# Reserve capacity and reliable eradication in inherited-state populations

21 September 2026

**Abstract**

We study finite-time eradication of a branching population with inherited cellular states under a shared intervention that also damages a finite regenerative reserve. The mission requires separate bounds on target survival and on reserve loss at *every* time during treatment, not merely at the deadline. Our main probabilistic tool is an excursion bound anchored at a freely chosen return level below capacity. It yields an explicit finite-product reserve certificate, it accommodates incomplete initial filling, and we show it is asymptotically exact at high renewal. Composing it with source-specific weighted drift inequalities for a six-state inheritance model and a bounded administration course with uncertain clearance, we prove that a single course, uniform over independent one-percent relative errors in the target rates, gives target-survival probability below 0.00856 and reserve-loss probability below $7 \times 10^{-6}$ from a 400-unit reserve that need only be filled to 278. The same course still meets both one-percent mission tolerances from any initial filling of 209 units, nine above the failure threshold, and a 232-unit reserve suffices at the same proportions. A capacity law with an explicit integer-rounding correction quantifies how additional reserve replaces fast renewal, while a complementary fluid-limit statement shows that capacity cannot compensate for an inadequate equilibrium reserve. We also prove the exact high-renewal limit $r_{\text{eff}}^d\,\mathbb{P}_K(\tau_H \leq T) \to (Km)^{d+1}T/d!$ for buffer $d$, and we state the limitations imposed by baseline loss and delayed replenishment. The probability arguments are conventional proofs; the decisive finite source and numerical inequalities are checked in Lean. The results are sufficient guarantees for an explicit synthetic source, not an optimal control frontier or a clinically calibrated regimen.

## 1 Introduction

An intervention can eliminate a target population and still fail its intended mission by depleting a healthy population before elimination is complete. This distinction matters especially for finite populations with inherited protective states: first moments need not determine extinction probabilities, and an adequate final healthy count does not imply that a functional reserve was maintained throughout treatment. We therefore ask for a single delivered intervention satisfying two probability constraints,

$$\mathbb{P}(|Z_T| > 0) \leq \delta, \qquad \mathbb{P}(\tau_H \leq T) \leq \alpha, \qquad \tau_H = \inf\{t \geq 0 : H_t < h_{\min}\}. \tag{1}$$

Here $Z$ records target counts by inherited state and $H$ is a finite functional reserve. Their outcomes may be dependent. The two inequalities imply joint success at least $1 - \delta - \alpha$ by the union bound,

$$\mathbb{P}(Z_T = 0,\ \tau_H > T) \geq 1 - \mathbb{P}(|Z_T| > 0) - \mathbb{P}(\tau_H \leq T),$$

so no independence assumption is needed in the final assembly. The second constraint is a genuine path requirement: recovery after a crossing does not erase the failure, and a small terminal shortfall probability does not imply it.

**Contributions.** The main probabilistic tool is elementary but useful: a healthy excursion need not start at the hard ceiling. Choosing its return level near the operating reserve substantially improves a finite-time bound and permits incomplete initial filling. Around this observation the paper establishes five things.

1. An anchored return-level certificate for reserve loss (Theorem 5) and its incomplete-preparation form (Corollary 6), valid for all adapted rates dominating a logistic comparison chain.
2. A capacity-sizing law with an explicit integer-rounding correction (Theorem 8), together with a fluid-limit statement (Proposition 9) showing that its regime hypothesis cannot be dropped: when the equilibrium reserve fraction falls below the threshold fraction, failure probability tends to one no matter how large the capacity.
3. An exact high-renewal asymptotic (Proposition 10): with buffer $d = K - h_{\min} \geq 1$ and a full initial reserve, $r_{\text{eff}}^{d}\, \mathbb{P}_K(\tau_H \leq T) \to (Km)^{d+1}T/d!$. The top-anchor product of Theorem 5 equals that constant divided by $r_{\text{eff}}^{d}$ *exactly*, for every $r_{\text{eff}}$, so the certificate is asymptotically sharp and the renewal-sizing rule (41) is order-optimal.
4. A delivered finite-mission theorem (Theorem 4) for a literal six-state inheritance source under a one-course administration with uncertain clearance, uniform over an explicit one-percent relative uncertainty box on the target rates.
5. Two consequences that quantify how little reserve the same course actually needs: the mission tolerances survive an initial filling of 209 out of 400 units (Corollary 13), and a 232-unit reserve at the same proportions suffices (Corollary 14).

**Relation to existing work.** Treatment optimization under phenotypic plasticity is established work. Gunnarsson and collaborators analyze switching, resistance and dosing, including asymptotically optimal dosing for specified plasticity laws [6, 7]. Healthy/cancer control with healthy-population constraints is also established [10]. Reversible, chromatin-associated drug-tolerant states and rare-cell reprogramming motivate treating protective cell states as inherited rather than as instantaneous phenotypes [11, 12]. Our question differs from all of these in its finite-population extinction event, its all-time stochastic reserve constraint, its inherited allocation source, and its one-course guarantee over an explicit uncertainty box. We do not claim to introduce toxicity-constrained control or to determine globally optimal treatment.

The probabilistic ingredients are classical. Scale functions and two-boundary ruin identities for birth–death chains are standard [8]; logistic-process extinction and its asymptotic regimes have a substantial literature [4]; density-dependent fluid limits are due to Kurtz [3, 9]; branching first-moment operators are textbook material [2]. The contribution here is not a new general claim about metastability but the explicit choice of a *return level*, the resulting finite-threshold certificate, and its composition with a delivered inherited-population course into a single uniform guarantee. Formulating the mission as a pair of marginal probability constraints places it near reach-avoid verification for controlled stochastic systems [1, 13], though we prove a sufficient certificate for one fixed input rather than synthesize an optimal policy.

**Benchmark lineage.** A 20-unit reserve provides a compact benchmark that we carry through the paper for comparison; it requires effective renewal at least 12 under a narrow uncertainty box. The larger-reserve constructions below demonstrate how capacity, initial filling and selectivity alter the sufficient renewal requirement, reducing it to 1. Table 1 collects these preparations. They are different physical preparations, not the same reserve relabelled or rescaled in time.

We first state the source and the finite-mission theorem, then prove the reserve bound, its capacity and asymptotic consequences, and the source-to-event implication. Design consequences and limitations follow. Only ingredients used for these conclusions are included; exploratory control searches are not premises of the paper.

## 2 Source, shared input and preparation

### 2.1 Inherited target population

The target has six types, ordered as

$$(UU, UR, RR, AU, AR, AA) \longleftrightarrow ((0,0),(0,1),(0,2),(1,0),(1,1),(2,0)).$$

The pair $(a, s)$ counts two kinds of marks on two sites; there are $u = 2 - a - s$ unmarked sites. With erasure parameter $e$, the chemical transitions are

$$\begin{aligned} &(a,s) \to (a+1,s):\ u(1/100 + a/2), &&(a,s) \to (a,s+1):\ u(1/100 + s/2),\\ &(a,s) \to (a-1,s):\ a(e + s/4), &&(a,s) \to (a,s-1):\ s(e + a/4). \end{aligned} \tag{2}$$

Zero-rate transitions outside the state space are omitted. The resulting row generator is $Q_e$, with diagonal entries chosen to make row sums zero. The nominal division rate is $b = 1/10$ and the nominal mortality vector is

$$d = (3/10, 3/10, 3/10, 1/100, 3/10, 1/100). \tag{3}$$

These are stipulated baseline mortalities in a synthetic model. They are not an unpriced clinical anticancer treatment. Healthy baseline mortality is likewise explicit below.

A dividing mother is removed. Her ordered daughters receive $(i, j)$ and $(a - i, s - j)$ old marks with probability

$$2^{-(a+s)}\binom{a}{i}\binom{s}{j}, \qquad 0 \le i \le a, \quad 0 \le j \le s, \tag{4}$$

and vacant sites are refilled with unmarked material. This is the complete allocation rule; there is no additional dilution event.

*Remark* 1 (Nature of the allocation law). Equation (4) is an *independent Bernoulli allocation of each individual mark to the first daughter, with complementary allocation to the second*. The first daughter's two mark counts are therefore independent binomials, $i \sim \mathrm{Bin}(a, 1/2)$ and $j \sim \mathrm{Bin}(s, 1/2)$. It is not a hypergeometric split: no sampling without replacement from a pooled set occurs. Complementarity is what makes the two daughters dependent, and that dependence is invisible to the first-moment operator below.

Its generating vector is

$$F(x) = \left(x_0^2,\ x_0x_1,\ \frac{x_0x_2 + x_1^2}{2},\ x_0x_3,\ \frac{x_0x_4 + x_1x_3}{2},\ \frac{x_0x_5 + x_3^2}{2}\right)^\top, \tag{5}$$

and the nonlinear extinction field is $\Phi(x) = Q_e x + d \odot (\mathbf{1} - x) + b(F(x) - x)$, with $D\Phi(\mathbf{1}) = A_e$ for the operator defined next. Write $L_{ij}$ for the expected number of type-$j$ daughters produced by a type-$i$ mother. For a column vector of weights $w$, the generator of $z \cdot w$ is $z \cdot A_e w$, where

$$A_e = Q_e + b(L - I) - \operatorname{diag}(d). \tag{6}$$

The $-I$ term removes the mother, and both daughters contribute to $L$. This first-moment operator uses the two-daughter expectation, not an independent-daughter replacement. A replacement kernel can preserve $L$, hence every mean equation, while changing extinction probabilities through $F$; a first-moment upper bound is insensitive to exactly that dependence.

*Remark* 2 (Memory and the Markov property). Cellular memory here is represented by persistent inherited mark configurations, which are part of the state. The expanded process $(Z, H, c)$ is Markov once its sufficient state and the delivered input are specified; the model is not intrinsically non-Markovian. What memory costs is dimension, not the Markov property. General history-dependent control policies need not leave the unaugmented count process Markov, but that is a separate control issue and is not used below.

### 2.2 Healthy reserve and delivery

The healthy source is a count process on $\{0, \dots, K\}$ with unit jumps, absorbing zero and hard capacity $K$. Its comparison rates are

$$\lambda_h = r_{\text{eff}} h(1 - h/K), \qquad \mu_h = mh. \tag{7}$$

For the actual source we allow predictable, locally bounded birth rates at least the displayed $\lambda_h$ and death rates at most the displayed $\mu_h$, subject to the same state space and unit-jump structure. These inequality directions are what an upper bound on failure requires, and they include renewal inhibition: if a nominal renewal rate $r$ retains at least a fraction $\chi_{\min} > 0$ under exposure, take $r_{\text{eff}} = r\chi_{\min}$. A count unit is an independently specified functional unit of this jump model. Enlarging $K$ changes the preparation or the physical scale; subdividing one correlated unit into several labels does not justify changing the noise law, and cannot legitimately improve a capacity certificate.

Administration $a(t)$ produces an effect-normalized concentration $c(t)$ by

$$\dot{c} = a - k_e c, \qquad c(0) = 0, \qquad e(t) = 1/100 + c(t). \tag{8}$$

The same concentration injures healthy units at per-unit rate

$$m_H(t) = \frac{3/10 + c(t)}{\sigma} + \nu. \tag{9}$$

Thus target and healthy effects cannot be prescribed independently: one input drives both. The parameter $\sigma$ describes differential response and $\nu$ is extra healthy mortality. The linear effect law and instantaneous logistic regeneration are stipulated modelling assumptions, not measured concentration–response curves; Section 8 records what each costs.

The mission starts with four $AA$ founders, has deadline $T = 120$, and requires $\delta = \alpha = 1/100$. The administered course is

$$a(t) = \begin{cases} 29/100, & 0 \le t \le 112, \\ 0, & t > 112. \end{cases} \tag{10}$$

The limits are $a \le 3/10$, $\int_0^{120} a \le 33$, $c \le 3/10$ and $\int_0^{120} c \le 33$. The final eight time units are a real washout tail and are retained in every healthy estimate. No target immigration occurs, so extinction is absorbing, including after the deadline. Indefinite healthy safety is not required and is not proved.

*Remark* 3 (Exact scope of the uncertainty set)*.* The target-rate uncertainty set consists of independent *fixed* multiplicative factors in $[99/100, 101/100]$ on each positive off-diagonal rate in (2), on each entry of (3), and on the division rate $b$. Zero rates remain zero, the $Q$ diagonal is reconstructed to keep row sums zero, and the allocation law (4) is fixed, as are the state space and the initial target preparation. The same factor applies throughout the course, and a single administration works for every choice in the box without observing that choice. This is a deterministic robustness specification, not a statistical confidence region, and it is narrower than "one-percent uncertainty in every biological parameter": in particular the daughter law is not varied.

The joint process can be constructed from independent driving Poisson random measures using predictable thinning. Under bounded controls the target count is dominated by a Yule process with per-cell rate 101/1000, and all other target intensities are bounded by a constant times its count on a finite horizon. The healthy space is finite. These bounds give nonexplosion and justify localization for the generator and expectation arguments below [3].

## 3 A robust finite-mission guarantee

**Theorem 4** (Delivered eradication with partially filled reserve)**.** *Consider the source in Section 2, including its one-percent target-rate uncertainty set of Remark 3. Let*

$$\begin{gathered} k_e \in [99/100, 101/100], \quad K = 400, \quad h_{\min} = 200, \\ r_{\text{eff}} \in [1,2], \quad \sigma \in [2,3], \quad \nu \in [0, 1/200], \quad H_0 \in \{278, \dots, 400\}. \end{gathered} \tag{11}$$

*One and the same administration (10), without observing any unknown parameter, obeys every stated limit and satisfies*

$$\mathbb{P}(|Z_{120}| > 0) < \frac{107}{12500} = 0.00856, \qquad \mathbb{P}(\tau_H \le 120) < \frac{7}{10^6}. \tag{12}$$

*Consequently*

$$\mathbb{P}(Z_{120} = 0,\ \tau_H > 120) > 0.991433. \tag{13}$$

*The result remains valid when the healthy rates vary predictably while preserving the comparison bounds in (7) with $r_{\text{eff}} \ge 1$ and the mortality ceiling established in (37).*

The conclusion consists of two separate marginal guarantees; the number in (13) is their joint consequence, not an extinction probability by itself. The role of the large reserve is explicit and is not a relabelling: the starting count need only be 278, but the model has 400 distinct units and a threshold of 200. The proof does not achieve the improved renewal requirement by changing the time unit or by subdividing the original 20-unit source; Table 2 records the dimensionless ratios that are actually changed.

Theorem 4 follows by combining the finite-threshold reserve theorem of Section 4 with the source-specific drift and delivery calculation of Section 6. Neither side of that argument assumes the desired probability bound as an input. Section 6.3 then shows that the hypotheses on $H_0$ and on $K$ in (11) are far from tight: Corollary 13 keeps both mission tolerances down to $H_0 = 209$, and Corollary 14 replaces $K = 400$ by $K = 232$.

# 4 A return-level certificate for reserve loss

**Theorem 5** (Anchored reserve bound)**.** *Fix integers* $1 \le h_{\min} \le M \le K$*, constants* $r_{\text{eff}}, m, T > 0$*, and put* $L = h_{\min} - 1$*. Define*

$$P_L = 1, \qquad P_j = \prod_{h=h_{\min}}^{j} \frac{Km}{r_{\text{eff}}(K-h)} \quad (L < j < M), \qquad D_M = \sum_{j=L}^{M-1} P_j. \tag{14}$$

*For the constant chain* (7)*, the probability, starting from* $M-1$*, of hitting* $L$ *before* $M$ *is*

$$q_M = \frac{P_{M-1}}{D_M}. \tag{15}$$

*For every initial count* $h_0 \ge M$*,*

$$\mathbb{P}_{h_0}(\tau_H \le T) \le \min\{1, MmTq_M\} \le \min\left\{1, MmT \prod_{h=h_{\min}}^{M-1} \frac{Km}{r_{\text{eff}}(K-h)}\right\}. \tag{16}$$

*The upper bounds also hold for adapted unit-jump rates with births at least the comparison rates and deaths at most the comparison rates, on the same capped state space.*

*Proof.* Let $p_h$ be the probability of hitting $L$ before $M$. At an interior state,

$$\lambda_h(p_{h+1} - p_h) = \mu_h(p_h - p_{h-1}), \qquad p_L = 1, \quad p_M = 0.$$

Successive differences are multiplied by $\mu_h/\lambda_h = Km/[r_{\text{eff}}(K-h)]$. Summing those differences between the two boundaries gives the standard scale-function solution [8],

$$p_h = \frac{\sum_{j=h}^{M-1} P_j}{D_M}. \tag{17}$$

In particular this proves (15). If $M = h_{\min}$, the departure already crosses the threshold and the empty-product convention gives $q_M = 1$.

Starting at or above $M$, every failure must be part of an excursion initiated by a jump $M \to M-1$. That departure counting process has intensity $Mm\mathbf{1}_{\{H_t = M\}}$, so its compensator, and hence its expected number by time $T$, is at most $MmT$. At each departure the strong Markov property gives eventual failure probability $q_M$. Summing these conditional probabilities over departures and applying the union bound gives the first inequality; counting a failure occurring after $T$ from an excursion begun before $T$ only enlarges the event. Since $D_M \ge P_L = 1$, the second inequality follows. Upward excursions above $M$ must revisit $M$ before reaching the lower boundary, so they do not escape this accounting.

For the adapted source, couple it above a constant comparison chain. At equal states use common jumps at the smaller rates; assign excess actual births and excess comparison deaths to the appropriate chain, which preserves the order. When the states differ, unit jumps cannot cross without first meeting. This nearest-neighbour coupling requires no monotonicity of the logistic birth rate in the state. The threshold-crossing event of the actual chain is therefore contained in that of the lower comparison chain, which proves the extension, including its boundary and initial-state assertions. □

**Corollary 6** (Incomplete or random preparation)**.** *For $h_{\min} \le h_0 < M$, the same comparison gives*

$$\mathbb{P}_{h_0}(\tau_H \le T) \le \min\left\{1, \frac{\sum_{j=h_0}^{M-1} P_j}{D_M} + MmTq_M\right\}, \tag{18}$$

*and the right-hand side is nonincreasing in $h_0$. For a random initial count, average the first term over its actual law, using zero for counts at least $M$ and one for counts below $h_{\min}$. In particular, if $\mathbb{P}(H_0 < M) \le \eta$, the bound is at most* $\min\{1, \eta + MmTq_M\}$.

*Proof.* For the constant comparison chain, first stop on reaching $L$ or $M$. Equation (17) bounds failure before reaching $M$; upon reaching $M$ the remaining time is no greater than $T$, so Theorem 5 applies at that endpoint, and the strong Markov property assembles the two stages. Monotonicity in $h_0$ holds because the numerator in (17) is a sum of nonnegative terms over $j \ge h_0$. Order coupling transfers the result to the adapted source, and averaging the conditional initial-state bounds proves the random-preparation statement. □

*Remark* 7 (Which of the two bounds to use). The two inequalities in (16) differ by the factor $D_M \ge 1$, which is a genuine improvement when the interior products decay slowly. At the operating instance of Theorem 4 one has $D_{278} \approx 2.59809$, so retaining the denominator improves the certificate by a factor of about 2.6. We state the main theorem with the cruder product because that is the inequality carried by the formal certificate; the sharper form is used in Corollaries 13 and 14, where it is evaluated in exact rational arithmetic. Proposition 10 shows that at high renewal the two agree to leading order, since then $D_M \to 1$.

### 4.1 Capacity at finite renewal

The product in (16) is itself a finite design rule. Its capacity dependence becomes transparent when the threshold is a fixed fraction of the reserve.

**Theorem 8** (Capacity law with explicit rounding)**.** *Fix $m, r_{\text{eff}} > 0$ and $0 < \theta < 1$ such that $x_* = 1 - m/r_{\text{eff}} > \theta$. For each integer $K$, set*

$$h_{\min} = \lceil K\theta \rceil, \qquad M = \lfloor Kx_* \rfloor,$$

*and assume $h_{\min} \le M$ and $H_0 \ge M$. Put*

$$f_\theta = \log \frac{r_{\text{eff}}(1-\theta)}{m},$$
$$\mathcal{I}(\theta; r_{\text{eff}}, m) = (1-\theta)\log\frac{r_{\text{eff}}(1-\theta)}{m} - (1-\theta) + \frac{m}{r_{\text{eff}}} > 0. \tag{19}$$

*Then the sources covered by Theorem 5 satisfy*

$$\mathbb{P}(\tau_H \le T) \le \min\{1, KmT\exp[-K\mathcal{I} + 2f_\theta]\}. \tag{20}$$

*For any positive sequence $T_K$ with $\log T_K = o(K)$, the corresponding upper exponential rate is at most $-\mathcal{I}$.*

*Proof.* The function $f(x) = \log[r_{\text{eff}}(1-x)/m]$ is decreasing and nonnegative on $[\theta, x_*]$. The negative logarithm of the product in (16) is a left Riemann sum for $Kf$,

$$S_K = \sum_{h=h_{\min}}^{M-1} f(h/K) \ge K\int_{h_{\min}/K}^{M/K} f(x)\,dx.$$

Each omitted end interval has length less than $1/K$ and integrand at most $f_\theta$. Consequently

$$S_K \geq K \int_\theta^{x_*} f(x)\,dx - 2f_\theta = K\mathcal{I} - 2f_\theta.$$

The integral is strictly positive because $x_* > \theta$; elementary integration gives (19). Substitute this estimate into (16) and use $M \leq K$. Dividing the logarithm of the resulting bound by $K$ gives the last assertion. □

Thus an explicit sufficient sizing condition is

$$K\mathcal{I} \geq \log \frac{KmT}{\alpha} + 2f_\theta, \qquad \lceil K\theta \rceil \leq \lfloor Kx_* \rfloor, \quad H_0 \geq \lfloor Kx_* \rfloor. \tag{21}$$

The constant $2f_\theta$ controls integer rounding; it is not a hidden $K$-dependent error. Merely assuming $h_{\min}/K \to \theta$ would generally give an $o(K)$ correction rather than this fixed one. Equation (20) is an upper-bound exponent, not a claim about an exact fixed-time large-deviation rate, and no matching lower bound is asserted in this regime. At the instance of Theorem 4 the envelope (20) evaluates to about $3.434 \times 10^{-5}$, looser than the exact product $6.938 \times 10^{-6}$: use the product for finite design and the action for interpretation.

## 5 Sharpness of the reserve certificate

Two statements delimit what Theorem 5 and Theorem 8 can be expected to give. The first shows that the regime hypothesis $x_* > \theta$ of the capacity law is not a technical artifact: below it, no amount of capacity helps. The second shows that in the opposite, high-renewal direction the anchored product is not merely an upper bound but is asymptotically exact.

### 5.1 Capacity cannot compensate for an inadequate equilibrium

**Proposition 9** (Failure of the subcritical-reserve regime)**.** *Fix $r_{\text{eff}} > m > 0$ and $\theta \in (0,1)$ with $x_* = 1 - m/r_{\text{eff}} < \theta$, and let*

$$t_{\text{cross}} = \frac{1}{r_{\text{eff}} - m} \log \frac{m\theta}{r_{\text{eff}}\theta - r_{\text{eff}} + m} \in (0, \infty).$$

*For each integer $K$ let $H^K$ be the constant chain (7) started full, $H_0^K = K$, with threshold $h_{\min} = \lceil K\theta \rceil$. If $T > t_{\text{cross}}$ then*

$$\mathbb{P}(\tau_H \leq T) \longrightarrow 1 \qquad (K \to \infty).$$

*Quantitatively, let $x$ solve $\dot{x} = r_{\text{eff}} x(1-x) - mx$ with $x(0) = 1$ and let $\varepsilon \in (0, \theta - x(T))$. Then*

$$\mathbb{P}(\tau_H \leq T) \;\geq\; 1 - \frac{4(r_{\text{eff}}/4 + m)\,T\,e^{2(r_{\text{eff}}+m)T}}{K\varepsilon^2}. \tag{22}$$

*Proof.* Put $X^K = H^K/K$. The chain is density dependent in the sense of Kurtz [9], and

$$X^K(t) = 1 + \int_0^t b(X^K(s))\,ds + \mathcal{M}_K(t), \qquad b(x) = r_{\text{eff}} x(1-x) - mx,$$

where $\mathcal{M}_K$ is a martingale whose jumps have size $1/K$ and whose predictable quadratic variation has rate $K^{-1}[r_{\text{eff}} x(1-x) + mx] \leq (r_{\text{eff}}/4 + m)/K$ on $[0,1]$. Doob's $L^2$ inequality gives $\mathbb{E}\sup_{[0,T]} \mathcal{M}_K^2 \leq$

$4(r_{\mathrm{eff}}/4+m)T/K$. On $[0,1]$ the drift $b$ is Lipschitz with constant $r_{\mathrm{eff}}+m$, so Gronwall's inequality yields $\sup_{[0,T]}|X^K-x| \le e^{(r_{\mathrm{eff}}+m)T}\sup_{[0,T]}|\mathcal{M}_K|$, and Chebyshev's inequality gives (22) after substitution.

The deterministic equation is logistic with intrinsic rate $r_{\mathrm{eff}}-m$ and carrying level $x_*$, so from $x(0)=1>x_*$ its solution decreases to $x_*$ and is

$$x(t)=\frac{x_*}{1-(m/r_{\mathrm{eff}})e^{-(r_{\mathrm{eff}}-m)t}}.$$

Solving $x(t)=\theta$ gives exactly $t=t_{\mathrm{cross}}$, which is positive because $\theta<1$, and finite because $x_*<\theta$. Hence $x(T)<\theta$ for $T>t_{\mathrm{cross}}$, and $\varepsilon$ may be chosen as stated. On the event $\sup_{[0,T]}|X^K-x|\le\varepsilon$ we have $H_T^K<K\theta\le h_{\min}$, so $\tau_H\le T$. The complement has probability at most the displayed remainder. □

Proposition 9 is a statement about one constant healthy source, not an all-policy therapeutic exclusion; a policy that reduces $m$ changes $x_*$. Its force is comparative. Theorem 8 makes failure exponentially small in $K$ when the equilibrium fraction $x_*$ exceeds the threshold fraction $\theta$, whereas here, with $x_*<\theta$, failure tends to one in $K$ for every horizon past a fixed crossing time. The sign of $x_*-\theta$, not the size of $K$, decides which regime the design is in, and the regime boundary is a property of rates rather than of scale. Adding units to a reserve whose equilibrium already sits below the functional threshold buys nothing.

## 5.2 The anchored bound is exact at high renewal

For a full reserve the natural anchor is the ceiling itself, $M=K$, and the buffer $d=K-h_{\min}$ is the number of units that must be lost in one excursion to fail.

**Proposition 10** (High-renewal limit and asymptotic sharpness)**.** *Fix integers $1\le h_{\min}<K$, put $d=K-h_{\min}\ge 1$, and fix $m,T>0$. For the constant chain (7) started at $H_0=K$:*

*(i) For every $r_{\mathrm{eff}}>0$ the top-anchor bound of Theorem 5 with $M=K$ is exactly*

$$KmT\,P_{K-1}=\frac{(Km)^{d+1}T}{d!\,r_{\mathrm{eff}}^d}. \tag{23}$$

*(ii) Moreover*

$$\lim_{r_{\mathrm{eff}}\to\infty} r_{\mathrm{eff}}^d\,\mathbb{P}_K(\tau_H\le T)=\frac{(Km)^{d+1}T}{d!}. \tag{24}$$

*Hence the certificate (16) is sharp to leading order as $r_{\mathrm{eff}}\to\infty$, and the sufficient renewal requirement (41) has the correct order in every parameter.*

*Proof.* (i) With $M=K$ and $L=h_{\min}-1$, reindex the product in (14) by $k=K-h$, which runs over $1,\dots,d$:

$$P_{K-1}=\prod_{h=h_{\min}}^{K-1}\frac{Km}{r_{\mathrm{eff}}(K-h)}=\prod_{k=1}^{d}\frac{Km}{r_{\mathrm{eff}}k}=\frac{(Km)^d}{r_{\mathrm{eff}}^d\,d!}.$$

Multiplying by $MmT=KmT$ gives (23). In particular Theorem 5 already yields the upper half of (24),

$$\limsup_{r_{\mathrm{eff}}\to\infty} r_{\mathrm{eff}}^d\,\mathbb{P}_K(\tau_H\le T)\le\frac{(Km)^{d+1}T}{d!}.$$

(ii) Write $q_r = q_K = P_{K-1}/D_K$ for the probability that an excursion leaving $K$ downward reaches $L$ before returning to $K$. Since $P_j$ is of order $r_{\text{eff}}^{-(j-L)}$ for $j > L$ while $P_L = 1$, we have $D_K \to 1$ and

$$q_r = \frac{(Km)^d}{d!\, r_{\text{eff}}^d} \left(1 + O(r_{\text{eff}}^{-1})\right). \tag{25}$$

Let $N_T$ be the number of downward departures from $K$ strictly before $\min(T, \tau_H)$. Its compensator is $Km \int_0^T \mathbf{1}_{\{H_s=K,\ s<\tau_H\}} ds$, so

$$\mathbb{E}N_T = Km\, \mathbb{E}\int_0^T \mathbf{1}_{\{H_s=K,\ s<\tau_H\}}\, ds. \tag{26}$$

Failure is absorbing, so at most one excursion ever fails. By the strong Markov property at each departure and Wald's identity for the departure counting process,

$$\mathbb{P}(\text{some excursion begun before } T \text{ eventually fails}) = q_r\, \mathbb{E}N_T. \tag{27}$$

The event $\{\tau_H \le T\}$ differs from the event in (27) only by excursions that begin before $T$ and fail after $T$. We control the two discrepancies in turn.

*Excursion durations.* Let $D$ be the duration of an excursion from $K-1$ until it hits $\{L, K\}$. On the finitely many interior states $h \in \{h_{\min}, \ldots, K-1\}$ the total jump rate satisfies $\lambda_h + \mu_h \ge r_{\text{eff}} h(K-h)/K \ge r_{\text{eff}} c_*$ with $c_* = \min_{h_{\min} \le h \le K-1} h(K-h)/K > 0$, so every holding time is stochastically dominated by an exponential of rate $r_{\text{eff}} c_*$. Moreover the upward one-step probability $\lambda_h/(\lambda_h + \mu_h) \to 1$ uniformly on this finite set, so a block of at most $d$ consecutive jumps returns the chain to $K$ with probability bounded below uniformly in $r_{\text{eff}}$ for $r_{\text{eff}}$ large. The number of jumps before absorption is therefore dominated by a geometric number of such blocks, with $r_{\text{eff}}$-independent mean, whence

$$\mathbb{E}D = O(r_{\text{eff}}^{-1}). \tag{28}$$

*Occupation time at the ceiling.* The total time spent away from $K$ before $\min(T, \tau_H)$ is the sum over departures of the corresponding durations; by Wald and $\mathbb{E}N_T \le KmT$ it has expectation at most $KmT\, \mathbb{E}D = O(r_{\text{eff}}^{-1})$. Time after absorption contributes at most $T\, \mathbb{P}(\tau_H \le T) = O(q_r)$. Hence the expectation in (26) equals $T - O(r_{\text{eff}}^{-1}) - O(q_r)$ and

$$\mathbb{E}N_T = KmT + O(r_{\text{eff}}^{-1}) + O(q_r). \tag{29}$$

*Terminal correction.* Let $D_f$ be the excursion duration conditional on failure. Conditioning a killed birth–death chain on hitting $L$ before $K$ is a Doob $h$-transform by the harmonic function $p_\cdot$ of (17), which leaves holding rates unchanged and replaces the downward one-step probability at $h$ by

$$\frac{\mu_h\, p_{h-1}}{(\lambda_h + \mu_h)\, p_h}.$$

Since $p_h \sim P_h$ and $P_{h-1}/P_h = r_{\text{eff}}(K-h)/(Km)$, while $\mu_h/(\lambda_h + \mu_h) = Km/[r_{\text{eff}}(K-h) + Km]$, this probability tends to 1 uniformly on the finite interior. Conditionally on failure the chain therefore makes $O(1)$ jumps at rates at least $r_{\text{eff}} c_*$, so $\mathbb{E}D_f = O(r_{\text{eff}}^{-1})$ by the same block argument. Using the departure intensity bound and the strong Markov property, the expected number of excursions that begin before $T$ and fail after $T$ is at most

$$Km\, q_r \int_0^T \mathbb{P}(D_f > T - s)\, ds \ \le\ Km\, q_r\, \mathbb{E}D_f = O(q_r r_{\text{eff}}^{-1}).$$

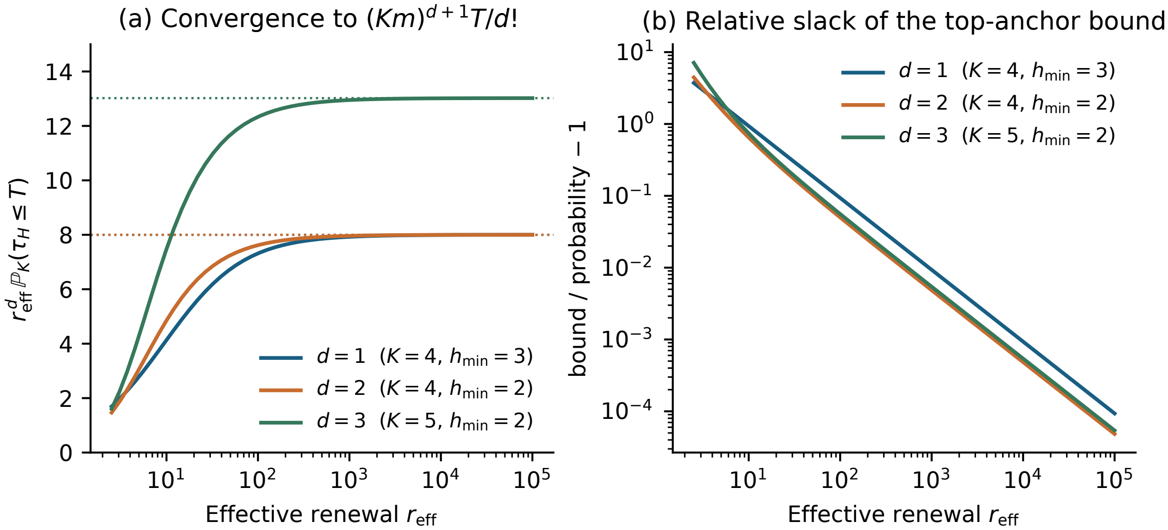


**Figure 1:** Proposition 10 against exact computation. (a) $r_{\mathrm{eff}}^d \mathbb{P}_K(\tau_H \le T)$ for the constant chain started full, at $m = 1/2$ and $T = 2$, obtained from the killed Kolmogorov equation; dotted lines mark the predicted constants $(Km)^{d+1}T/d!$, which equal 8, 8 and 625/48 respectively. (b) Relative slack of the top-anchor bound (23), which decays like $1/r_{\mathrm{eff}}$. Because the probabilities reach $10^{-15}$ and below, these curves were evaluated at 60 decimal digits; in double precision the subtraction $1 - \sum_j p_{Kj}(T)$ loses all significant figures and can appear to violate the proved bound. The bound exceeds the exact probability at every plotted point.

Combining this with (27), (29) and (25),

$$\mathbb{P}(\tau_H \le T) = q_r\,\mathbb{E}N_T + O(q_r r_{\mathrm{eff}}^{-1}) = \frac{(Km)^{d+1}T}{d!\,r_{\mathrm{eff}}^d}\left(1 + O(r_{\mathrm{eff}}^{-1})\right),$$

which is (24). □

*Remark* 11 (Depleted preparations have a different exponent). The exponent in (24) is a property of a *full* reserve. If instead $h_{\min} \le h_0 < K$, the initial scale-function term of Corollary 6 is itself of order $r_{\mathrm{eff}}^{-(h_0 - h_{\min}+1)}$, so the failure probability is of order $r_{\mathrm{eff}}^{-\min\{d,\, h_0 - h_{\min}+1\}}$. One may not transfer the fully filled exponent to a depleted preparation: a reserve that begins one unit above the threshold gains only a single power of renewal, however large the buffer $d$ is.

*Remark* 12 (What is not claimed). Proposition 10 is a homogeneous statement with fixed $K, m, T$. It does not license replacing $m$ by a time-varying $m(t)$, interchanging the limit with an arbitrary schedule, or deriving a Jensen-type optimal scheduling law. A time-varying version would need inhomogeneous conditional-excursion estimates and an $o(r_{\mathrm{eff}}^{-d})$ remainder that is uniform over rate profiles; renewal-dependent oscillations are a concrete obstruction. We therefore leave the time-varying coefficient open and do not use it anywhere below.

# 6 Source-specific proof of the delivered mission

## 6.1 A weight that remains valid under rate uncertainty

Use

$$w = (14, 11, 10, 84, 43, 107)^\top, \qquad w_{\min} = 10. \tag{30}$$

Direct rational evaluation of (6) at the endpoints gives

$$\begin{aligned} A_e w &\le -\frac{99}{1000} w && (29/100 \le e \le 31/100), \\ A_e w &\le \frac{67}{500} w && (1/100 \le e \le 31/100). \end{aligned} \tag{31}$$

The endpoint checks suffice because each component of $A_e w/w$ is affine in $e$, so its maximum over an interval is attained at an endpoint. The reconstructed row values and their checked inequalities are given in Appendix A. In particular the largest weighted drift ratios at $e = .01, .29, .31$ are, respectively,

$$\frac{559}{4200}, \qquad -\frac{533}{5350}, \qquad -\frac{73}{700}.$$

If $\widetilde{A}_e$ belongs to the one-percent box of Remark 3, then changing off-diagonal chemistry rates also changes their diagonal row sums. Accounting for both changes gives, for each type,

$$\begin{aligned} |((\widetilde{A}_e - A_e)w)_i| &\le \frac{1}{100}\left[\sum_{j\neq i}(Q_e)_{ij}|w_j - w_i| + \frac{1}{10}\,|(Lw)_i - w_i| + d_i w_i\right] \\ &\le \frac{27}{2000}\, w_i, \qquad 1/100 \le e \le 31/100. \end{aligned} \tag{32}$$

The three terms are the chemistry, division and mortality contributions. The bracket divided by $w_i$ is affine and nondecreasing in $e$, so it is maximal at $e = 31/100$, where its largest value is $1427/1075 < 27/20$. The envelope therefore covers every point of the box, not a sample of perturbed models. From (31)–(32),

$$\widetilde{A}_e w \le g\, w \quad \text{on } [.01, .31], \qquad \widetilde{A}_e w \le -\gamma\, w \quad \text{on } [.29, .31], \qquad g = \frac{59}{400}, \quad \gamma = \frac{171}{2000}. \tag{33}$$

### 6.2 Delivery, extinction and the healthy path

*Proof of Theorem 4.* For the fixed course (10), concentration is

$$c(t) = \begin{cases} \dfrac{29}{100k_e}(1 - e^{-k_e t}), & 0 \le t \le 112, \\ c(112)e^{-k_e(t-112)}, & 112 < t \le 120. \end{cases} \tag{34}$$

Consequently $0 \le c(t) \le 29/99 < 3/10$ uniformly. For $4 \le t \le 112$, bounding $1/k_e$ and $1 - e^{-k_e t}$ separately from below,

$$c(t) \ge \frac{29}{101}\left(1 - e^{-99/25}\right) > \frac{29}{101} \cdot \frac{49}{50} = \frac{1421}{5050} > \frac{28}{100}, \tag{35}$$

where $e^{99/25} > 50$ follows from a positive Taylor sum through degree 15. Hence the effective erasure $e = 1/100 + c$ lies in $[.01, .31]$ throughout and in $[.29, .31]$ on $[4, 112]$: the first four units lie in the growth interval of (33) and the next 108 units lie in its contraction interval.

For $Y_t = Z_t \cdot w$, localization at target count $N$ together with the generator identity yields the corresponding expectation inequalities; the Yule domination of Section 2 provides the integrable finite-horizon bound needed to let $N \to \infty$. Gronwall's inequality and $Y_t \ge w_{\min}\mathbf{1}_{\{|Z_t|>0\}}$ then give

$$\begin{aligned} \mathbb{P}(|Z_{112}| > 0) &\le \frac{4 \cdot 107}{10} \exp(4g - 108\gamma) \\ &= \frac{214}{5} \exp\left(-\frac{2161}{250}\right) < \frac{214/5}{5000} = \frac{107}{12500}. \end{aligned} \tag{36}$$

Indeed $2161/250 > 43/5$, and the positive Taylor sum for $e^{43/5}$ through degree 30 exceeds 5000. No target immigration is present, so extinction is absorbing and survival at 120 implies survival at 112; it is unnecessary to assume contraction during washout.

For healthy units, (11) and the uniform concentration bound imply

$$m_H(t) \le \frac{3/10 + 29/99}{2} + \frac{1}{200} < \frac{61}{200} =: m. \tag{37}$$

Apply Theorem 5 with $K = 400$, $h_{\min} = 200$, $M = 278$, $r_{\text{eff}} = 1$, $m = 61/200$ and the full horizon $T = 120$. Its simpler product bound is

$$278 \cdot \frac{61}{200} \cdot 120 \prod_{h=200}^{277} \frac{122}{400 - h} < \frac{7}{10^6}, \tag{38}$$

an exact rational inequality. Every initial count allowed by (11) is at least the anchor. The bound counts crossings followed by later recovery, and it already includes the entire washout period. Combining it with (36) proves the two marginal statements in (12) and, by the union bound, (13).

Finally, the total administered amount is $112(29/100) = 812/25 = 32.48 < 33$. Integrating (8) gives

$$k_e \int_0^{120} c(t)\,dt = \int_0^{120} a(t)\,dt - c(120), \qquad \int_0^{120} c(t)\,dt \le \frac{3248}{99} < 33. \tag{39}$$

The amplitude limits $a \le 3/10$ and $c \le 29/99 < 3/10$ have already been checked. The post-stop concentration integral is at most $2900/9801$ and was not removed from the healthy law. Every step holds for all parameters in (11) and all target-rate factors in the declared box under the same fixed administration. This proves the theorem. □

The proof uses only three properties of the input: an effective-erasure band, a healthy mortality ceiling, and the administration constraints. For other response laws, including saturating ones, the same implication holds whenever the actual delivered trajectory satisfies those three requirements. Such a statement is conditional on a checked response band; it is not evidence that an arbitrary Hill law realizes this example, and if saturation prevents entry into the band then this weight certifies nothing about that course either way.

### 6.3 How little reserve the same course needs

The hypotheses on $H_0$ and $K$ in (11) were chosen to make the healthy margin comfortable, not to make the mission feasible. Two exact-arithmetic consequences quantify the difference. Both keep the administration (10), the target uncertainty box, the four $AA$ founders, $T = 120$, $r_{\text{eff}} \ge 1$, $\sigma \in [2, 3]$, $\nu \in [0, 1/200]$ and $k_e \in [.99, 1.01]$; hence the target bound (36) is unchanged and only the healthy side moves.

**Corollary 13** (Mission tolerance survives a nearly depleted reserve)**.** *In the setting of Theorem 4 with $K = 400$ and $h_{\min} = 200$, replace the requirement $H_0 \ge 278$ by $H_0 \ge h_0$. Then, evaluating (18) with $M = 278$ in exact rational arithmetic,*

$$\mathbb{P}(\tau_H \le 120) < \begin{cases} 9.827 \times 10^{-3}, & h_0 = 209, \\ 8.191 \times 10^{-5}, & h_0 = 221, \\ 2.908 \times 10^{-6}, & h_0 = 240. \end{cases}$$

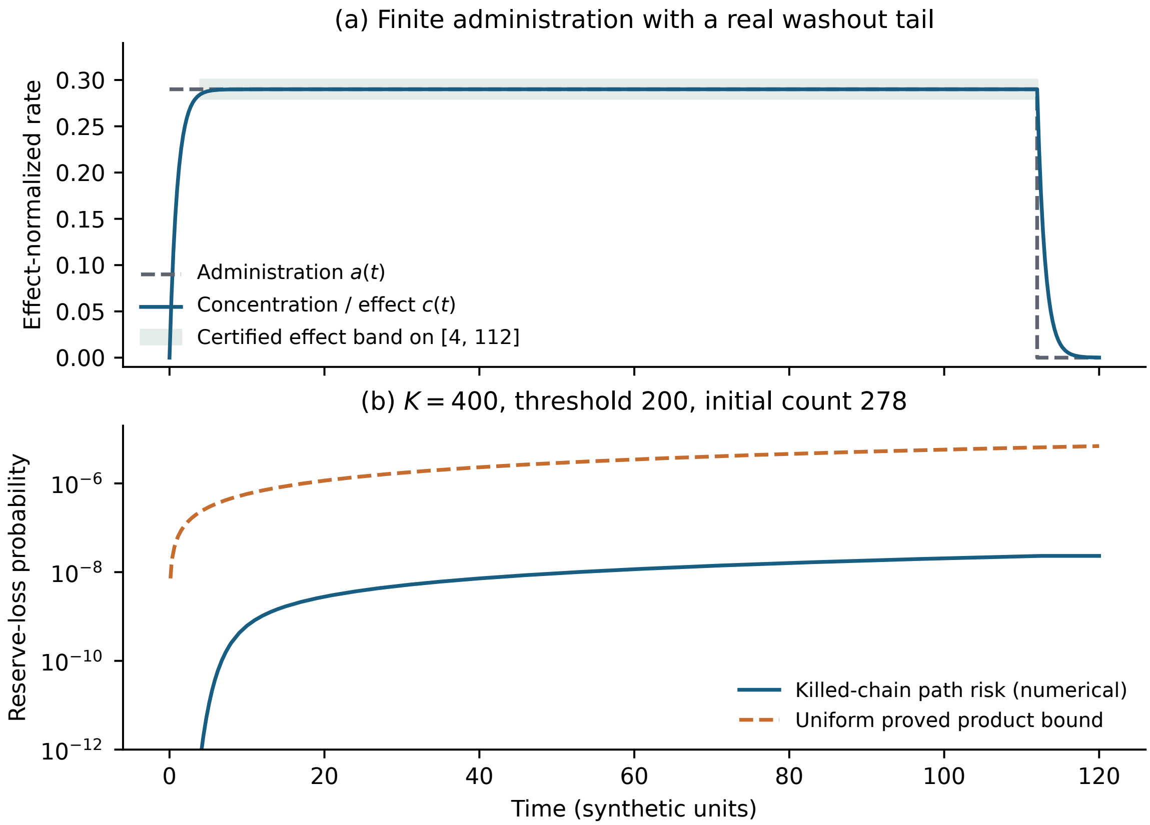


**Figure 2:** The delivered course and reserve risk. The concentration curve uses nominal clearance $k_e = 1$; the band and probability envelope are uniform over the theorem's box. The healthy curve solves the killed forward equation numerically at $K = 400$, $h_{\min} = 200$, $H_0 = 278$, $r_{\text{eff}} = 1$, $\sigma = 2$, $\nu = .005$ for that one parameter trajectory. It is an illustrative floating computation, not a Monte Carlo estimate and not the certificate. The dashed bound uses the worst constant mortality, is uniform over the box, and retains all threshold crossings; the two curves are therefore different probabilistic objects, which is why the visible gap is not a measure of looseness at a fixed parameter.

*In particular both mission tolerances $\delta = \alpha = 1/100$ are met from any initial count $h_0 \geq 209$, with joint success greater than $0.98161$; and $h_0 \geq 240$ gives joint success greater than $0.991437$. The value $209$ is the exact minimum for which this certificate meets $\alpha = 1/100$: at $h_0 = 208$ the same expression exceeds $1/100$.*

*Proof.* The right-hand side of (18) is nonincreasing in $h_0$ by Corollary 6, so it suffices to evaluate it at the displayed counts. With $K = 400$, $h_{\min} = 200$, $M = 278$, $r_{\text{eff}} = 1$, $m = 61/200$, $T = 120$ one has $D_{278} = 2.5980909040\ldots$ and $MmTq_M = 2.6704763715\ldots \times 10^{-6}$, and the initial scale term $\sum_{j \geq h_0} P_j/D_M$ evaluates to $9.8243 \times 10^{-3}$, $7.9238 \times 10^{-5}$ and $2.3748 \times 10^{-7}$ at $h_0 = 209, 221, 240$. Adding the two terms gives the displayed bounds. The target bound (36) is unaffected, and the union bound gives the joint statements. All quantities are ratios of integers and were evaluated exactly. $\square$

**Corollary 14** (A smaller sufficient reserve at the same tolerances)**.** *Keep every assumption of Theorem 4 except the reserve preparation, and take instead*

$$K = 232, \qquad h_{\min} = 116, \qquad M = 161, \qquad H_0 \geq 161.$$

*Then* $\mathbb{P}(\tau_H \le 120) < 9.781 \times 10^{-3} < 1/100$ *and* $\mathbb{P}(|Z_{120}| > 0) < 107/12500$*, so both mission tolerances hold and*

$$\mathbb{P}(Z_{120} = 0,\ \tau_H > 120) > 0.98165.$$

*Proof.* The threshold is again half of capacity and the anchor is again the worst-case equilibrium count, $\lfloor 232(1 - .305)\rfloor = 161$. With $r_{\text{eff}} = 1$, $m = 61/200$ and $T = 120$, the first bound of (16) gives

$$\mathbb{P}(\tau_H \le 120) \le \frac{MmT\,P_{M-1}}{D_M} = \frac{161 \cdot \frac{61}{200} \cdot 120}{D_{161}} \prod_{h=116}^{160} \frac{232 \cdot \frac{61}{200}}{232 - h} = 9.7801\ldots \times 10^{-3},$$

with $D_{161} = 2.6271217\ldots$, again an exact rational quantity. The mortality ceiling (37) and the target argument are untouched by the change of reserve, so (36) still applies and the union bound gives the joint statement. □

*Remark* 15 (The denominator is doing real work here). Corollary 14 uses the sharper of the two bounds in (16). The cruder product $MmTP_{M-1}$ alone equals $2.569 \times 10^{-2}$ at $K = 232$ and does not meet the tolerance; the smallest capacity it certifies in this proportional family is $K = 252$, where it gives $9.831 \times 10^{-3}$. Within the same family, $K = 232$ is the smallest capacity certified by the sharper bound. Neither statement is a minimum-capacity theorem: both are sufficient constructions from one proof method, and no converse excludes smaller reserves.

# 7 Design consequences and comparison of preparations

## 7.1 Capacity, initial filling and retained renewal

Table 1 compares the certified preparations under the same administration, the same amount 32.48 and the same horizon 120. The original small-reserve estimate uses the top anchor and a narrower target box: chemistry factors in $[.9999, 1.0001]$, division within $10^{-5}$ of .1, and each baseline death within $10^{-5}$ of its nominal value. Its weighted error is at most $.001w$, giving target risk below .00535. The wide-box results spend some of that probability margin to allow one-percent relative error in all these target rates. They also change capacity and selectivity, so the reduction of the renewal requirement from 12 to 1 is *not* a comparison at otherwise identical parameters; it is a statement about what a different preparation buys.

The three wide-box rows separate two resources that are easy to conflate. The *lower filling* row shows that the 278-unit anchor of Theorem 4 buys margin rather than feasibility: the certificate degrades gracefully, and it is only near 209, nine units above the failure threshold, that the mission tolerance is actually reached. The *smaller reserve* row shows the complementary trade: shrinking capacity by 42% at fixed proportions consumes essentially the same margin. Increasing $K$ and increasing $H_0/K$ enter different terms of the guarantee, the first through the product in (16) and the second through the initial scale term in (18). Figure 3 displays the two dependences side by side.

For $K = 400$, the return level 278 is close to the equilibrium count under the worst comparison source, $400(1 - .305) = 278$. Anchoring there rather than at the ceiling is what makes the certificate usable: an excursion is not required to return to an atypical hard ceiling that a treated reserve rarely visits. Theorem 5 remains valid above this level even when the initial reserve is not full, and Corollary 6 quantifies depleted or random preparation rather than silently resetting the reserve.

If intervention inhibits renewal, only the retained lower rate enters the certificate. In particular a nominal $r = 2$ with retained fraction $\chi_{\min} = 1/2$ meets $r_{\text{eff}} \ge 1$. This is a quantitative sufficient

**Table 1:** Certified preparations. All rows have $\nu \in [0, .005]$, clearance $k_e \in [.99, 1.01]$, four $AA$ founders, horizon 120 and the same delivered course (10). Healthy renewal may exceed the displayed minimum. The displayed probabilities are separate marginal upper bounds, not exact risks.

| Preparation | $K$ | $h_{\min}$ | $H_0 \geq$ | $r_{\text{eff}} \geq$ | $\sigma$ | Healthy loss | Source |
|---|---|---|---|---|---|---|---|
| Original benchmark | 20 | 10 | 20 | 12 | $[1, 2]$ | $5.0 \times 10^{-4}$ | |
| Same capacity | 20 | 10 | 20 | 9 | $[1, 2]$ | $8.5 \times 10^{-3}$ | |
| Larger reserve | 400 | 200 | 278 | 1 | $[2, 3]$ | $7.0 \times 10^{-6}$ | Thm. 4 |
| Lower filling | 400 | 200 | 240 | 1 | $[2, 3]$ | $2.9 \times 10^{-6}$ | Cor. 13 |
| Minimal filling | 400 | 200 | 209 | 1 | $[2, 3]$ | $9.8 \times 10^{-3}$ | Cor. 13 |
| Smaller reserve | 232 | 116 | 161 | 1 | $[2, 3]$ | $9.8 \times 10^{-3}$ | Cor. 14 |

| Preparation | Target uncertainty | Target survival |
|---|---|---|
| Original benchmark; same capacity | Narrow mixed relative/absolute box | 0.00535 |
| All wide-box rows | Independent one-percent relative box | 0.00856 |

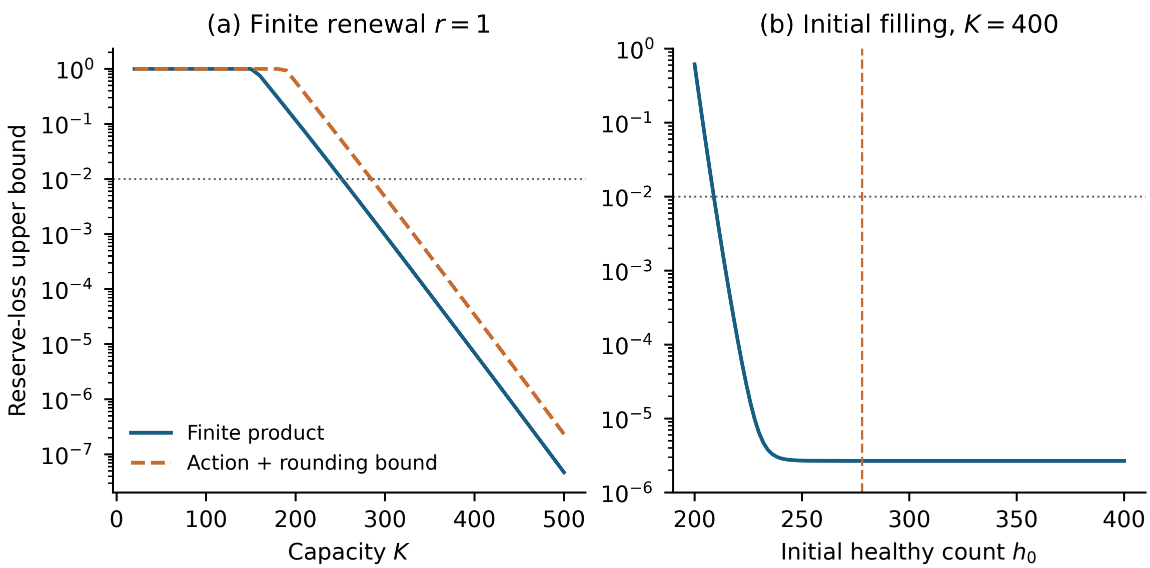


**Figure 3:** Two distinct reserve resources. (a) Finite-product and action bounds for $m = .305$, $r_{\text{eff}} = 1$, threshold $\lceil K/2 \rceil$ and initial count at least $\lfloor .695K \rfloor$, through $T = 120$. (b) The initial-state scale-function bound at $K = 400$, threshold 200, anchor 278. Increasing capacity and increasing initial filling change different terms of the guarantee. The horizontal dotted line is risk .01; the sharp knee in (b) near $h_0 \approx 210$ is the content of Corollary 13. Curves are evaluations of proved bounds, clipped at one.

condition on renewal retained *during* exposure; it does not assert that any particular intervention has that property, and renewal measured without treatment is not the relevant quantity.

## 7.2 Founder burden and finite mission limits

The weight (30) gives $z_0 \cdot w/w_{\min} \leq 10.7n$ for any $n$ initial target cells, regardless of their types. A four-unit ramp followed by a contraction interval of length $\ell$ therefore has the sufficient target bound

$$\mathbb{P}(|Z_{4+\ell}| > 0) \leq 10.7\, n \exp(4g - \gamma\ell), \qquad \ell \geq \gamma^{-1} \left[\log \frac{10.7n}{\delta} + 4g\right], \tag{40}$$

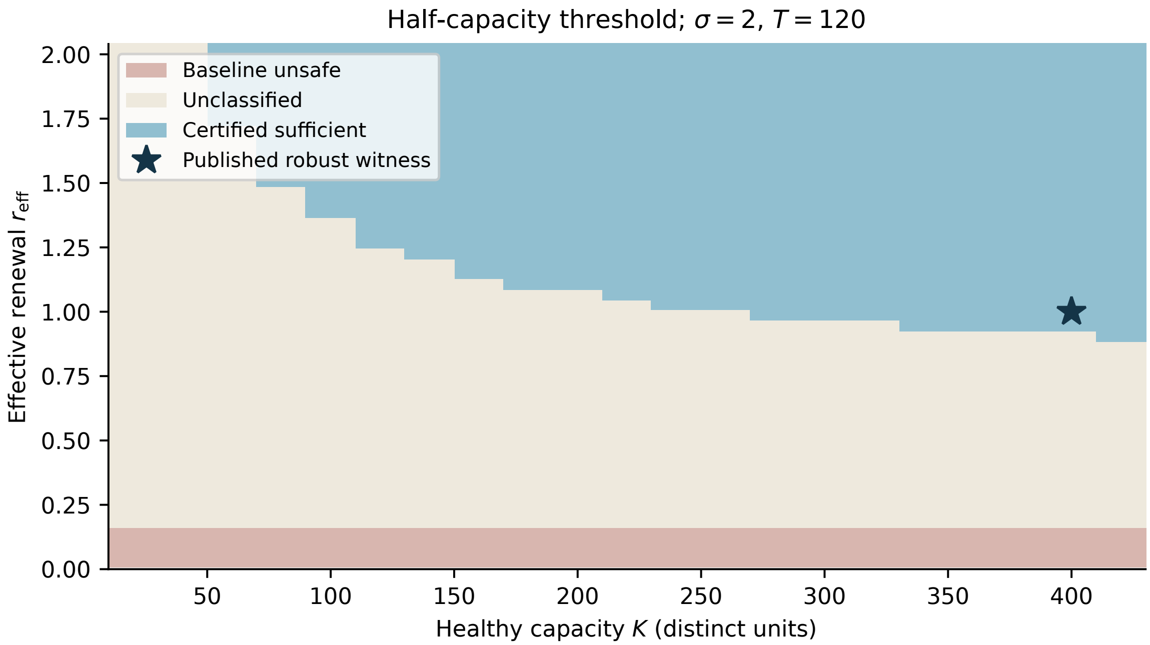


**Figure 4:** A sufficient design map, not an optimal frontier. At each displayed $(K, r_{\mathrm{eff}})$ a return level is selected and its finite-product inequality is checked rationally at $m = .305$, $h_{\min} = K/2$, $T = 120$; the sufficient region assumes initial count at least that selected level. Combined with the wide-box target certificate, blue points satisfy both .01 tolerances. The red band is already unsafe from baseline mortality .15 at $\sigma = 2$ even when initially full; its classification uses the expectation bound of Section 8. The remaining region is unclassified, not proved infeasible. The star is Theorem 4.

with zero sufficing when the displayed requirement is negative. For the fixed course of Theorem 4 the coarser rational bound is $107n/50000$, which certifies $n \leq 4$ at $\delta = .01$. The dependence on burden is only logarithmic, but the proof does not cover arbitrary burden at fixed amount and deadline, because a longer course is a different course: administration lasts $4 + \ell$ and requires amount $.29(4 + \ell)$, concentration-exposure allowance at least $.29(4 + \ell)/.99$, and a horizon including the declared washout. All three limits must be enlarged explicitly, and the enlarged horizon then feeds back into the reserve requirement. Substituting that horizon into (21) links capacity to founder burden through the logarithmic contraction duration. Alternatively the top-anchor product gives the sufficient renewal requirement

$$r_{\mathrm{eff}} \geq Km \left( \frac{KmT}{\alpha\, d!} \right)^{1/d}, \qquad d = K - h_{\min} \geq 1, \quad H_0 = K. \tag{41}$$

By Proposition 10 this rule has the right order: the bound it inverts is asymptotically exact as $r_{\mathrm{eff}} \to \infty$, so no argument of this type can improve its $\alpha^{-1/d}$ or $(d!)^{-1/d}$ dependence by more than a factor tending to one. It remains a sufficient resource law rather than a necessary one, and capacity, initial filling, renewal, selectivity and target burden should be kept as separate coordinates in any design comparison.

# 8 What the guarantee does and does not establish

## 8.1 Baseline failure is different from treatment-induced failure

The original small-reserve excluded example has $K = H_0 = 20$, $h_{\min} = 10$, $\sigma \le 2$ and $r \le .001$. For any nonnegative intervention, healthy per-unit mortality is at least .15 while births are at most $rh$. Therefore

$$\mathcal{G}H \le -(.15 - .001)H = -.149H, \qquad \mathbb{P}(H_{120} \ge 10) \le \frac{\mathbb{E}H_{120}}{10} \le 2e^{-17.88} < 10^{-7},$$

so path failure exceeds $1 - 10^{-7}$ under every policy. This is a genuine all-policy obstruction, but it reflects *baseline* reserve collapse: the reserve in that region dies without any treatment at all. It is therefore no evidence of a sharp treatment-induced selectivity boundary. More generally, when mortality is at least $m_0 > r_{\mathrm{eff}}$ and births are at most $r_{\mathrm{eff}}h$, survival above threshold is at most $(H_0/h_{\min})e^{-(m_0 - r_{\mathrm{eff}})T}$, and this elementary test is what labels the baseline-unsafe band of Figure 4.

## 8.2 The safe-baseline converse is open

The interesting question left open is the converse in the complementary region: if the untreated reserve is safe, is there a source for which *no* admitted policy achieves both tolerances? We do not resolve it, and we record why the natural approach fails, since the obstruction is concrete.

A standard route seeks a generator subsolution certifying an upper bound on joint success. With $W = 1 - q^z$ for a product supersolution $q$, the target generator satisfies $\mathcal{G}_C W \ge -(67v/73 + k)W$, where $v$ and $k$ are the delivered rate levels; one then looks for $e^{-\eta(T-t)} f(h) W$ with $0 \le f \le 1$ and $\mathcal{G}_H f + (\eta - 67v/73 - k)f \ge 0$ on all safe states. Because the constraint is affine in the rate levels, exact vertex testing on the admissible rectangle covers all trajectories. The attempt fails for a structural reason rather than a numerical one. Under the normalization $f(K) = 1$, births vanish at the ceiling, so the $K$-state row forces

$$\eta \;\ge\; \frac{67 \cdot \frac{3}{10}}{73} + \frac{1}{10} = 0.37534\ldots,$$

which is incompatible with the discount range in which the remaining rows admit a solution. Raising the failure penalty cannot repair this row, since it does not appear there. Eighteen exact linear programs over three renewal values and six discounts returned no separating objective beyond the zero function. This is an algebraic obstruction to that ansatz, not a theorem that a therapeutic window does or does not exist; closing the question needs a subsolution retaining more source, control or budget information. The sufficient mission theorem and the capacity law do not depend on it, and we therefore make no global frontier claim and no constant-versus-varying optimality claim.

## 8.3 Immediate regeneration is a substantive assumption

At a hard ceiling with $h_{\min} = K$ the first healthy death is already failure. Since $\lambda_K = 0$, renewal cannot act before that event, and per-unit mortality at least $m_0$ gives failure probability exactly $1 - e^{-Km_0T}$ for constant mortality, regardless of $r_{\mathrm{eff}}$. With a positive buffer the top-anchor bound instead tends to zero as $r_{\mathrm{eff}}$ grows, at the exact rate given by Proposition 10. The discontinuity at $d = 0$ belongs to the specified jump model.

**Table 2:** Scales and quantities that would require empirical identification. No clinical units are inferred.

| Quantity | Nominal value or ratio | Meaning / required information |
|---|---|---|
| Renewal / division | $r_{\text{eff}}/b = 10$ at the new lower bound | Retained renewal during exposure, not untreated renewal alone |
| Clearance / division | $k_e/b \in [9.9, 10.1]$ | Actual clearance and effect-site response |
| Healthy mortality / division | $m/b \le 3.05$ | Shared injury and healthy baseline loss |
| Target mortality / division | $d_i/b \in \{.1, 3\}$ nominally | State-specific baseline event rates |
| Reserve threshold | $h_{\min}/K = .5$ | Functional threshold and independently specified count units |
| Initial reserve | $H_0/K \ge .695$, or $\ge .5225$ by Cor. 13 | Actual filling or preparation distribution |
| Target inheritance | Equation (4) | Joint daughter allocation, not only mean trajectories |
| Uncertainty | Relative errors $\le 1\%$ on the rates of Rem. 3 | A mathematical box requiring external justification |
| Maturation | No delay in the main source | Transit time and initial pipeline if present |

Delayed replenishment can impose a risk floor even with a buffer. Suppose no replacement can complete during an initial latency $\tau \le T$, there is no preloaded maturation pipeline, $H_0 = K$, and mortality is at least $m_0$ per unit during that interval. Pure-death comparison gives

$$\mathbb{P}(\tau_H \le T) \ge \mathbb{P}\{\text{Bin}(K, e^{-m_0\tau}) < h_{\min}\}, \tag{42}$$

independently of subsequent renewal, by coupling the actual deaths above independent rate-$m_0$ deaths before the first replenishment is available; recovery after a crossing cannot undo the path event. At $K = 20$, $h_{\min} = 10$ and $e^{-m_0\tau} = 1/2$ this floor exceeds 0.41. A loaded pipeline changes the hypothesis and may remove the floor. Myelosuppression models already distinguish proliferating, maturation-transit and circulating compartments [5]; immediate logistic replenishment should not be equated with such recovery dynamics, and (42) is the precise price of that idealization.

## 8.4 Dimensionless meaning and empirical requirements

Table 2 records the main scales in units of nominal target division $b = .1$. A time-unit conversion leaves these ratios unchanged, so "renewal one" is a statement about a ratio, not evidence of a physiological timescale: the new renewal ratio is ten rather than the original 120, and it remains an uncalibrated assumption.

The most informative experiments are those that test a comparison assumption rather than sweep schedules: whether renewal remains above the needed floor under exposure, whether the delivered response actually enters the contraction band, and whether replenishment has a latency incompatible with the mission. Any one of these can invalidate the model faster, and more informatively, than a large dosing grid. Joint daughter-state observations matter for the same reason: matching mean inheritance is not sufficient, because a replacement kernel with the same $L$ can change extinction probabilities.

The intervention acts on a literal inherited-state population, while the reserve law is a birth–death mechanism with its own probabilistic origin. The capacity results are therefore not *caused* by epigenetic inheritance, and should not be described that way. The distinct contributions are the finite-reserve certificate, the source-specific inherited-population contraction, and their composition under one delivered input. Together these identify properties that a hypothetical intervention and preparation would need. They do not establish the safety of any available drug, prescribe a regimen, or predict a patient's outcome.

# 9 Evidence, scope and reproducibility

Different statements in this paper rest on different kinds of evidence, and it is worth separating them explicitly.

**Table 3:** Evidence supporting each class of statement. These classes describe how a claim was checked, not how biologically applicable it is: a mechanically checked inequality can be exactly correct about an unsuitable source.

| Class | Statements | Verification |
|---|---|---|
| Conventional proof | Process construction and nonexplosion; localization; scale function and excursion counting (Thm. 5, Cor. 6); order coupling; Riemann-sum capacity law (Thm. 8); fluid limit (Prop. 9); high-renewal limit (Prop. 10); latency floor (42) | Text of this paper |
| Exact finite computation | Rational drift rows and error envelope; reserve products; Taylor sums; Cor. 13 and Cor. 14 | Rational arithmetic, no floating point |
| Mechanized | Source interval inequalities, relative-error envelope, reserve product (38), delivery arithmetic and the assembled margins of Thm. 4 | Lean 4.30.0, pinned Mathlib |
| Numerical illustration | Figures 2, 3 and 4; the exact-probability curves of Fig. 1 | Deterministic solvers; never used in a proof |

The probability construction, stopping arguments, order comparison, capacity limit and both sharpness statements of Section 5 are conventional mathematical proofs, given in full above. Lean 4.30.0 checks the finite source-specific drift bounds over real erasure intervals, the relative-error envelope, the rational reserve product and the numerical exponential and risk margins that enter Theorem 4; the assembled declaration is `TherapeuticWindows.publication_certificate`, compiled against a pinned Mathlib with warnings treated as errors. This is a modular algebra verification. It does not formalize the process construction, the stopping theory, the coupling or the asymptotics, and it does not by itself establish the probabilistic meaning of the inequalities it checks; that meaning is supplied by the proofs in the text.

Two scope statements should be read alongside that certificate. First, Corollaries 13 and 14, and Propositions 9 and 10, are *not* part of the compiled declaration. The corollaries are exact rational evaluations of an already-proved inequality and are reproduced by the arithmetic script in the package; the propositions are conventional proofs. Second, no figure is ever used in place of an inequality: neither a sampled dosing grid nor a numerical healthy mean supplies an exclusion theorem, and the curve in Figure 2 is a numerically integrated killed forward equation for one parameter trajectory, not a rare-event Monte Carlo estimate. The unresolved safe-baseline converse of Section 8.2 is left unclassified rather than asserted in either direction.

**Code and data availability.** The source package contains nine Lean modules: `Source`, `Clock`, `Reserve`, `Robustness`, `Delivery` and `Root` for the benchmark result, together with `AnchoredReserve`, `WideBox` and `Publication`, which carry the anchored product, the relative-error envelope and the assembled declaration. It also contains the Python code that reconstructs the source and evaluates every rational quantity quoted above; the exact witness and figure data, including the high-precision evaluations behind Figure 1; build instructions with toolchain pins; and a reproduction driver. Figure 1 is computed at sixty decimal digits because the quantities plotted fall below $10^{-15}$, where double precision loses every significant figure to cancellation and can appear to contradict a proved bound. A persistent archive location and version identifier accompany the released version of this manuscript.

# A Finite certificates and source reconstruction

This appendix records the algebra used in the main proof, rather than the exploratory calculations that preceded it. In the state order of Section 2, the chemical action on a column vector is

$$\begin{aligned}
(Q_e x)_0 &= (x_1 - x_0)/50 + (x_3 - x_0)/50,\\
(Q_e x)_1 &= (51/100)(x_2 - x_1) + (1/100)(x_4 - x_1) + e(x_0 - x_1),\\
(Q_e x)_2 &= 2e(x_1 - x_2),\\
(Q_e x)_3 &= (51/100)(x_5 - x_3) + (1/100)(x_4 - x_3) + e(x_0 - x_3),\\
(Q_e x)_4 &= (e + 1/4)(x_1 + x_3 - 2x_4),\\
(Q_e x)_5 &= 2e(x_3 - x_5).
\end{aligned}$$

The two-daughter expectation from (4) is

$$Lx = \left(2x_0,\ x_0 + x_1,\ \frac{x_0 + x_2}{2} + x_1,\ x_0 + x_3,\ \frac{x_0 + x_4 + x_1 + x_3}{2},\ \frac{x_0 + x_5}{2} + x_3\right)^{\top},$$

so that $Lw = (28, 25, 23, 98, 76, 289/2)$ for the weight (30). Combining these with (6) gives the exact row ratios below. The last column is the bracket in (32) divided by $w_i$, evaluated at $e = .31$; it is maximal there because every rate entering it is nondecreasing in $e$.

| Type | $(A_{.01}w)_i/w_i$ | $(A_{.29}w)_i/w_i$ | $(A_{.31}w)_i/w_i$ | Error envelope / $w_i$ |
|---|---|---|---|---|
| UU | $-\frac{73}{700}$ | $-\frac{73}{700}$ | $-\frac{73}{700}$ | $\frac{353}{700}$ |
| UR | $-\frac{103}{550}$ | $-\frac{61}{550}$ | $-\frac{29}{275}$ | $\frac{323}{550}$ |
| RR | $-\frac{21}{125}$ | $-\frac{14}{125}$ | $-\frac{27}{250}$ | $\frac{123}{250}$ |
| AU | $\frac{559}{4200}$ | $-\frac{421}{4200}$ | $-\frac{491}{4200}$ | $\frac{451}{1050}$ |
| AR | $-\frac{363}{2150}$ | $-\frac{237}{2150}$ | $-\frac{114}{1075}$ | $\frac{1427}{1075}$ |
| AA | $\frac{111}{5350}$ | $-\frac{533}{5350}$ | $-\frac{579}{5350}$ | $\frac{477}{2675}$ |

Each displayed ratio is affine in $e$, so the endpoint values above certify the inequalities (31) and (32) on the whole intervals. Reading off column maxima,

$$\max_i \frac{(A_{.01}w)_i}{w_i} = \frac{559}{4200} < \frac{67}{500}, \qquad \max_i \frac{(A_{.29}w)_i}{w_i} = -\frac{533}{5350} < -\frac{99}{1000},$$

and the envelope maximum $1427/1075 < 27/20$ gives the perturbation term $27/2000$.

**Reserve inequality.** The certificate of Theorem 4 is the single finite rational product

$$\frac{278 \cdot 61 \cdot 120}{200} \prod_{j=0}^{77} \frac{122}{200-j} < \frac{7}{10^6},$$

whose decimal display is $6.9381403704212035 \times 10^{-6}$; the complete reduced rational value is supplied in the data package. Retaining the denominator $D_{278} = 2.5980909040121194\ldots$ gives the sharper value $2.6704763715\ldots \times 10^{-6}$ used in Corollary 13, whose initial scale terms are

$$\frac{1}{D_{278}} \sum_{j \ge h_0} P_j = \begin{cases} 9.8243173 \times 10^{-3}, & h_0 = 209, \\ 7.9237537 \times 10^{-5}, & h_0 = 221, \\ 2.3748386 \times 10^{-7}, & h_0 = 240, \end{cases}$$

and whose $h_0 = 208$ value exceeds $10^{-2}$. For Corollary 14, $D_{161} = 2.6271217287\ldots$ at $K = 232$ and the sharpened bound is $9.7801264\ldots \times 10^{-3}$, against $2.5693583\ldots \times 10^{-2}$ for the cruder product.

**Exponential inequalities.** The two exponential bounds used for the delivered wide-box witness follow from the positive partial sums

$$\sum_{j=0}^{15} \frac{(99/25)^j}{j!} > 50, \qquad \sum_{j=0}^{30} \frac{(43/5)^j}{j!} > 5000.$$

All terms are positive and each sum is less than the corresponding exponential. The exponent identity and marginal assembly are

$$108\left(\frac{99}{1000} - \frac{27}{2000}\right) - 4\left(\frac{67}{500} + \frac{27}{2000}\right) = \frac{2161}{250}, \qquad 1 - \frac{107}{12500} - \frac{7}{10^6} = \frac{991433}{10^6}.$$

**The benchmark rows.** The narrow-box comparison in Table 1 uses $A_e w \le -.09w$ on $[.29, .31]$, $A_e w \le .14w$ on $[.01, .31]$, and weighted perturbation at most $.001w$. Its net exponent is $108(.089) - 4(.141) = 9.048 > 9$, while $e^9 > 8000$ by the positive Taylor sum through degree 25; its target upper bound is therefore $(214/5)/8000 = .00535$. Its healthy bounds come directly from the top-anchor product at $K = M = 20$,

$$20 \cdot \frac{61}{100} \cdot 120 \cdot \frac{(20(61/100)/r_{\mathrm{eff}})^{10}}{10!} \quad \text{at } r_{\mathrm{eff}} = 12 \text{ and } r_{\mathrm{eff}} = 9,$$

which give $4.7595 \times 10^{-4}$ and $8.4518 \times 10^{-3}$ respectively. By Proposition 10 this top-anchor expression is asymptotically exact in $r_{\mathrm{eff}}$, so the factor $(12/9)^{10} \approx 17.8$ separating the two rows is not an artifact of the bound: at $d = 10$ the benchmark genuinely depends on renewal through its tenth power, which is why so much renewal is needed when the reserve carries no spare capacity.

# References


[1] Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008. doi: 10.1016/j.automatica.2008.03.027.

[2] Krishna B. Athreya and Peter E. Ney. *Branching Processes*. Springer, Berlin, 1972. doi: 10.1007/978-3-642-65371-1.

[3] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986. doi: 10.1002/9780470316658.

[4] Eric Foxall. Extinction time of the logistic process. *Journal of Applied Probability*, 58(3): 637–676, 2021. doi: 10.1017/jpr.2020.112. URL https://arxiv.org/abs/1805.08339.

[5] Lena E. Friberg, Anja Henningsson, Hugo Maas, Laurent Nguyen, and Mats O. Karlsson. Model of chemotherapy-induced myelosuppression with parameter consistency across drugs. *Journal of Clinical Oncology*, 20(24):4713–4721, 2002. doi: 10.1200/JCO.2002.02.140. URL https://pubmed.ncbi.nlm.nih.gov/12488418/.

[6] Einar Bjarki Gunnarsson, Subhajyoti De, Kevin Leder, and Jasmine Foo. Understanding the role of phenotypic switching in cancer drug resistance. *Journal of Theoretical Biology*, 490: 110162, 2020. doi: 10.1016/j.jtbi.2020.110162. URL https://arxiv.org/abs/2001.07690.

[7] Einar Bjarki Gunnarsson, Benedikt Vilji Magnússon, and Jasmine Foo. Optimal dosing of anti-cancer treatment under drug-induced plasticity, 2025. URL https://arxiv.org/abs/2412.16391v2. arXiv:2412.16391, version 2.

[8] Samuel Karlin and Howard M. Taylor. *A First Course in Stochastic Processes*. Academic Press, New York, 2 edition, 1975.

[9] Thomas G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970. doi: 10.2307/3212147.

[10] Camille Pouchol, Jean Clairambault, Alexander Lorz, and Emmanuel Trélat. Asymptotic analysis and optimal control of an integro-differential system modelling healthy and cancer cells exposed to chemotherapy, 2016. URL https://arxiv.org/abs/1612.04698. arXiv:1612.04698.

[11] Sydney M. Shaffer, Margaret C. Dunagin, Stefan R. Torborg, Eduardo A. Torre, Benjamin Emert, Clemens Krepler, Marilda Beqiri, Katrin Sproesser, Patricia A. Brafford, Min Xiao, Elliott Eggan, Ioannis N. Anastopoulos, Cesar A. Vargas-Garcia, Abhyudai Singh, Katherine L. Nathanson, Meenhard Herlyn, and Arjun Raj. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658):431–435, 2017. doi: 10.1038/nature22794. URL https://pubmed.ncbi.nlm.nih.gov/28607484/.

[12] Sreenath V. Sharma, Diana Y. Lee, Bihua Li, Margaret P. Quinlan, Fumiyuki Takahashi, Shyamala Maheswaran, Ultan McDermott, Nancy Azizian, Lee Zou, Michael A. Fischbach, Kwok-Kin Wong, Kathleyn Brandstetter, Ben Wittner, Sridhar Ramaswamy, Marie Classon, and Jeff Settleman. A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1):69–80, 2010. doi: 10.1016/j.cell.2010.02.027. URL https://pubmed.ncbi.nlm.nih.gov/20371346/.

[13] Sean Summers and John Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010. doi: 10.1016/j.automatica.2010.08.006.