eight is the least feasible capacity, exactly as in Theorem 3.1, with no change of k_b or k_g . The grid classification does not transfer: with $\kappa = 0.1$ and $k_b = 0.001$ (so that capacity ten has the nominal clock), capacity eight's window becomes $[4.2796, 4.3528]$ minutes, which contains the frame 4.3 (numerically blank .00969904, miss .04952343), so eight is grid-feasible in this family.*

Proof. Proposition 4.1 with $c = \kappa R$ applied capacity by capacity; the numerical window follows from Table 1's method at normalized time 3.44. $\square$

In this family a pre-threshold initial count n means $X(0) = n$ and $P(0) = R - n$; a protocol that adds n targets to a fixed inventory R has total $R + n$ and is a different loaded source, a distinction that matters when n/R is appreciable. Other suggested mechanisms (Michaelis–Menten substrate dependence, reporter saturation, finite surface occupancy, droplet-volume variation) change the source in different ways: a reporter ceiling alone changes the observation map, not the birth rates, and a saturating factor $(R - z)/(K + R - z)$ with substrate far above K has little rate effect. They must be modelled separately rather than relabelled to preserve the present law.

4.4 Initial detection and the case for large thresholds

At threshold five and mean four, $\mathbb{P}(N \geq 5) \approx 0.371163$: about 37% of loaded reactions are positive at time zero. This is part of the stated source and limits a literal interpretation in which a handful of captured molecules must amplify substantially before detection. Depletion is also substantial near detection: at $z = 4$ the resource multiplier is 0.2 for capacity five and 0.6 for capacity ten. The next section shows that neither feature is responsible for the phenomenon: at $h = 10^8$ the initially positive fraction is below 10^{-10^7} , and the contrast between $R = h$ and $R = 2h$ persists with explicit constants.

5 Large thresholds: holding times and two regimes

Throughout this section $g = 1$ (dimensionless time gt), $a = b/g > 0$, and $h \rightarrow \infty$. Everything is conventional mathematics; Section 7 lists the formal obligations that a Lean development of this section would have to discharge.

5.1 Holding-time representation

Let $T_{R,h,n}$ be the threshold time of the chain with rates (2) started at $n < h$.

Lemma 5.1 (Holding times; conventional). *Let $E_n, E_{n+1}, \dots, E_{h-1}$ be independent unit exponential variables. Then*

$$T_{R,h,n} \stackrel{d}{=} \sum_{z=n}^{h-1} \frac{E_z}{q_R(z)} = \frac{R}{R+a} \left[\sum_{z=n}^{h-1} \frac{E_z}{z+a} + \sum_{z=n}^{h-1} \frac{E_z}{R-z} \right]. \quad (9)$$

Proof. A pure-birth chain holds in state z for an $\text{Exp}(q_R(z))$ time and moves to $z+1$; successive holding times are independent by the strong Markov property [6, 17]. Equivalently, thinning the common Poisson clock of (4) at state z with probability $q_R(z)/\nu$ produces the same law. The identity $\frac{R}{(z+a)(R-z)} = \frac{R}{R+a} \left[\frac{1}{z+a} + \frac{1}{R-z} \right]$ is partial fractions. $\square$

This replaces any h -state matrix by a sum of independent terms. The first sum in (9) is the undepleted threshold time; the second is the resource delay. They share the same E_z and are not independent at finite h .

5.2 The startup law

Lemma 5.2 (Beta law of the undepleted time; conventional). *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, with $W_n \sim \text{Gamma}(n+a, 1)$,*

$$A_{h,n} - \log h \Rightarrow -\log W_n \quad (h \rightarrow \infty). \quad (10)$$

Moreover there exist, on an extension of the probability space, $Z \sim \text{Gamma}(h+a, 1)$ and $W_n \sim \text{Gamma}(n+a, 1)$ with $A_{h,n} = \log Z - \log W_n$ exactly.

Proof. For $s \geq 0$, $\mathbb{E}e^{-sA_{h,n}} = \prod_{z=n}^{h-1} \frac{z+a}{z+a+s} = \frac{\Gamma(n+a+s)\Gamma(h+a)}{\Gamma(n+a)\Gamma(h+a+s)}$, which is the s th moment of $\text{Beta}(n+a, h-n)$; the moments determine a law on $[0, 1]$. For the coupling, 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 [16] the product of independent $\text{Gamma}(h+a)$ and $\text{Beta}(n+a, h-n)$ variables is $\text{Gamma}(n+a)$. Since $Z/(h+a) \rightarrow 1$ in probability, (10) follows. (The variables Z and W_n are not independent; nothing below uses independence between them.) $\square$

The law (10) is the classical logarithmic limit of the Yule process with immigration [15]: a reaction that starts with n units behaves, after $\log h$ units of dimensionless time, like a $\text{Gamma}(n+a)$ “effective initial mass”. Its $n = 0$ member, $\text{Gamma}(a)$ with $a = 0.01$, is extremely spread; this spread is the whole difficulty of separating blanks from loaded reactions by timing.

5.3 Growing reserve: depletion becomes a common delay

Write $T = A + D$ with $D_{R,h,n} = T_{R,h,n} - A_{h,n} = \sum_{z=n}^{h-1} \frac{zE_z}{(z+a)(R-z)}$ under the coupling of Lemma 5.1.

Theorem 5.3 (Reserve regime; conventional). *Let $R = R(h) \geq h + 1$ with $R - h \rightarrow \infty$, and let $d_h = \mathbb{E}D_{R,h,0} = \sum_{z=0}^{h-1} \frac{z}{(z+a)(R-z)}$. Then for every fixed n ,*

$$T_{R,h,n} - \log h - d_h \Rightarrow -\log W_n, \quad W_n \sim \text{Gamma}(n+a, 1).$$

If $R/h \rightarrow \rho > 1$ then $d_h \rightarrow d_\rho = \log \frac{\rho}{\rho-1}$. Consequently, for deadlines $t = \log h + d_h + s$,

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

the error tradeoff of the undepleted source after a common additive delay.

Proof. $\text{Var } D_{R,h,n} \leq \sum_{z < h} (R-z)^{-2} \leq \sum_{j > R-h} j^{-2} \leq (R-h)^{-1} \rightarrow 0$, and $|\mathbb{E}D_{R,h,n} - d_h| \leq \sum_{z < n} (R-z)^{-1} \rightarrow 0$, so $D_{R,h,n} - d_h \rightarrow 0$ in probability by Chebyshev. Combine with (10). For $\rho > 1$, $d_h = \sum_{z < h} \frac{1}{R-z} - \sum_{z < h} \frac{a}{(z+a)(R-z)}$; the first sum is $H_R - H_{R-h} \rightarrow \log \frac{\rho}{\rho-1}$ and the second is $O(\log h/(R-h)) \rightarrow 0$. The limits of F_0 and M_λ follow from the continuity of the limit laws and dominated convergence over n (loads $n \geq h$ are absorbed and have vanishing probability). $\square$

5.4 Fixed headroom: persistent terminal randomness

Theorem 5.4 (Headroom regime; conventional). *Let $R = h + m$ with a fixed integer $m \geq 0$. Then for every fixed n ,*

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

where $W_n \sim \text{Gamma}(n+a, 1)$ and $Y_{m+1} \sim \text{Gamma}(m+1, 1)$ are independent. For $m = 0$, $-\log Y_1$ is a standard Gumbel variable. Consequently, for deadlines $t = 2 \log h + s$,

$$F_{0,h+m}(t) \rightarrow \mathbb{P}(W_0 Y_{m+1} \geq e^{-s}), \quad M_{h+m,\lambda}(t) \rightarrow \sum_{n \geq 0} \mathbb{P}(N = n) \mathbb{P}(W_n Y_{m+1} \leq e^{-s}).$$

Proof. Put $c = R/(R+a)$ and $k = \lfloor h/2 \rfloor$. By (9), $T/c = U_n + V + C_n$ 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 sets of exponentials and are independent. By Lemma 5.2 with h replaced by k , $U_n - \log k \Rightarrow -\log W_n$. The Mellin 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 by the same beta-gamma argument $V - \log(R-k-m) \Rightarrow -\log Y_{m+1}$ with $Y_{m+1} \sim \text{Gamma}(m+1, 1)$. The remainder has $\mathbb{E}C_n = \sum_{z < k} \frac{1}{R-z} + \sum_{z \geq k} \frac{1}{z+a} + o(1) = \log \frac{R}{R-k} + \log \frac{h}{k} + o(1) \rightarrow 2 \log 2$ 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 $(1-c) \cdot O(\log h) = O(\log h/h) \rightarrow 0$, the claim follows. For $m = 0$, $Y_1 \sim \text{Exp}(1)$ and $-\log Y_1$ is standard Gumbel [5]. $\square$

The mechanism is specific: the last $m+1$ births before exhaustion take holding times of order one in dimensionless time, and their sum, $\sum_{j=1}^{m+1} E'_j/j$ up to a shift, is an independent $\text{Gamma}(m+1)$ variable in logarithmic time that blurs the separation between blank and loaded reactions. When the unused inventory grows without bound these terminal fluctuations concentrate and disappear. Random slowdown near an absorbing boundary is a known phenomenon in absorption-time asymptotics [11]; the contribution here is its exact form for this source and its consequence for feasibility. For $m = 0$ the product law has the closed form

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

5.5 Limiting tradeoffs and the decisive role of headroom

In both regimes the limiting window is the set of shifts s for which the limiting blank and miss are acceptable; since the blank limit is increasing in s and the miss limit decreasing, it is nonempty if and only if the limiting miss at the blank boundary is at most β . Table 2 evaluates that quantity at $a = 0.01$, $\lambda = 4$, $\alpha = 0.01$ (numerical quadrature of the gamma and product-gamma laws; (11) and a direct quadrature agree to 10^{-14} at $m = 0$).

	fixed headroom $R = h + m$							reserve
	$m = 0$	1	2	3	5	10	20	$R - h \rightarrow \infty$
Limiting miss at blank boundary (%)	10.914	6.108	5.027	4.639	4.346	4.146	4.052	3.962
Blank boundary y	0.194	0.447	0.707	0.969	1.497	2.820	5.469	0.265

Table 2: Limiting tradeoffs (numerical). “Reserve” is Theorem 5.3 ($R - h \rightarrow \infty$, equivalently the undepleted source); m is the fixed headroom of Theorem 5.4. The blank boundary y solves $\mathbb{P}(W_0 Y_{m+1} \geq y) = \alpha$ (respectively $\mathbb{P}(W_0 \geq y) = \alpha$); the deadline it corresponds to is $2 \log h - \log y$ (respectively $\log h + d_h - \log y$).

Three conclusions follow. First, in the reserve regime the limiting miss 3.96% is strictly below 5%: for all sufficiently large h a window exists, with margin. Second, with zero, one or two units of headroom the limiting miss exceeds 5% (barely so at $m = 2$), so no window exists for large h ; with three or more it does. Third, the limiting tradeoffs depend only on (a, λ, m) : two assay families with the same background-to-growth ratio and the same loading face the same limiting design constraints, whatever their thresholds. The margin at $m = 2$ is 0.03 percentage points and should not be promoted to a finite-threshold claim without a quantitative estimate of the kind given next; the reviewer-style condition “ $R/h = O(1)$ ” is not the right dichotomy, since $R/h \rightarrow \rho > 1$ and $R = h + m$ both satisfy it while giving opposite conclusions.

5.6 An explicit theorem for every threshold at least 10^8

The limits above are qualitative. The following statement is quantitative and covers every large threshold with explicit deadlines and constants.

Theorem 5.5 (Large thresholds; conventional). *Let $a = b/g = 1/100$, $\lambda = 4$, $\alpha = 1/100$, $\beta = 1/20$, and let $h \geq 10^8$ be an integer. Define*

$$\mu_+(h) = \sum_{z=0}^{h-1} \frac{z}{(z+a)(2h-z)}, \quad t_+(h) = \log(h+a) + \mu_+(h) - \log \frac{3}{10}, \quad (12)$$

$$\mu_-(h) = \sum_{z=0}^{k-1} \frac{1}{h-z} + \sum_{z=k}^{h-1} \frac{1}{z+a}, \quad t_-(h) = \frac{h}{h+a} \left[\log(k+a) + \log(\ell+1) + \mu_-(h) - \log \frac{3}{20} \right], \quad (13)$$

with $k = \lfloor h/2 \rfloor$ and $\ell = h - k$ (times in units of $1/g$). Then

- (a) at capacity $R = 2h$, $F_{0,2h}(t_+) < 0.00958$ and $M_{2h,4}(t_+) < 0.04894$, so $t_+(h) \in \mathcal{W}_{2h,4}$;
- (b) at capacity $R = h$, $F_{0,h}(t_-) > 0.01006$ and $M_{h,4}(t_-) > 0.0555$, so $t_-(h)$ is a separating time and $\mathcal{W}_{h,4} = \emptyset$.

The proof is in Appendix C. It combines Lemma 5.1, the exact couplings of Lemma 5.2 applied to the blocks of Theorem 5.4, Chebyshev bounds with $\varepsilon = 0.01$ on the gamma variables and on the concentrated remainders (each contributing $O(1/h)$ to a union bound, which is where $h \geq 10^8$ enters), and four scalar enclosures: an upper bound for $\mathbb{P}(\text{Gamma}(a) \geq 0.293)$ through the exponential integral, an upper bound for the loaded mixture through $e^{-4}I_0(2\sqrt{1.228})$, a lower bound for $\mathbb{P}(\text{Gamma}(a)Y_1 \geq 0.155)$ through a truncated integral of the form (11), and the elementary lower bound $\mathbb{P}(W_n Y_1 \leq y) \geq y/(n+a+y)$. The enclosures are evaluated with directed rounding in Appendix D. The constant 10^8 is a convenient sufficient threshold for these estimates, not a claim that the phenomenon begins there: Theorem 3.1 covers $h = 5$, and the numerical results below cover the range between.

5.7 Numerical corroboration at $h = 10^8$ and convergence

The Laplace transform of $T_{R,h,n}$ is the finite product $\varphi_n(s) = \prod_{z=n}^{h-1} q_R(z)/(q_R(z) + s)$, which costs $O(h)$ to evaluate; the distribution function $\mathbb{P}(T_{R,h,n} \leq t)$ is recovered by fixed-Talbot inversion [1, 2, 22]. At $h = 5$ and $h = 20$ the inversion reproduces the matrix-exponential values of Table 1 and of the earlier threshold-twenty check to 1.5×10^{-12} ; near and above the blank boundary, node counts 32 to 56 agree to about 10^{-7} . Table 3 gives the results at $h = 10^8$, and Figure 4 the convergence of the boundary miss to the limits of Table 2 over eight orders of magnitude in h . The script `large_threshold.py` reproduces both.

Capacity	Deadline	Blank	Bound	Miss	Bound
$R = 2h$	$t_+(10^8) = 20.3178$	0.009072	< 0.00958	0.042852	< 0.04894
$R = h$	$t_-(10^8) = 38.7385$	0.011712	> 0.01006	0.091891	> 0.0555
$R = 2h$	$t_b = 20.4417$	0.01	—	0.039625	limit 0.03962
$R = h$	$t_b = 38.4831$	0.01	—	0.109139	limit 0.10914

Table 3: The stress test at threshold $h = 10^8$ (numerical; times in units of $1/g$; $\mu_+ = 0.6931$, $\mu_- = 1.3863$). The first two rows evaluate the source at the explicit deadlines of Theorem 5.5 and compare with its rigorous bounds; the last two locate the blank boundary t_b and compare the miss there with the limit laws of Table 2. The initially positive fraction $\mathbb{P}(N \geq 10^8)$ is below 10^{-10^7} .

5.8 What the large-threshold behaviour means for diagnostic detection

- (1) *The effect is not a small-count artefact.* Its cause is the order-one holding time of the last few births before resource exhaustion, which is independent of h ; Theorem 5.5 exhibits

h	$R = h$	$R = h + 1$	$R = h + 2$	$R = h + 3$	$R = 2h$	$R = 10h$
5	8.957	6.145	5.259	4.832	4.416	3.701
10	11.607	7.250	6.004	5.436	4.420	3.869
20	12.062	7.319	6.037	5.486	4.273	3.932
100	11.437	6.602	5.440	4.996	4.054	3.963
10^3	11.002	6.188	5.094	4.698	3.975	3.963
10^4	10.926	6.119	5.037	4.648	3.964	3.963
10^6	10.914	6.108	5.028	4.640	3.963	3.962
10^8	10.914	—	—	—	3.962	—
limit	10.914	6.108	5.027	4.639	3.962	3.962

Table 4: Miss probability (%) at the blank boundary as the threshold grows (numerical, Laplace inversion), with the limits of Theorems 5.3 and 5.4. Rows $h = 5$ reproduce the windows of Section 3 ($R = 5, 6, 7, 8$ and 10); the headroom needed for a window is three at $h = 5$ and in the limit but four at $h = 10$ to 50 .




Figure 4: Left: miss probability at the blank boundary as a function of the threshold h for fixed headroom $R = h, \dots, h + 3$ and fixed ratio $R = 2h, 10h$ (numerical, Laplace inversion; squares at $h = 10^8$); dashed lines are the limits of Table 2. Right: the blank boundary itself grows like $\log h$ in the reserve regime and like $2\log h$ at saturation.




Figure 5: Limiting miss at the blank boundary as a function of the headroom m (numerical). With these parameters at most two remaining units leave no window in the limit; the reserve limit is the undepleted value.

it at $h = 10^8$, where the fraction of reactions positive at time zero is negligible and the Poisson loading is a vanishing fraction of the threshold.

- (2) *Absolute headroom, not the capacity ratio, decides.* A capacity that exceeds the detection threshold by a growing margin restores the undepleted design rules after a common delay d_h ; a capacity that exceeds it by a fixed number of units does not, and the limiting tradeoff is then a function of $(b/g, \lambda, m)$. For a real assay the question becomes measurable: how many effective units remain when the detector fires. Increasing h alone cannot answer it.
- (3) *Deadlines grow logarithmically, windows do not shrink.* In units of $1/g$ the blank boundary is $\log h + d_h - \log y$ in the reserve regime and $2 \log h - \log y_m$ at fixed headroom (Figure 4, right); at $h = 10^8$ and $g = 1$ per minute these are about 20 and 38 minutes. The width of a nonempty window converges to a positive limit, so the timing tolerance available to the instrument does not deteriorate with threshold, and an observation grid of a few seconds is not the obstruction at large h that it is at $h = 5$. In the fixed-rate-constant family of Proposition 4.2 the time unit is $V/(k_g R)$, so physical deadlines and widths scale like $(\log R)/R$ and $1/R$; whether a real assay follows that family is part of the open calibration question.
- (4) *The reserve limit is strictly feasible; saturation is not.* With $a = 0.01$ and $\lambda = 4$ the undepleted limit has miss 3.96% at the blank boundary, so any assay design that keeps a growing reserve will have a window at large thresholds. A design that consumes essentially all of its effective resource at detection will not, whatever the threshold, and no re-timing repairs it.
- (5) *The universal tradeoff is a benchmark.* Since the limit laws are explicit gamma and product-gamma laws, inference or optimisation software for stochastic assay design can be tested against them without any molecular calibration.

6 Data boundary, a discriminating experiment and sampling confidence

Public example. The saved digital-NAAT example distributed with the Ismagilov laboratory’s melt-curve analyser [12] contains 258 wells and 185 amplification frames. Reapplying

its strict crossing rule (baseline-corrected intensity greater than 250) reproduces all 29 saved crossings with no mismatch; the other 229 wells are retained as right-censored non-crossers, giving an all-well detection fraction $29/258 \approx 0.1124$ rather than the positive-only distribution that ends at one. This is a replay of the observation rule, not a fit or validation of the source: the example has no verified matched blank/loading labels, resource levels or independent run identifiers, and the unobserved future of the non-crossers is not identified by the recorded frames. Three candidate dilution and formulation workbooks in a related deposit [14] returned HTTP 403 during this investigation; their controls remain unassessed, which does not establish that they are unsuitable. Melt information is available only after amplification and cannot be used by a rule that must decide earlier.

A discriminating experiment. The prediction that separates the finite-resource source from a common clock change is a change in the sensitivity–blank tradeoff under a controlled resource intervention, not merely a shift of all crossing times. A study should retain matched blanks and known loading, all trajectories including non-crossers, and independent run labels across at least two measured resource conditions; fit a declared subset, freeze the source and decision rule, and test a withheld condition against an empirical crossing-time model with common acceleration. Section 5.8 adds a second, sharper prediction: the effect should depend on the headroom at detection and disappear when the reserve is large. Outcomes that would change the model include a pure common-clock shift, unequal target and background growth, delayed activation or extinction in the low-loading tail, and readout noise or run heterogeneity larger than the tolerance of Theorem 3.2.

Sampling confidence is separate from operating probability. For a prespecified deadline and independent identically distributed reactions, zero errors in n trials give the exact one-sided upper confidence bound $1 - \delta^{1/n}$ at level $1 - \delta$ [3, 8]. Allocating $\delta = 0.025$ to each error limit, $0.99^{368} \leq 0.025$ and $0.95^{72} \leq 0.025$, so zero errors in 368 blanks and 72 positive reactions give simultaneous one-sided confidence bounds at total level 95% by the union bound. This planning calculation covers neither unmodelled batch variation nor selection of the deadline from the same observations, and the 1% and 5% model limits and the 95% statistical confidence must not share an unlabelled axis. An OR rule across m independent blank reactions has false-positive probability $1 - (1 - \alpha)^m$, not α ; the paper makes no specimen-level claim.

7 Verification and reproducibility

Lean-verified. The following are theorems of the frozen Lean 4 project (toolchain `leanprover/lean4:v4.30.0`, pinned Mathlib manifest with SHA-256 prefix `a8df0b60`), compiled with warnings treated as errors and with their declaration interfaces authenticated; each depends only on the axioms `propext`, `Classical.choice` and `Quot.sound`. No admitted proof, extra probability axiom or native computation is involved. The reusable probability layer is `FiniteCopy/FiniteKernel` and `FiniteCopy/PoissonKernel`; the development is in `proofs/DiagnosticWindows`.

- The normalized six-state source (3), its nonnegativity and row sums, the latch indicator and the fact that supported moves increase the count by at most one (`birthKernel`, `birth_aggregates`); the literal rate formula (2) at each capacity (`capacityR_rates`).
- The infinite loading mixture (6) and its exact finite reduction (`loadedSurvival`, `loading_finite`); the empty-loading floor (7) (`empty_loading_floor`).
- The spectral identities and exact five-exponential survival formulas for capacities five to ten (`spectral_survival`, `capacityR_survival`); time monotonicity from dissipation

(spectral_antitone).

- The exponential enclosure scheme (`exp_enclosure16`) and its 11+51 rational certificates; the certified inequalities of Table 1 and of Theorem 3.2.
- Theorem 3.1: `minimum_capacity_certificates` (parts (a), (b)) and `capacity10_interval` (part (c)); the original two-capacity contrast `main`.
- Rate order, loading order and the scaling identity (`birth_survival_order`, `loading_antitone`, `scaled_survival`), and the existence equivalence of Proposition 4.1 (`clock_feasibility_iff`).
- Theorem 3.2: `robust_guarantee`.

The strict receipts for `Windows.lean` (`minimum_capacity_certificates`; source SHA-256 prefix 08b9fdef, 14 dependency modules, 31 s), `Robustness.lean` (`robust_guarantee`; 67d6b291, 18 modules, 79 s) and `Main.lean` (`main`; 493d2290, 103 s) accompany the sources. Appendix E maps each claim to its declaration.

Conventional. Lemmas 2.1 and 2.2 as stated for general R (their instances are Lean-verified); Proposition 4.2; all of Section 5 and Appendices C–D. A Lean development of Section 5 would have to formalize the identification of the uniformized law with independent holding times, the beta law and beta–gamma coupling, the variance bounds, and the four enclosures; the last are within reach of the existing enclosure scheme, the first three are not yet formalized.

Numerical. The decimal displays in every table; the window endpoints and $\lambda_{\min}$ values; the undepleted reference at 2.8 minutes; the mass-action grid witness; Table 2, Table 3 and Figures 4–5; the earlier threshold-twenty check (capacity 20 at 8.0 minutes: blank .01341831, miss .08764661; capacity 40 at 4.8: .00908692, .04637034), now superseded as a stress test by Theorem 5.5 but retained as a data point in Figure 4.

Reproducibility. The paper source directory contains `check_paper.py`, which recomputes the rational eigenvector decompositions, evaluates every witness of Table 1 in 40-digit arithmetic, replays the directed-rounding enclosures of Appendix D and reports 41 checks; `large_threshold.py`, which performs the Laplace inversion and the limit-law quadratures; and `make_figures.py`. The problem workspace of the repository retains the certificate generators, the strict receipts with dependency manifests, the public-example replay with download hashes, and the original figure scripts. The repository location will be supplied with the author information.

8 Conclusion and open problems

For an explicit finite-resource amplification source, exact certificates settle whether a usable readout deadline exists, what a prescribed observation grid does to that answer, and how much fixed-rate, loading and timing uncertainty a feasible design tolerates; these are Lean-verified at threshold five. A common clock change cannot create a window; a state-dependent resource change can. At large thresholds the holding-time representation reduces the question to explicit gamma and product-gamma laws: a growing reserve restores the undepleted rules after a deterministic delay, a fixed headroom does not, and for every $h \geq 10^8$ the contrast between saturation and a doubled capacity is proved with explicit constants.

Open problems. (1) A measured interpretation of effective capacity, background and growth in a real assay, and a test of the headroom prediction under a controlled resource intervention. (2) A quantitative finite-threshold classification of the minimal headroom, which the numerical

results suggest is three at $h = 5$ and in the limit but four at intermediate thresholds. (3) Formalization of the holding-time and beta-gamma arguments, which would bring Section 5 to the same evidence level as Section 3. (4) Extensions with activation delay, loss, unequal target and background growth, readout noise, time-dependent rates or adaptive stopping, each of which changes the source and needs its own analysis. (5) Specimen-level rules that aggregate several reactions.

A Exact source calculation and rational enclosures

Write $r_z = q_R(z)/\nu$ and $P = P_R$. For each of the six capacities the five transient rates are distinct, and the uncalled indicator is a sum of five eigenvectors,

$$f = \sum_{i=0}^4 V_i, \quad PV_i = e_i V_i, \quad e_i = 1 - r_i.$$

To construct V_i , set its entries above i (including at the absorbing state) to zero and use the backward recurrence

$$V_i(z) = \frac{q_R(z)}{q_R(z) - q_R(i)} V_i(z+1) \quad (z < i), \quad (14)$$

choosing $V_i(i)$ from $i = 4$ down to 0 so that $\sum_j V_j(i) = 1$. Every coefficient is rational; Lean checks the decomposition and each eigenvector equation against the literal kernel by direct substitution. Induction gives $P^n V_i = e_i^n V_i$, and summing the Poisson series in (4) yields

$$S_R(t, z) = \sum_{i=0}^4 e^{-q_R(i)t} V_i(z). \quad (15)$$

The components may have signed entries; positivity is a property of the kernel, not of this representation. Near-coincident rates produce large cancelling coefficients: at capacity five, where $q_5(2) = 1.206$ and $q_5(3) = 1.204$, the loaded coefficients exceed 4900 in absolute value (over 11,000 at capacity seven). Floating-point agreement was therefore not accepted as a certificate.

Lemma A.1 (Time monotonicity; Lean-verified). *$S_R(t, z)$ is nonincreasing in $t \geq 0$ for every z .*

Proof. $Pf \leq f$ because f is nonincreasing on the states and P moves upward. In the uniformization clock $u = \nu t$ put $d = Pf - f = \sum_i (e_i - 1)V_i \leq 0$. Differentiating (15),

$$\frac{d}{du} \sum_i e^{u(e_i-1)} V_i(z) = \sum_i e^{u(e_i-1)} (e_i - 1) V_i(z) = \sum_{n \geq 0} e^{-u} \frac{u^n}{n!} (P^n d)(z) \leq 0,$$

the last inequality by kernel positivity. Nonnegative Poisson loading weights preserve the order. $\square$

Rational enclosures. At a rational deadline every exponent in (15) is rational. For each required exponent x let $y = x/16$, so $|y| \leq 1$, and put $A(y) = \sum_{j=0}^{11} y^j/j!$ and $B(y) = 13|y|^{12}/(12 \cdot 12!)$. The Taylor remainder bound available in Mathlib gives $A(y) - B(y) \leq e^y \leq A(y) + B(y)$. Rational l, u are chosen outward at precision 10^{-12} with $0 \leq l \leq A(y) - B(y)$ and $A(y) + B(y) \leq u$; then $l^{16} \leq e^x \leq u^{16}$, and a final outward rounding at 10^{-9} gives compact endpoints. Each step, including the sign-sensitive propagation of the enclosures through (15) and (6), is checked in Lean. The robustness endpoints use loading 4 for the blank bound and the exact mean 399/100 for the miss bound; no displayed decimal is a premise.

B Comparison lemmas for the source law

Let P_r be a pure-birth kernel on $\{0, \dots, 5\}$ with jump probabilities $0 \leq r_z \leq 1$ and absorbing state five.

Lemma B.1 (Rate order; Lean-verified). *If a is nonincreasing on the states then so is $P_r a$. If $r_z \leq s_z$ for all z then $P_s^n f \leq P_r^n f$ for all n , hence $S_s(t, z) \leq S_r(t, z)$ and, after loading, $M_{s,\lambda} \leq M_{r,\lambda}$. Also $S_r(t, z)$ is nonincreasing in z .*

Proof. $(P_r a)(z) = (1 - r_z)a(z) + r_z a(z + 1)$ lies between $a(z + 1)$ and $a(z)$; for $x < y$, $(P_r a)(x) \geq a(x + 1) \geq a(y) \geq (P_r a)(y)$. No increasing-rate assumption is needed. If $r_z \leq s_z$ then $(P_s a)(z) - (P_r a)(z) = (s_z - r_z)(a(z + 1) - a(z)) \leq 0$; positivity and induction give the order for powers, and nonnegative Poisson weights give it for (4) and (6). $\square$

Lemma B.2 (Loading order; Lean-verified). *For fixed t , $\lambda \mapsto M_{R,\lambda}(t)$ is nonincreasing on $[0, \infty)$.*

Proof. With $a_n = S(t, n)$, $a_5 = 0$, nonnegative and nonincreasing,

$$\frac{d}{d\lambda} \left[e^{-\lambda} \sum_{n=0}^4 \frac{\lambda^n}{n!} a_n \right] = e^{-\lambda} \left[\sum_{n=0}^3 \frac{\lambda^n}{n!} (a_{n+1} - a_n) - \frac{\lambda^4}{24} a_4 \right] \leq 0. \quad \square$$

Lemma B.3 (Rate scaling; Lean-verified). *If $cr_z \leq 1$ for all z then $P_{cr} = cP_r + (1 - c)I$, its action on an eigenvector of eigenvalue e_i is multiplication by $1 + c(e_i - 1)$, and $S_{cr}(u, z) = S_r(cu, z)$; the same identity holds after loading. For arbitrary $c > 0$, uniformizing the rates cq_z at cu leaves the normalized kernel unchanged and scales its Poisson clock, which gives Proposition 4.1.*

C Proof of Theorem 5.5

Fix $a = 1/100$, $\varepsilon = 1/100$, $n_* = 25$ and $C = 10,000 + (20,000/199)^2$. All times are in units of $1/g$.

Concentration inequalities. For $Z \sim \text{Gamma}(s, 1)$, the event $|\log(Z/s)| > \varepsilon$ implies $|Z - s| > s(1 - e^{-\varepsilon})$, so by Chebyshev

$$\mathbb{P}(|\log(Z/s)| > \varepsilon) \leq \frac{1}{s(1 - e^{-\varepsilon})^2} \leq \frac{(20,000/199)^2}{s}, \quad (16)$$

using $1 - e^{-\varepsilon} \geq \varepsilon - \varepsilon^2/2 = 199/20,000$. For $h \geq 10^8$ put

$$\delta_+ = C/h < 0.000202, \quad \delta_- = (401/100)C/h < 0.000807, \quad (17)$$

and note $\mathbb{P}(N > 25) \leq e^{-4 \frac{4^{26}}{26!}} (1 - \frac{4}{27})^{-1} < 3 \times 10^{-13}$ and $25/(h - 24) < 3 \times 10^{-7}$.

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 $\mathbb{E} \log X = -\gamma \geq -1$ gives $\mathbb{E}X^a \geq 1 - a$. (ii) From $w^a \leq 1 - a + aw$,

$$\mathbb{P}(\text{Gamma}(a, 1) \geq y) \leq 0.01 E_1(y) + \frac{0.01}{99} e^{-y}, \quad E_1(y) = \int_y^\infty \frac{e^{-w}}{w} dw. \quad (18)$$

(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,

$$\sum_{n \geq 0} \mathbb{P}(N = n) \mathbb{P}(\text{Gamma}(n+a, 1) \leq y) \leq e^{-4} \sum_{n \geq 0} \frac{(4y)^n}{(n!)^2} = e^{-4} I_0(4\sqrt{y}). \quad (19)$$

(iv) For $Y \sim \text{Exp}(1)$ independent of $W \sim \text{Gamma}(a, 1)$, since $0.97^{100} < 0.05$ gives $w^a \geq 0.97$ on $w \geq 0.05$,

$$\mathbb{P}(WY \geq y) = \frac{1}{\Gamma(a)} \int_0^\infty w^{a-1} e^{-w-y/w} dw \geq 0.0097 \int_{0.05}^{10} \frac{e^{-w-y/w}}{w} dw. \quad (20)$$

(v) For $W_n \sim \text{Gamma}(n+a, 1)$, by $1 - e^{-u} \geq u/(1+u)$ and the convexity of $w \mapsto y/(w+y)$,

$$\mathbb{P}(W_n Y \leq y) = \mathbb{E}[1 - e^{-y/W_n}] \geq \mathbb{E}\left[\frac{y}{W_n + y}\right] \geq \frac{y}{n+a+y}. \quad (21)$$

(a) Capacity $2h$. By Lemma 5.1, $T_{2h,h,n} = A_{h,n} + D_n$ with $D_n = \sum_{z=n}^{h-1} \frac{zE_z}{(z+a)(2h-z)}$, $\text{Var } D_n \leq \sum_{z < h} (2h-z)^{-2} \leq 1/h$, and $|\mathbb{E}D_n - \mu_+(h)| \leq 25/(h-24)$ for $n \leq 25$. By Lemma 5.2, $A_{h,n} = \log Z - \log W_n$ with $Z \sim \text{Gamma}(h+a, 1)$, $W_n \sim \text{Gamma}(n+a, 1)$. Apply Chebyshev to D_n (probability at most $\text{Var } D_n/\varepsilon^2 \leq 10^4/h$) and (16) to Z ; outside an event of probability at most δ_+ ,

$$|T_{2h,h,n} - \log(h+a) - \mu_+(h) + \log W_n| \leq 2\varepsilon + \frac{25}{h-24} < 0.020001. \quad (22)$$

At t_+ of (12), on the good event, $T_{2h,h,0} \leq t_+$ forces $W_0 \geq 0.3e^{-0.020001} > 0.293$, and $T_{2h,h,n} > t_+$ forces $W_n < 0.3e^{0.020001} < 0.307$. Hence

$$\begin{aligned} F_{0,2h}(t_+) &\leq \mathbb{P}(\text{Gamma}(a, 1) \geq 0.293) + \delta_+, \\ M_{2h,4}(t_+) &\leq \sum_{n \geq 0} \mathbb{P}(N = n) \mathbb{P}(\text{Gamma}(n+a, 1) \leq 0.307) + \delta_+ + 3 \times 10^{-13}, \end{aligned}$$

where loads $n > 25$ are covered by the Poisson tail and loads $n \geq h$ are absorbed. By (18), (19) and the enclosures of Appendix D,

$$0.01E_1(0.293) + \frac{0.01}{99}e^{-0.293} < 0.009372, \quad e^{-4} \sum_{n \geq 0} \frac{1.228^n}{(n!)^2} < 0.048731,$$

so $F_{0,2h}(t_+) < 0.009372 + 0.000202 < 0.00958$ and $M_{2h,4}(t_+) < 0.048731 + 0.000202 + 3 \times 10^{-13} < 0.04894$.

(b) Capacity h . With $c_h = h/(h+a)$, $k = \lfloor h/2 \rfloor$, $\ell = h - k$, Lemma 5.1 gives $T_{h,h,n}/c_h = U_n + V + C_n$ as in the proof of Theorem 5.4, with $e^{-U_n} \sim \text{Beta}(n+a, k-n)$ and $e^{-V} \sim \text{Beta}(1, \ell)$ independent. Apply the coupling of Lemma 5.2 to each block: $U_n = \log Z_1 - \log W_n$ and $V = \log Z_2 - \log Y$ 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. Every coefficient in C_n is at most $1/k$ and there are at most h of them, so $\text{Var } C_n \leq h/k^2 \leq 4h/(h-1)^2 \leq 4.01/h$; likewise $1/(k+a) + 1/(\ell+1) \leq 2/k \leq 4.01/h$; and $|\mathbb{E}C_n - \mu_-(h)| \leq 25/(h-24)$ for $n \leq 25$. Chebyshev on C_n and (16) on Z_1, Z_2 give, outside an event of probability at most δ_- ,

$$|T_{h,h,n}/c_h - \log(k+a) - \log(\ell+1) - \mu_-(h) + \log(W_n Y)| < 0.030001. \quad (23)$$

At t_- of (13), on the good event, $W_0Y \geq 0.155 > 0.15e^{0.030001}$ forces $T_{h,h,0} \leq t_-$, and $W_nY \leq 0.145 < 0.15e^{-0.030001}$ forces $T_{h,h,n} > t_-$. Hence

$$F_{0,h}(t_-) \geq \mathbb{P}(\text{Gamma}(a, 1)Y \geq 0.155) - \delta_-,$$

$$M_{h,4}(t_-) \geq e^{-4} \sum_{n=0}^5 \frac{4^n}{n!} \frac{0.145}{n + 0.155} - \delta_-,$$

using (21) and keeping only the nonnegative loads $n \leq 5$. By (20) and Appendix D, $\mathbb{P}(\text{Gamma}(a, 1)Y \geq 0.155) > 0.01087$ and the finite sum exceeds 0.056383, so $F_{0,h}(t_-) > 0.01087 - 0.000807 > 0.01006$ and $M_{h,4}(t_-) > 0.056383 - 0.000807 > 0.0555$. Both limits are strictly violated at t_- , and Lemma 2.1 excludes every deadline. $\square$

D Directed arithmetic for the four enclosures

The four scalar bounds used in Appendix C are evaluated in `check_paper.py` with Python's `decimal` module at 50 digits and explicit rounding directions, so that each printed bound is an outward bound of the corresponding exact quantity. For $0 \leq u \leq 1$ the Taylor polynomial of degree eleven lies below e^u and adding $13u^{12}/(12 \cdot 12!)$ gives an upper bound; scaling by 32 and squaring five times encloses e^x for $0 \leq x \leq 32$, and reciprocation encloses e^{-x} . The exponential integral is bounded above by an upper Riemann sum on $[0.293, 20]$ with step 0.005 plus the tail $e^{-20}/20$; the series $\sum 1.228^n/(n!)^2$ is summed to $n = 20$ with a geometric tail; the Poisson tail uses its first omitted term and a geometric ratio $4/27$; the integral in (20) is bounded below by a lower Riemann sum on $[0.05, 10]$ with step 0.005, using the convexity of $w + 0.155/w$ on each subinterval. The results are

Quantity	Directed bound
capacity $2h$, blank upper bound, after adding δ_+	0.0095725
capacity $2h$, miss upper bound, after adding δ_+ and the Poisson tail	0.0489313
capacity h , blank lower bound, after subtracting δ_-	0.0100703
capacity h , miss lower bound, after subtracting δ_-	0.0555773

and the script asserts the coarser inequalities < 0.00958 , < 0.04894 , > 0.01006 , > 0.0555 used in the theorem. An independent evaluation with `mpmath` at 40 digits agrees with each bound to the printed digits; the exact limiting quantities are more favourable still ($\mathbb{P}(\text{Gamma}(a) \geq 0.293) = 0.009247$, $\mathbb{P}(\text{Gamma}(a)Y \geq 0.155) = 0.011485$). The arithmetic is h -independent except through the allowances $\delta_{\pm}$, which decrease with h .

E Claims-to-declarations map

Claim	Lean declaration(s) in namespace <code>DiagnosticWindows</code>
Source kernel, positivity, row sums, latch	<code>birthKernel</code> , <code>uncalled_indicator</code> , <code>birth_aggregates</code> (<code>BirthSource</code>)
Loading mixture, exact finite reduction, floor (7)	<code>loadedSurvival</code> , <code>loading_finite</code> , <code>empty_loading_floor</code> (<code>Loading</code>)
Spectral survival, monotonicity from dissipation	<code>spectral_survival</code> (<code>Spectral</code>), <code>spectral_antitone</code> (<code>Monotonicity</code>)
Rational eigenvectors, rates, survival formulas $R = 5, \dots, 10$	<code>capacityR_rates</code> , <code>capacityR_sum</code> , <code>capacityR_eigen</code> , <code>capacityR_survival</code> (<code>Capacity</code> , <code>ExtraCapacity</code>)
Exponential enclosures	<code>exp_enclosure16</code> (<code>ExponentialBounds</code>); 11 certificates (<code>ExpCertificates</code>); 51 certificates <code>post_exp_i</code> (<code>WindowCertificates</code>)
Table 1 inequalities	<code>capacity5_separating_errors</code> , <code>capacity10_errors</code> (<code>Numerical</code>); <code>rR_blank_bound</code> , <code>rR_miss_bound</code> , <code>r8early_miss_bound</code> , <code>r8late_blank_bound</code> , <code>interval_early_miss_bound</code> , <code>interval_late_blank_bound</code> (<code>WindowCertificates</code>)
Theorem 3.1(a),(b)	<code>capacityR_no_deadline</code> , <code>capacity8_grid_impossible</code> , <code>minimum_capacity_certificates</code> (<code>Windows</code>)
Theorem 3.1(c)	<code>capacity10_interval</code> (<code>Windows</code>)
Original two-capacity contrast	<code>main</code> (<code>Main</code>)
Lemma B.1	<code>birth_step_antitone</code> , <code>birth_survival_order</code> , <code>birth_survival_antitone_state</code> (<code>Comparison</code>)
Lemma B.2	<code>loading_antitone</code> , <code>loaded_rate_order</code> (<code>LoadingOrder</code>)
Lemma B.3, Proposition 4.1	<code>scaled_survival</code> , <code>scaled10_loading</code> , <code>clock_feasibility_iff</code> (<code>RateScale</code>)
Theorem 3.2	<code>robust_blank_bound</code> , <code>robust_miss_bound</code> (<code>WindowCertificates</code>); <code>robust_guarantee</code> (<code>Robustness</code>)

References

- [1] Joseph Abate and Peter P. Valkó. Multi-precision Laplace transform inversion. *International Journal for Numerical Methods in Engineering*, 60(5):979–993, 2004. doi: 10.1002/nme.995.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] Paul Erdős and Alfréd Rényi. On a classical problem of probability theory. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, 6:215–220, 1961.
- [6] William Feller. *An Introduction to Probability Theory and Its Applications*, volume I. Wiley, New York, third edition, 1968.
- [7] Winfried K. Grassmann. Transient solutions in Markovian queueing systems. *Computers & Operations Research*, 4(1):47–53, 1977. doi: 10.1016/0305-0548(77)90007-7.
- [8] 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.

- [9] 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.
- [10] 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.
- [11] 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.
- [12] Ismagilov Laboratory. Digital NAAT two-channel melt-curve analyzer: MATLAB source and saved example. https://github.com/IsmagilovLab/Digital_NAAT_2Ch_MeltCurve_Analyzer, 2020. Accessed September 2026.
- [13] 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.
- [14] Erik Jue. Dataset: Two-phase wash to solve the ubiquitous contaminant-carryover problem in commercial nucleic-acid extraction kits. CaltechDATA, <https://data.caltech.edu/records/8d1ar-e8e29>, 2019.
- [15] David G. Kendall. On the generalized “birth-and-death” process. *Annals of Mathematical Statistics*, 19(1):1–15, 1948. doi: 10.1214/aoms/1177730285.
- [16] Eugene Lukacs. A characterization of the gamma distribution. *Annals of Mathematical Statistics*, 26(2):319–324, 1955. doi: 10.1214/aoms/1177728549.
- [17] James R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, 1997. doi: 10.1017/CBO9780511810633.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.