Reliable predictions from small cell populations: preparation dependence, measurement and certified count bounds

17 September 2026

Abstract

Prediction from small cell populations depends on how the population was prepared and how it is observed, not only on the dynamics of individual lineages. We study upper prediction regions $\{0, \dots, k\}$ for the cell count at a fixed horizon in two-type branching models with inherited states, and ask which observations and assumptions justify a given cutoff k. First, we construct two founder preparations, one drawing a growth class independently for each founder and one giving every founder in a well the same class, that have identical observation laws for every single-founder experiment yet different minimal 95% two-founder cutoffs, six and seven. Ordered-endpoint error accounting yields finite-sample marginal coverage for a data-selected cutoff without independence between calibration and the future population; an exact binomial calculation attains coverage at least 245/256 from 640 independent calibration wells, and a continuous preparation family identifies $\rho = 3/5$ as the dependence level at which the minimal cutoff changes. Second, independent cell detection can erase or reverse the signal of a low-count statistic at efficiency one half, while a bounded generating-function statistic separates the preparations for every positive efficiency; thinning-invariant moment identities identify an unknown constant efficiency together with the preparation, and a feasible-set rule with bounded statistics keeps the coverage guarantee when the efficiency is unresolved. Third, one-sided stochastic comparisons give uniform recorded-count bounds under demographic, founder and false-object uncertainty: recorded count at most four with probability at least 0.9503 for all birth rates up to 0.105 and all death rates at least 0.05 per time unit at horizon seven, and at most five with probability at least 0.9567 for arbitrary death and switching rates when births are at most 0.1. Fourth, for a reference inherited-state rate class with per-type ℓ^1 uncertainty, a finite successful-history certificate proves recorded count at most three with probability at least 0.955935, with three the smallest endpoint valid over the class. We separate theorems checked in Lean 4 from conventional extensions and exact computations, and report a documented lineage-data analysis that does not establish biological coverage. The results give explicit conditions for trustworthy count predictions and for measurements that justify smaller prediction regions.

1 Introduction

1.1 Prediction before mechanism

A low-founder proliferation assay places a small number of cells in a well, waits a fixed time T, and records a count. A prediction region for that count is a set $\{0, \dots, k\}$ together with a guarantee $\mathbb{P}(N_T \le k) \ge 0.95$, where N_T is the true terminal count and the probability is over the preparation, the dynamics and, when a rule is learned from data, the calibration data. Such a region is useful precisely when it does *not* require the experimenter to identify every

hidden state or transition rate: a smaller region signals unusual outgrowth, a valid region protects against false alarms.

This article is about the conditions under which such a region can be trusted when the preparation, the hidden cellular states and the detector may differ from the setting in which the model was fitted. The thread has four parts.

(i) A model that reproduces the expected count, and even the expected birth and death fluxes, can assign the wrong probability to small and large counts when phenotypic states are inherited. This is the subject of a companion manuscript [8]; we do not repeat it, but the present results start from the same observation that averages do not determine tails.

(ii) Even a correct single-founder count law need not transfer to a different founder preparation. We give two preparations with *identical* observation laws for every single-founder experiment and different minimal 95% two-founder cutoffs (Section 3). Increasing the number of founders does not remove the discrepancy.

(iii) The measurement that repairs this must address the dependence relevant to the decision. Exact binomial and generating-function calculations show which two-founder statistics identify the preparation, how many wells are sufficient for a stated coverage, what a common detection efficiency does to the signal, and how to proceed when the efficiency is unknown (Sections 4 and 5).

(iv) Some upper-count guarantees can be proved without resolving the mechanism at all, by one-sided comparison with a scalar birth–death process or with a finite family of successful histories (Sections 6 and 7).

1.2 Contributions

The main results are stated for two-type continuous-time branching processes with inherited type, and the constants are exact rationals unless marked as displays.

- **Exact preparation ambiguity** (Theorem 3.1). Preparations I and W have equal single-family observation laws and two-founder cumulative distribution functions $F_I(k) = 1 - (k+5)/(4 \cdot 2^k)$ and $F_W(k) = 1 - (k+1)/2^{k+1}$; their minimal 95% cutoffs are six and seven. For m founders the variances are $\frac{5}{4}m$ and $m + \frac{1}{4}m^2$, and the independent-preparation 95% quantile has coverage tending to $1/2$ under W (Proposition 3.3). A continuous family $P_\rho = (1-\rho)P_I + \rho P_W$ has minimal cutoff six exactly when $\rho \le 3/5$ (Proposition 3.4).

- **Direction-aware calibration** (Lemma 4.1, Theorem 4.2, Proposition 4.4). For any rule with values in $\{6, 7\}$, coverage under I is at least $245/256$ and under W at least $31/32$ minus the probability of selecting six, with no assumption relating calibration and the future well. The count-≤ 2 rule with 4000, 1280 or 1024 wells and the counts-3-or-4 rule with 640 or 576 wells give uniform coverage $245/256$ or $19/20$, with exact integer certificates. A confidence rule for the dependence parameter certifies $\rho \le 3/5$ with uniform coverage (Proposition 4.5).

- **The detector is part of the measurement** (Propositions 5.1–5.4). Under independent thinning with efficiency d, the count-≤ 2 gap is $d^3(2d-1)/(1+d)^4$, zero at $d = 1/2$; the statistic z^Y has gap $u^2(1-u)^2/(4(1+u)^2)$ with $u = d(1-z)$, positive for all $d > 0$. The pair (efficiency, preparation) is identified by the first two factorial moments; a bounded feasible-set rule retains the coverage guarantee with unknown d; complete erasure is an obstruction.

- **Uniform safety without resolving switching** (Theorem 6.1, Propositions 6.2–6.6). A recorded count at most four has probability at least 0.952037 over a demographic box

at horizon seven with arbitrary switching, arbitrary deletion, multiple-founder mass at most 0.001 and false-object probability at most 0.01; the proof uses only upper birth and lower death bounds, and pathwise domination extends it to births at most 0.105 (cutoff four, bound above 0.9503) and to arbitrary deaths with births at most 0.1 (cutoff five, bound above 0.9567).

- **Robust inherited-state prediction** (Theorem 7.1). Over a reference rate class with per-type ℓ^1 radius 0.0005 per day, founder resistant probability within 0.001 of 5/24, horizon $\log 2/0.10001$ days, and the same preparation and false-object budgets, the recorded count is at most three with probability at least $647025972840560666 1/676851426700000000 = 0.955935\ldots$, and three is the smallest endpoint valid over the class. Endpoint four admits a 1% inference-failure allowance.

Each claim is labelled by its evidence level (Section 9): "Lean-checked" means that the statement, at the stated scope, is a compiled theorem in the accompanying Lean 4 development; "conventional" means a written proof in this article, with exact rational arithmetic replayed by a script that regenerates every printed constant and figure.

1.3 Relation to existing work

Browning et al. [6] study the practical identifiability of phenotypic adaptation from low-cell-count experiments with a stochastic model and several observation designs, and document the difficulty of distinguishing heterogeneity from alternative phenotype structures; see also Browning et al. [5] for identifiability of stochastic models in systems biology and Iyer et al. [12] for the role of inheritable states in drug-persister correlations. Our question is narrower and complementary: can two *exactly* equivalent single-founder observation laws imply different count cutoffs after a preparation change, and which additional observation supports the appropriate cutoff? The phenomena motivating inherited-state models, persistence and drug tolerance without genetic change, are documented in Balaban et al. [4], Shaffer et al. [22], Sharma et al. [23]; the classical Luria–Delbrück argument that fluctuations reveal mechanism goes back to Luria and Delbrück [18].

The mathematical components are established: branching generating functions and the linear birth–death law [2, 10, 13, 14, 26], Hoeffding's inequality [11], exact binomial confidence limits [7], stochastic domination by coupling [17], finite-state truncation of master equations [20], and asymmetric error control in ordered decisions [16]. Distribution-free prediction [1, 25] supplies marginal guarantees from exchangeable calibration data; hierarchical data need attention to the sampling unit [15]. Those methods do not by themselves justify transfer between founder preparations or a latent-count bound under an unspecified detector, and conversely the model-specific theorems here are not distribution-free. Our contribution is the explicit source-to-decision combination, the exact examples, the improved calibration and detector results, the broad one-sided demographic envelopes, and the verification boundary; we do not claim a general identifiability theorem, an optimal test, or biological validation.

1.4 Evidence conventions

Two kinds of evidence appear. *Lean-checked* theorems are compiled in Lean 4 [9] against mathlib [24] with warnings treated as errors, and are named in the text by their declaration; they concern the encoded source process and the stated hypotheses, not the biological adequacy of either. *Conventional* results have complete written proofs here; exact rational inequalities in them are additionally replayed by the accompanying script. Decimal values are downward-rounded displays of exact rationals unless they are marked numerical. A file hash identifies bytes, not correctness.

2 Source model, prediction targets and contracts

2.1 The branching source

The latent state is $X_t = (S_t, R_t) \in \mathbb{N}^2$ with total count $N_t = S_t + R_t$. A cell of type $a \in \{S, R\}$ divides into two cells of its type at rate b_a, dies at rate d_a, and switches type at rate q_a; the six rates are finite and nonnegative, with units inverse to the time unit. Cells evolve independently conditional on their initial states and on a possibly shared environment. Division increases the count by one, death decreases it by one, switching leaves it unchanged. Births are dominated by a Yule process of rate $\max(b_S, b_R)$ per cell and the remaining event rates are bounded on bounded population sets, so the process is nonexplosive [21]. The labels S and R are model types; nothing below requires a biological interpretation of them.

2.2 Targets

Let N_T be the true terminal count and let the recorded count be $Y_T = D + A$ with $D \le N_T$ the number of genuine cells detected and $A \in \mathbb{N}$ the number of false objects. An exceedance of cutoff k begins at $k + 1$. For a fixed source law P the *minimal* 95% *cutoff* is $k_{.95}(P) = \min\{k \in \mathbb{N} : P(N_T \le k) \ge 19/20\}$. For a data-selected cutoff $K = K(C)$ from calibration data C, *marginal coverage* means $\mathbb{P}(N_T \le K) \ge 19/20$ with the probability including C, and *certification of the model-specific cutoff* means $\mathbb{P}(K = k_{.95}) \ge 1 - \beta$. Neither implies coverage conditional on the realized value of K; that is a separate statement (Section 4.4).

2.3 Three contracts

The theorems use three source contracts, which must not be interchanged.

Contract 2.1 (Menu calibration)**.** Exactly two founders per well; the source is one of two specified preparations I and W; normalized horizon $T = \log 2$ in units of the fast birth rate; calibration wells are independent with perfect terminal counts; the future well has the same source marginal but may depend on the calibration wells.

Contract 2.2 (Uniform demographic safety)**.** Preparation is empty, single-founder or multiple-founder with multiple-founder probability at most γ; horizon seven time units; per-cell rates in a stated set; arbitrary finite switching; detection may delete genuine cells arbitrarily; false objects appear with probability at most η.

Contract 2.3 (Inherited-state rate class)**.** A reference rate vector with per-type ℓ^1 uncertainty radius 0.0005 per day, founder resistant probability within 0.001 of 5/24, a possibly shared static environment, physical horizon $T_0 = \log 2/0.10001 \approx 6.93078$ days, the same preparation and false-object budgets as Contract 2.2, and arbitrary deletion.

The two-founder example is normalized: assigning fast birth rate one per day makes its horizon about 0.693 days; assigning 0.1 per day makes it about 6.93 days. Neither choice establishes biological calibration, and "horizon seven" in Contract 2.2 means seven time units. The cutoffs three, four, six and seven proved below belong to different contracts and cannot be compared without stating the contract.

3 Exact preparation ambiguity

3.1 Construction

At the reference rates $(b_S, d_S, q_S; b_R, d_R, q_R) = (0, 0, 0; 1, 0, 0)$ a *slow* family stays at one cell and a *fast* family is a Yule process of rate one per cell. Its count Z at $T = \log 2$ satisfies

$$\mathbb{P}(Z = j) = 2^{-j} \quad (j \geq 1), \qquad g_0(z) = z, \qquad g_1(z) = \frac{z}{2 - z}, \tag{1}$$

where g_0, g_1 are the slow and fast terminal generating functions; indeed the forward equations $p_j' = (j-1)p_{j-1} - jp_j$ from one cell are solved by $p_j(t) = e^{-t}(1 - e^{-t})^{j-1}$.

In preparation I each founder independently receives the slow or fast class with probability $1/2$ each. In preparation W one class is drawn for the well and given to every founder. Conditional on the classes, families evolve independently in both preparations. "Shared class" names a dependence in the preparation; it does not assert that a persistent environmental cause has been identified.

Theorem 3.1 (Equal single-founder observations, unequal predictions)**.** *For any measurable observation of a single family's chronological trajectory, preparations I and W have equal observation laws. For two founders the terminal count distribution functions are, for $k \geq 2$,*

$$F_I(k) = 1 - \frac{k+5}{4 \cdot 2^k}, \qquad F_W(k) = 1 - \frac{k+1}{2^{k+1}}, \tag{2}$$

and the minimal 95% cutoffs are $k_{.95}(P_I) = 6$ and $k_{.95}(P_W) = 7$.

Proof. Let P_0, P_1 be the conditional family laws on a common measurable family space and $M = (P_0 + P_1)/2$. The two-family laws are

$$P_I = M \otimes M, \qquad P_W = \tfrac{1}{2}P_0 \otimes P_0 + \tfrac{1}{2}P_1 \otimes P_1.$$

Both first-family marginals equal M, so every measurable observation of a single family has the same law under I and W. This is an identity of laws, not merely of moments, and it applies to the full chronological state, event and waiting-time trace.

For two founders, the initial-class mixtures are

$$I: \ \tfrac{1}{4}(\text{slow}, \text{slow}) + \tfrac{1}{2}(\text{slow}, \text{fast}) + \tfrac{1}{4}(\text{fast}, \text{fast}), \qquad W: \ \tfrac{1}{2}(\text{slow}, \text{slow}) + \tfrac{1}{2}(\text{fast}, \text{fast}),$$

so the terminal generating functions are

$$G_I(z) = \frac{(g_0(z) + g_1(z))^2}{4}, \qquad G_W(z) = \frac{g_0(z)^2 + g_1(z)^2}{2}. \tag{3}$$

The all-slow count is two. The mixed count is $1 + Z$, with $\mathbb{P}(1 + Z \leq k) = 1 - 2^{-(k-1)}$. The all-fast count $Z_1 + Z_2$ has the negative binomial law $\mathbb{P}(Z_1 + Z_2 = j) = (j-1)2^{-j}$ for $j \geq 2$, and writing $j = k + i$,

$$\sum_{j>k} (j-1)2^{-j} = 2^{-k}\Big[(k-1)\sum_{i\geq 1} 2^{-i} + \sum_{i \geq 1} i\, 2^{-i}\Big] = (k+1)2^{-k}.$$

Hence $F_I(k) = \frac{1}{4} + \frac{1}{2}(1 - 2^{-(k-1)}) + \frac{1}{4}(1 - (k+1)2^{-k}) = 1 - (k+5)/(4 \cdot 2^k)$ and $F_W(k) = \frac{1}{2} + \frac{1}{2}(1 - (k+1)2^{-k}) = 1 - (k+1)/2^{k+1}$. In particular

$$F_I(5) = \tfrac{59}{64} < \tfrac{19}{20} \leq F_I(6) = \tfrac{245}{256}, \qquad F_W(6) = \tfrac{121}{128} < \tfrac{19}{20} \leq F_W(7) = \tfrac{31}{32},$$

and monotonicity of a distribution function excludes every smaller cutoff. □

Table 1: Exact two-founder model probabilities $\mathbb{P}(N_T \le k)$ at $T = \log 2$. These are model calculations, not empirical observations.

Cutoff k	Independent classes (I): $F_I(k)$		Shared class (W): $F_W(k)$	
2	9/16	$= 0.5625$	5/8	$= 0.625$
5	59/64	$= 0.921875$	29/32	$= 0.90625$
6	245/256	$= 0.95703125$	121/128	$= 0.9453125$
7	125/128	$= 0.9765625$	31/32	$= 0.96875$

Remark 3.2 (What is Lean-checked). The single-family identity is compiled for the implemented chronological trace as `chronological_single_observation_equal`, and the four inequalities defining the two minimal cutoffs as `chronological_minimal_endpoints`, both in `YuleAmbiguity.lean`. The formal proof does not assume the generating functions as an oracle: pure-birth paths keep the slow count fixed and never decrease the fast count, so killing paths once the fast count exceeds seven loses no successful history for the required thresholds, and a finite backward equation with exponential-sum solutions is checked by integer coefficient identities (`yule_finite_exact`, `source_cdf_formula`). The general-k formula (2) is conventional. A genealogical-tree sample space is not separately implemented; the abstract mixture identity applies to it whenever it is supplied.

This is a *preparation-transfer obstruction*: the single-founder laws agree, and the target uses two founders under different dependence rules. It is not a contradiction in which one law has two quantiles.

3.2 More founders do not remove shared uncertainty

Proposition 3.3. *For m founders both preparations have mean $3m/2$, while*

$$\mathrm{Var}_I(N_m) = \tfrac{5}{4}m, \qquad \mathrm{Var}_W(N_m) = m + \tfrac{1}{4}m^2. \tag{4}$$

Let q_m be the minimal 95% cutoff of N_m under I. Then $q_m/m \to 3/2$ and $P_W(N_m \le q_m) \to 1/2$.

Proof. The one-founder mixture law M has mean $\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 2 = \frac{3}{2}$ and second moment $\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 6 = \frac{7}{2}$, hence variance $\frac{5}{4}$; under I the m families are independent, which gives the first formula. Under W the count is m with probability $\frac{1}{2}$ and a sum of m independent geometric($\frac{1}{2}$) variables (mean $2m$, variance $2m$) with probability $\frac{1}{2}$, so $\mathbb{E}N_m^2 = \frac{1}{2}m^2 + \frac{1}{2}(2m + 4m^2)$ and $\mathrm{Var}_W(N_m) = m + \frac{1}{4}m^2$. Under I the strong law gives $N_m/m \to \frac{3}{2}$, so for every $\varepsilon > 0$, $P_I(N_m \le (\frac{3}{2} + \varepsilon)m) \to 1$ and $P_I(N_m \le (\frac{3}{2} - \varepsilon)m) \to 0$; hence $q_m/m \to \frac{3}{2}$. Under W, the slow branch has $N_m = m \le q_m$ eventually, while the fast branch has $N_m/m \to 2 > \frac{3}{2}$, so $P_W(N_m \le q_m) \to \frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 0$. □

The law of large numbers operates correctly within each conditional population; the shared draw remains. This is persistent shared-condition uncertainty, not a failure of a limit theorem, and it is not evidence about any patient outcome.

3.3 A continuous dependence family

Proposition 3.4. *For $0 \le \rho \le 1$ let $P_\rho = (1-\rho)P_I + \rho P_W$: with probability ρ the well uses a shared class draw, otherwise independent draws. Every single-founder observation law is unchanged, the founders' fast-class indicators have correlation ρ, and*

$$F_\rho(6) = \frac{245 - 3\rho}{256}, \qquad F_\rho(7) = \frac{250 - 2\rho}{256}, \qquad P_\rho(N_T \in \{3,4\}) = \frac{19 - 5\rho}{64}. \tag{5}$$

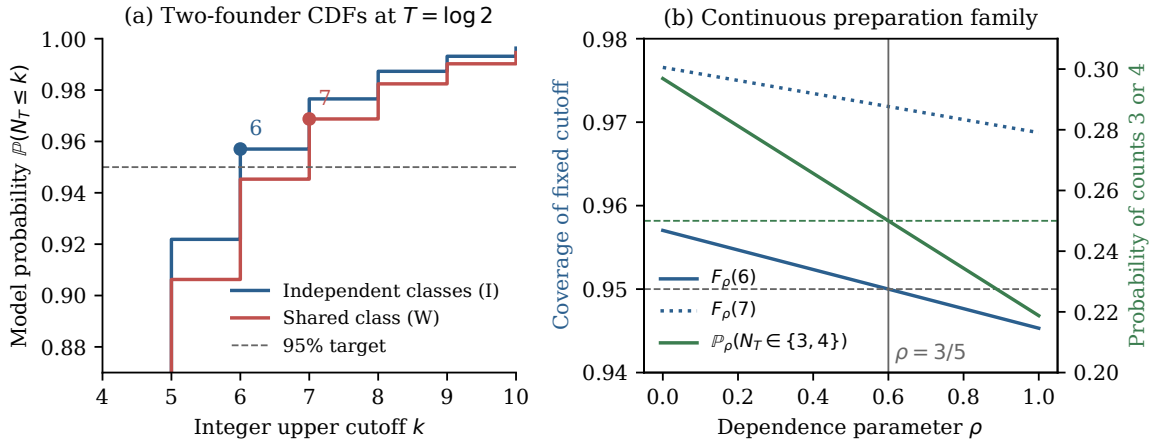

Figure 1: (a) Exact two-founder distribution functions of Theorem 3.1; the marked points are the minimal 95% cutoffs six and seven. (b) The continuous family of Proposition 3.4: the coverage of the fixed cutoff six crosses 19/20 and the probability of counts three or four crosses 1/4 at the same $\rho = 3/5$. All values are model probabilities.

The minimal 95% *cutoff of* P_ρ *is six if* $\rho \le 3/5$ *and seven if* $\rho > 3/5$*; equivalently, it is six exactly when* $P_\rho(N_T \in \{3,4\}) \ge 1/4$.

Proof. Linearity of the mixture gives (5) from Table 1 and from $P_I(\{3,4\}) = 19/64$, $P_W(\{3,4\}) = 7/32$ (Lemma 4.3 below). $F_\rho(6) \ge 19/20$ holds exactly when $3\rho \le 245 - 243.2$, that is $\rho \le 3/5$; $F_\rho(5) = (1-\rho)\frac{59}{64} + \rho\frac{29}{32} < 19/20$ for every ρ; and $F_\rho(7) \ge 31/32 > 19/20$ for every ρ. Finally $(19 - 5\rho)/64 \ge 1/4$ is again $\rho \le 3/5$. The single-founder marginal fast probability is 1/2 in both components, and the joint indicator law is $(1-\rho)$ times the product law plus ρ times the diagonal law, which gives correlation ρ. □

Figure 1 shows the two distribution functions and the continuous family. The observation criterion $P_\rho(N_T \in \{3,4\}) \ge 1/4$ turns the binary ambiguity into a quantitative question: the measurement need only establish that dependence lies below the level that changes the decision, not identify a mechanism.

4 Information for the decision

Throughout this section Contract 2.1 holds. Let $N_1, \dots, N_n$ be the perfectly counted terminal counts of n independent two-founder calibration wells, and let N_T be the future two-founder count, with the same source marginal as the calibration wells but otherwise arbitrary dependence on them.

4.1 Ordered-endpoint accounting

Lemma 4.1 (Nested endpoints)**.** *For any rule* K *with values in* $\{6,7\}$*, measurable with respect to the calibration data, and any future count with the declared source marginal,*

$$P_I(N_T \le K) \ge F_I(6) = 245/256, \tag{6}$$

$$P_W(N_T \le K) \ge \max\{F_W(6),\, F_W(7) - P_W(K=6)\} = \max\{121/128,\, 31/32 - e_W\}, \tag{7}$$

where $e_W = P_W(K=6)$*. More precisely, under either source,* $\mathbb{P}(N_T \le K) = F(7) - \mathbb{P}(N_T = 7, K = 6)$.

Proof. Since $K \geq 6$, $\{N_T \leq 6\} \subseteq \{N_T \leq K\}$, giving both bounds at six. Since $\{N_T \leq 7\} = \{N_T \leq K\} \cup \{N_T = 7, K = 6\}$ disjointly, $\mathbb{P}(N_T \leq K) = F(7) - \mathbb{P}(N_T = 7, K = 6) \geq F(7) - \mathbb{P}(K = 6)$. Insert the source values from Table 1. □

No independence between calibration and the future count enters. An error under I enlarges the region and cannot damage coverage; the dangerous decision is selecting six when W is true. Given only e_W, the loss $\mathbb{P}_W(N_T = 7, K = 6)$ is at most $\min\{e_W, P_W(N_T = 7)\} = \min\{e_W, 3/128\}$, and this extremal overlap is attainable over unrestricted couplings, so (7) is the best bound of its kind. Standard testing permits asymmetric error control [16]; the point here is that the ordering of the endpoints makes one error harmless for coverage. The inclusions are compiled as `nested_endpoint_lower` and `upper_endpoint_loss` in `Postproof.lean`.

4.2 The low-count rule

Let $X_i = \mathbf{1}\{N_i \leq 2\}$ and $C_n = \sum_i X_i$. Theorem 3.1 gives success probabilities $p_I = 9/16$ and $p_W = 5/8$, with midpoint $19/32$. Define

$$K_n = \begin{cases} 6, & C_n/n < 19/32, \\ 7, & C_n/n \geq 19/32. \end{cases} \tag{8}$$

Theorem 4.2 (Adaptive coverage)**.** *Rule (8) has uniform marginal coverage at least* $245/256 = 0.95703125$ *under both preparations in each of the following cases: (i)* $n = 4000$*, with the concentration proof compiled in Lean; (ii)* $n = 1280$*, with the exact binomial certificate below. At* $n = 1024$ *the same argument gives uniform coverage at least* $19/20$*. The future count need not be independent of calibration.*

Proof. (i) The threshold is $1/32$ from either mean. Hoeffding's one-sided inequality for independent $[0, 1]$ variables [11] gives $e_W, e_I \leq \exp(-2n/32^2) = \exp(-125/16) \leq 1/2000$, the last step because the Taylor polynomial of degree 20 of $\exp(125/16)$ already exceeds 2000. Then $\min\{245/256, 31/32 - 1/2000\} = 245/256$ and Lemma 4.1 applies. This is compiled as `adaptive_prediction_sharp`, which imports the source-connected concentration theorem `classification_error`.

(ii) By independence and the Bernoulli means, $C_n \sim \text{Binomial}(n, p)$ under either source: each subset of j successful wells has probability $p^j(1-p)^{n-j}$ and the $\binom{n}{j}$ such events partition $\{C_n = j\}$. Hence

$$e_W(n) = P_W(C_n \leq \lceil 19n/32 \rceil - 1) = 8^{-n} \sum_{j=0}^{\lceil 19n/32 \rceil - 1} \binom{n}{j} 5^j 3^{n-j}. \tag{9}$$

The integer inequalities of Appendix A give $e_W(1024) \leq 3/160$ and $e_W(1280) \leq 3/256$. Since $31/32 - 3/160 = 19/20$ and $31/32 - 3/256 = 245/256$, Lemma 4.1 completes the proof. □

The first three rows of Table 2 show the exact errors. These are sufficient sample sizes for this rule and this menu, not minimal ones, and adjacent sizes can behave nonmonotonically because the integer threshold jumps. The 1280-well rule does not have 99.95% label accuracy: its two errors are about 1.0% and 1.3%. The conservative 4000-well concentration bound supplies the stronger label guarantee. Coverage and label accuracy are different targets.

4.3 A better separating event

Lemma 4.3. *For* $j \geq 3$ *the two-founder mass functions are* $p_I(j) = (j+3)/2^{j+2}$ *and* $p_W(j) = (j-1)/2^{j+1}$*, so* $p_I(j) - p_W(j) = (5-j)/(4 \cdot 2^j)$*; also* $p_I(2) = 9/16$*,* $p_W(2) = 5/8$*. Consequently the event* $\{3, 4\}$ *attains the total variation distance:*

$$P_I(N_T \in \{3,4\}) = \frac{19}{64}, \qquad P_W(N_T \in \{3,4\}) = \frac{7}{32}, \qquad \text{TV}(P_I, P_W) = \frac{5}{64}. \tag{10}$$

Table 2: Calibration errors of the two rules from exact binomial sums. Decimals are downward-rounded displays; the certificates are integer inequalities. e_W is the probability of selecting six under W, e_I the probability of selecting seven under I. The last column is the uniform coverage lower bound from Lemma 4.1.

n	Event	Selects six when	e_W	e_I	Uniform bound
1024	$\{N \le 2\}$	$C_n \le 607$	0.01833462	0.02336813	19/20
1280	$\{N \le 2\}$	$C_n \le 759$	0.00994686	0.01284612	245/256
4000	$\{N \le 2\}$	$C_n \le 2374$	0.00002262	0.00003456	245/256
576	$\{3, 4\}$	$H_n \ge 149$	0.01280508	0.01905814	0.95594491
640	$\{3, 4\}$	$H_n \ge 165$	0.01055929	0.01284164	245/256

Proof. Difference (2) at k and $k-1$. The mass differences are positive exactly at $j \in \{3, 4\}$, negative at $j = 2$, zero at $j = 5$ and negative beyond, so $\{3, 4\}$ is the set of positive differences and its probability gap is the total variation distance. □

The count-≤ 2 event has gap $1/16 = 4/64$; the event $\{3, 4\}$ has the largest gap $5/64$ among all events. Define

$$H_n = \sum_{i=1}^{n} \mathbf{1}\{N_i \in \{3, 4\}\}, \qquad K'_n = \begin{cases} 6, & H_n \ge \lceil 33n/128 \rceil, \\ 7, & \text{otherwise}, \end{cases} \tag{11}$$

where 33/128 is the midpoint of 7/32 and 19/64. The orientation is reversed relative to (8): a high frequency of wells with three or four cells supports I.

Proposition 4.4 (Halving the sufficient calibration size)**.** *Under Contract 2.1, rule* (11) *with* $n = 640$ *independent calibration wells has uniform marginal coverage at least* 245/256*, and with* $n = 576$ *at least* 19/20*, without independence between calibration and the future count.*

Proof. Under W, $H_n \sim \text{Binomial}(n, 7/32)$ and $e_W = P_W(H_n \ge \lceil 33n/128 \rceil)$. With $B_{n,c} = \sum_{j \ge c} \binom{n}{j} 7^j 25^{n-j}$, Appendix A verifies the integer inequalities

$$160\, B_{576,149} \le 3 \cdot 32^{576}, \qquad 256\, B_{640,165} \le 3 \cdot 32^{640}, \tag{12}$$

that is $e_W(576) \le 3/160$ and $e_W(640) \le 3/256$; here $\lceil 33 \cdot 576/128 \rceil = 149$ and $\lceil 33 \cdot 640/128 \rceil = 165$. Lemma 4.1 gives the two bounds; the I bound 245/256 holds for every rule with values in $\{6, 7\}$. □

The new rule halves the sufficient size of Theorem 4.2(ii) at nearly the same I error (Table 2, Figure 2). It requires the full counts rather than the bit $\mathbf{1}\{N_i \le 2\}$. The event and threshold are fixed analytically before any data are seen; selecting among several such rules on the calibration sample would invalidate an unadjusted certificate. The largest single-event gap does not make the repeated Bernoulli test optimal among all procedures using complete counts: the likelihood ratio is $p_I(2)/p_W(2) = 9/10$ and $p_I(j)/p_W(j) = (j+3)/(2(j-1))$ for $j \ge 3$, and a likelihood-ratio or sequential procedure could be more efficient, but it needs its own exact error calculation and stopping rule; no optimality claim is made. A worked input: among 640 qualifying wells, $H = 164$ selects seven and $H = 165$ selects six.

4.4 What coverage alone does not identify

Always returning seven is uniformly safe without calibration, since $F_I(7) = 125/128$ and $F_W(7) = 31/32$. Calibration therefore determines whether a smaller fixed cutoff is supported; it is not needed for the existence of a safe cutoff.

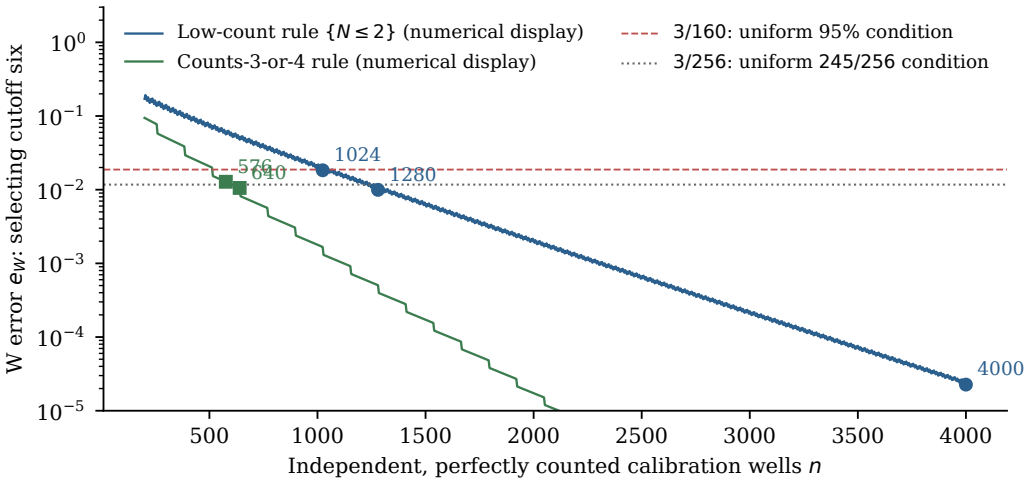

Figure 2: The W error controls the coverage guarantee. Curves are numerical renderings of the binomial tails for display; the marked sample sizes are the exact integer certificates of Table 2. The horizontal levels are sufficient conditions, not lower bounds on necessary sample size.

An independent external randomizer returning seven with probability q and six otherwise has W coverage $(1-q)\frac{121}{128} + q\frac{31}{32}$, which is at least 19/20 exactly when $q \geq 1/5$; at $q = 1/5$ the mean endpoint is 6.2, the W coverage is exactly 19/20 and the I coverage is 123/128. This rule learns nothing about the model. Its coverage conditional on returning six under W is only 121/128. A calibrated rule is therefore not claimed to be necessary or optimal for unconditional coverage with a small mean endpoint; its value is a data-based certification of the model-specific minimum with explicit error probabilities.

If the future count is additionally independent of calibration, conditioning on K gives the exact coverages

$$P_I(N_T \leq K) = \frac{245}{256} + \frac{5e_I}{256}, \qquad P_W(N_T \leq K) = \frac{31}{32} - \frac{3e_W}{128}. \tag{13}$$

The second expression exceeds 19/20 as soon as $e_W \leq 4/5$. The large difference between this and the arbitrary-dependence requirement $e_W \leq 3/160$ explains why a sample-size statement must name its goal: robust coverage under arbitrary dependence, model-label recovery, or selection of the smallest justified endpoint. Thousands of wells are not necessary merely to obtain some 95% marginal rule. Neither (7) nor (13) implies coverage conditional on $K = 6$, which under W is 121/128 whatever the calibration.

4.5 Certifying a dependence tolerance

Proposition 4.5. *Under Contract 2.1 with the source P_ρ of Proposition 3.4 for unknown $\rho \in [0,1]$, let $\theta = P_\rho(N_T \in \{3,4\})$ and let $L(C)$ be a lower confidence bound for a Bernoulli probability from n independent trials with failure probability at most β uniformly in the probability, for instance the exact Clopper–Pearson bound [7] applied to H_n. The rule "six if $L(C) \geq 1/4$, otherwise seven" has marginal coverage at least $\min\{19/20, 31/32 - \beta\}$ for every ρ; in particular $\beta \leq 3/160$ gives uniform 95% coverage over the whole family, without calibration–future independence.*

Proof. If $\rho \leq 3/5$ then $F_\rho(6) \geq 19/20$, so both endpoints cover by the inclusion $\{N_T \leq 6\} \subseteq \{N_T \leq K\}$. If $\rho > 3/5$ then $\theta < 1/4$, so selecting six requires $L(C) \geq 1/4 > \theta$, an event of probability at most β; the identity in Lemma 4.1 then gives coverage at least $F_\rho(7) - \beta \geq 31/32 - \beta$. □

Seven is reported as a conservative fallback, not as proof that $\rho > 3/5$. Near the boundary, strong confidence that six is justified requires substantial data; this is a real statistical boundary of the problem.

5 The detector is part of the identifying measurement

This section changes the observation contract: conditional on the terminal population, each terminal cell is detected independently with the same probability $d \in [0, 1]$, and there are no false objects. Then $Y \mid N \sim \text{Binomial}(N, d)$ and every generating function is evaluated at $w = 1 - d + dz$:

$$G_I^{(d)}(z) = \frac{(g_0(w) + g_1(w))^2}{4}, \qquad G_W^{(d)}(z) = \frac{g_0(w)^2 + g_1(w)^2}{2}. \tag{14}$$

5.1 Sign reversal and a bounded repair

Proposition 5.1 (The low-count signal vanishes at one half)**.**

$$P_W(Y \le 2) - P_I(Y \le 2) = \frac{d^3(2d-1)}{(1+d)^4}. \tag{15}$$

At $d = 1/2$ the Bernoulli law of $\mathbf{1}\{Y \le 2\}$ is identical under I and W; for $d < 1/2$ the ordering of the two means reverses.

Proof. Subtracting the two expressions in (14) gives $\frac{1}{4}[g_0(w) - g_1(w)]^2 = \frac{1}{4}\big[w(1-w)/(2-w)\big]^2$. Thinning the fast law gives detected-count coefficients $(1-d)/(1+d)$ and $2d^i/(1+d)^{i+1}$ for $i \ge 1$, and the slow law gives $1-d$ and d; convolving the appropriate mixtures and summing the coefficients of z^0, z^1, z^2 yields (15). (The accompanying script performs this coefficient computation exactly at several rational d.) □

Even above one half the perfect-count midpoint cannot be reused, because both means have moved. The repair is to use a bounded statistic whose expectation is the generating function itself.

Proposition 5.2 (Bounded repair)**.** *For known $d > 0$ and fixed $0 < z < 1$, the statistic $z^Y \in [0, 1]$ separates the preparations: with $u = d(1-z) \in (0, 1)$,*

$$\Delta(d, z) = \mathbb{E}_W z^Y - \mathbb{E}_I z^Y = \frac{u^2(1-u)^2}{4(1+u)^2} > 0. \tag{16}$$

For n independent detected wells, classification by the midpoint of the two known expectations has each one-sided error at most $\exp(-n\Delta^2/2)$.

Proof. $\mathbb{E}z^Y = G^{(d)}(z)$, so the gap is $\frac{1}{4}[g_0(w) - g_1(w)]^2$ with $w = 1 - u$, which is (16). Each expectation is $\Delta/2$ from the midpoint; apply Hoeffding's inequality to the empirical mean of $[0, 1]$-valued variables. □

At $z = d = 1/2$ the two expectations are $729/1600$ and $369/800$, with gap $9/1600$: the vanished low-count signal does not mean equality of the complete observed laws. The repair is constructive but not sample-efficient: obtaining $e_W \le 3/256$ from the Hoeffding bound at $d = z = 1/2$ requires $n \ge 2\ln(256/3)/\Delta^2$, that is 281,067 wells. This is a sufficient size for this statistic and this concentration method, not a necessary one; strictly positive information is not the same as a practically efficient assay. Within the family (16) the gap increases in u up to $u^* = \sqrt{2} - 1$ and decreases beyond, so for $d > u^*$ the choice $z = 1 - u^*/d$ maximizes it, and for $d \le u^*$ the supremum over $0 < z < 1$ is approached as $z \to 0$, where $z^Y \to \mathbf{1}\{Y = 0\}$. This optimizes the expectation gap, not the error exponent or the experimental cost. Figure 3(a) shows both gaps.

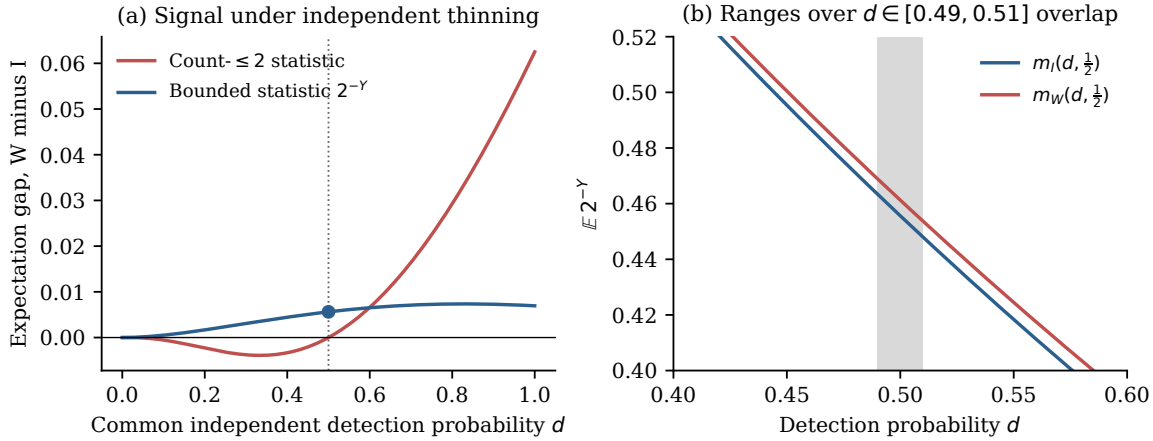

Figure 3: (a) Exact expectation gaps under independent thinning: the count-≤ 2 statistic loses its signal at $d = 1/2$ and reverses sign below, while the bounded statistic 2^{-Y} separates for every $d > 0$; both gaps vanish at complete erasure. (b) The two generating-function expectations at $z = 1/2$ as functions of d: over an efficiency interval $[0.49, 0.51]$ their ranges overlap, so a single midpoint threshold cannot be certified (Section 5.3). No detector efficiency has been estimated in this figure.

5.2 Complete erasure

At $d = 0$ both recorded laws are the point mass at zero. Any function of that law returns the same answer for I and W and cannot return both latent minimal cutoffs. This obstruction is compiled as `deleted_count_cannot_identify_latent_cutoff` in `DetectionBoundary.lean`. Propositions 5.1 and 5.2 are conventional; no general thinning-source theorem is claimed in Lean. Unknown false objects, cell-dependent detection or shared fluctuations in detection invalidate their assumptions.

5.3 Unknown constant efficiency

The midpoint rule of Proposition 5.2 assumes known d. Uncertainty in d does not, however, make the preparations unidentifiable within the fixed menu; three questions must be separated: identifiability from the complete law, a finite-data confidence procedure, and sample efficiency.

Proposition 5.3 (Joint identification)**.** *Both latent preparations have* $\mathbb{E}N = 3$*, and* $\mathbb{E}_I[N(N-1)] = 17/2$*,* $\mathbb{E}_W[N(N-1)] = 9$*. Under common independent thinning* $\mathbb{E}Y = 3d$*,* $\mathbb{E}_I[Y(Y-1)] = \frac{17}{2}d^2$*,* $\mathbb{E}_W[Y(Y-1)] = 9d^2$*. Hence for unknown constant* $d > 0$ *the complete observed law identifies both* d *and the preparation, and the thinning-invariant ratio*

$$\frac{\mathbb{E}[Y(Y-1)]}{(\mathbb{E}Y)^2} = \begin{cases} 17/18, & I, \\ 1, & W, \end{cases} \tag{17}$$

equals $(17+\rho)/18$ *under* P_ρ*, identifying* ρ *and* d *jointly in that family.*

Proof. $\mathbb{E}_I[N(N-1)] = \mathrm{Var}_I\, N + 9 - 3$ with $\mathrm{Var}_I\, N = 5/2$ from (4), and similarly $\mathrm{Var}_W\, N = 3$. Binomial thinning multiplies the r-th factorial moment by d^r. The mean identifies $d = \mathbb{E}Y/3$; equal observed laws at d_I and d_W would force $d_I = d_W$ and then $\frac{17}{2}d^2 = 9d^2$, impossible for $d > 0$. Linearity in ρ gives the last claim. □

These are population identities; plugging sample moments into (17) gives neither an exact confidence interval nor an unbiased estimator, and the argument uses the exact latent scale

$\mathbb{E}N = 3$, the fixed horizon and the known menu. If birth rates, founder number or latent scale are also unknown, mean-based detector identification no longer follows. A finite-data rule that keeps the coverage guarantee with unresolved d is nevertheless available with bounded statistics.

Proposition 5.4 (A safe feasible-set rule)**.** *Fix $L \geq 2$, $0 < z < 1$ and $\beta_1, \beta_2 > 0$. From n independent recorded wells of one preparation at common unknown efficiency d compute*

$$\bar{T} = \frac{1}{n}\sum_i \min(Y_i, L), \qquad \bar{V} = \frac{1}{n}\sum_i z^{Y_i}, \qquad a = L\sqrt{\frac{\ln(2/\beta_1)}{2n}}, \qquad e = \sqrt{\frac{\ln(2/\beta_2)}{2n}},$$

and with $R_L = (L+2)/2^L$ the detector interval $\mathcal{D}(C) = [\frac{\bar{T}-a}{3}, \frac{\bar{T}+a+R_L}{3}] \cap [0,1]$. Writing $m_M(d,z) = G_M^{(d)}(z)$ for $M \in \{I, W\}$, retain model M if some $d \in \mathcal{D}(C)$ satisfies $|m_M(d,z) - \bar{V}| \leq e$. Choose six only when I is retained and W is excluded; otherwise choose seven. Then

$$P_I(N_T \leq K) \geq \frac{245}{256}, \qquad P_W(N_T \leq K) \geq \frac{31}{32} - (\beta_1 + \beta_2), \tag{18}$$

so the rule is uniformly at least 95% when $\beta_1 + \beta_2 \leq 3/160$ and at least 245/256 when $\beta_1 + \beta_2 \leq 3/256$.

Proof. Realize the thinning so that $Y \leq N$ pathwise. Then $0 \leq \mathbb{E}Y - \mathbb{E}\min(Y, L) = \mathbb{E}(Y - L)_+ \leq \mathbb{E}(N - L)_+$, and summing the tails of Lemma 4.3,

$$\mathbb{E}_I(N-L)_+ = \frac{L+6}{2^{L+1}} \leq \mathbb{E}_W(N-L)_+ = \frac{L+2}{2^L} = R_L \qquad (L \geq 2). \tag{19}$$

Hoeffding's two-sided inequality bounds $|\bar{T} - \mathbb{E}\min(Y, L)| \leq a$ except with probability β_1 and $|\bar{V} - m_M(d,z)| \leq e$ except with probability β_2. On the intersection, $3d = \mathbb{E}Y \in [\bar{T}-a, \bar{T}+a+R_L]$, so $d \in \mathcal{D}(C)$, and the true model satisfies the retention test. The true pair is therefore excluded with probability at most $\beta_1 + \beta_2$. Under W, output six requires that exclusion; apply Lemma 4.1. The two statistics may use the same observations, and no calibration–future independence is used. □

The retention test is a finite calculation: for fixed z both $m_I(\cdot, z)$ and $m_W(\cdot, z)$ are continuous and decreasing in d, so a model is feasible exactly when its range over the endpoints of $\mathcal{D}(C)$ meets $[\bar{V} - e, \bar{V} + e]$; an empty detector interval is reported explicitly. The rule solves validity with an uncertain constant detector at a conservative price: it does not guarantee frequent selection of six at small n. At $L = 10$ the truncation remainder is 12/1024, and sampling error can dominate it.

Detector controls. If an independent control experiment yields $d \in [d_-, d_+]$ with failure probability at most β_D, a midpoint test is usable when $G = m_W(d_+, z) - m_I(d_-, z) > 0$; the endpoints follow because both means decrease in d. For $z = 1/2$, $[0.499, 0.501]$ gives $G \approx 0.0040965 > 0$, while $[0.49, 0.51]$ gives $G \approx -0.0096619$ and the single-threshold certificate fails (Figure 3(b)). Since a random control interval makes G random, an unconditional guarantee should predeclare a gap requirement g^*, use the midpoint test only when $G \geq g^*$ and otherwise return seven; with controls independent of the classification wells and a stable common d,

$$e_W \leq \beta_D + \exp(-n(g^*)^2/2), \tag{20}$$

and Lemma 4.1 applies to this deterministic budget. The conservative fallback can increase the I label error but not harm coverage.

Remaining boundaries. As $d \to 0$ both recorded laws concentrate at zero: $\mathbb{P}(Y > 0) \leq \mathbb{E}Y = 3d$, so the total variation distance between the two n-well data laws is at most $6nd$ by comparison with the all-zero law and subadditivity over products. No fixed n achieves uniform label accuracy over all $d > 0$ without a positive efficiency floor; pointwise identifiability and uniform distinguishability differ. If efficiency varies randomly between wells, even independently of N, ratio (17) acquires the factor $\mathbb{E}[d^2]/(\mathbb{E}d)^2$, and a factor 18/17 makes the I value equal the W constant-efficiency value: the moment diagnostic is confounded, although this does not assert equality of the full observed laws. Type-dependent detection, shared detection failures, population-dependent detection and false objects require new observation models.

6 Uniform safety without resolving switching

This section uses Contract 2.2. The switching rates are arbitrary finite nonnegative numbers, so the theorems hold without identifying the memory mechanism. Deletion is arbitrary, which makes an upper-count event easier to satisfy but can destroy all information about outgrowth; coverage of the recorded count and informativeness of the detector are reported separately.

6.1 The compiled rate-box theorem

Theorem 6.1 (Recorded-count safety)**.** *Let the per-cell rates satisfy*

$$b_S, b_R \in [1/10,\, 501/5000], \qquad d_S, d_R \in [1/20,\, 501/10000], \qquad T = 7, \tag{21}$$

with arbitrary finite switching rates and an arbitrary initial type for a single founder. Let the preparation be empty, single-founder or multiple-founder with multiple-founder probability at most γ and an arbitrary multiple-founder law; let $D \leq N_7$ be the genuine detected count and $\mathbb{P}(A > 0) \leq \eta$. Then

$$\mathbb{P}(D + A \leq 4) \geq \frac{963}{1000}(1 - \gamma) - \eta. \tag{22}$$

In particular $\gamma \leq 0.001$ and $\eta \leq 0.01$ give $\mathbb{P}(D + A \leq 4) \geq 0.952037$.

Proof. For a function f of the total count $n = s + r$, the typed generator is

$$\mathcal{L}f(s,r) = B(s,r)[f(n+1)-f(n)]+D_0(s,r)[f(n-1)-f(n)], \quad B = b_S s+b_R r,\ D_0 = d_S s+d_R r, \tag{23}$$

with vanishing rates at zero; switching cancels because it leaves n unchanged. If f is decreasing, then $B \leq \lambda n$ with $\lambda = 501/5000$ and $D_0 \geq \mu n$ with $\mu = 1/20$ give

$$\mathcal{L}f(s,r) \geq \lambda n[f(n+1) - f(n)] + \mu n[f(n-1) - f(n)]. \tag{24}$$

Retain the counts zero through five and send an exit above five to an absorbing cemetery state six with payoff zero; count five must be retained because histories can return from five to four. On the retained states use the scalar rates $\lambda n, \mu n$ and the payoff $h(n) = \mathbf{1}\{n \leq 4\}$, $h(6) = 0$. The total scalar rate is at most $5(\lambda + \mu) = 0.751 < 1$, so uniformization at rate one gives a stochastic matrix $P = I + Q$. It preserves nonnegative decreasing functions that vanish at six: for adjacent live states the coefficients of successive differences are nonnegative because $1 - \lambda n - \mu(n + 1) \geq 0$ for $0 \leq n \leq 4$, and the cemetery boundary is immediate. Hence the scalar backward solution $v(t) = e^{tQ}h$ is decreasing for every t.

Apply (24) to $v(t, \cdot)$ on the finite stopped typed state space: v is a subsolution of the typed backward equation with the same initial payoff, so positivity of the finite Markov semigroup bounds the typed success probability below by $v(t, n)$. Equivalently, uniformize at a common clock rate bounding the typed rates, which may depend on the finite switching

Table 3: Sufficient operating budgets from (22). These are mathematical allowances, not measured errors.

Condition	Sufficient requirement	Consequence
General 95% target	$0.963\gamma + \eta \le 0.013$	coverage ≥ 0.95
$\gamma = 0.001$	$\eta \le 0.012037$	coverage ≥ 0.95
$\eta = 0.01$	$\gamma \le 1/321$	coverage ≥ 0.95
Compiled reference budgets	$\gamma = 0.001,\ \eta = 0.01$	coverage ≥ 0.952037

rates, and induct over clock steps using the preserved class of decreasing scalar test functions. Restoring the discarded histories can only increase the success probability, and the stopped-clock correspondence and nonexplosion transfer the inequality to the original source. Appendix B verifies $v(7,1) \ge 963/1000$ by an exact rational certificate. Thus every allowed single-founder law, including any initial-type mixture, has $\mathbb{P}(N_7 \le 4) \ge 963/1000$. With preparation weights e, w, m for empty, single and multiple founders,

$$\mathbb{P}(N_7 \le 4) \ge e + w\,(0.963) \ge (1-m)(0.963) \ge (1-\gamma)(0.963),$$

and finally $\{N_7 \le 4\} \subseteq \{D+A \le 4\} \cup \{A > 0\}$ gives (22) by subadditivity, with no independence assumption on the detector. □

The theorem is compiled as `observed_four` in `RateBox.lean`, with the scalar comparison in `SourceLower.lean` and the rational certificate as `finite_scalar_certificate`. The compiled source bound is 0.963; the certificate value is $0.96355\ldots$ and the full scalar law gives more (Section 6.3), but those are not substituted for the formal bound. Table 3 lists sufficient budgets. The theorem does not prove that four is a minimal latent or recorded cutoff over the box, and it is not a detector-quality guarantee. The budgets γ and η are asserted allowances, not measured preparation or instrument performance.

6.2 The printed box is narrower than the proof needs

Proposition 6.2. *The conclusion (22) holds, with the same certificate, under the one-sided conditions*

$$0 \le b_S, b_R \le 501/5000, \qquad 1/20 \le d_S, d_R < \infty, \tag{25}$$

with arbitrary finite switching, the same preparation and observation budgets and the same horizon.

Proof. The comparison (24) uses only $B \le \lambda n$ and $D_0 \ge \mu n$; the lower birth limit and the upper death limit in (21) are never used. Finite-state stopping permits a clock depending on the actual finite rates of each model, so no common upper bound on death or switching rates is needed, and the birth cap gives nonexplosion for each model. The scalar certificate is unchanged. □

The relative width 0.2% of each interval in (21) is therefore not the tolerance of the safety argument: it admits slow or nondividing cells and arbitrarily large death rates. The compiled theorem assumes the smaller box; the extension is conventional and is not relabelled as Lean-checked.

6.3 Full scalar domination

Lemma 6.3 (Pathwise domination)**.** *If every cell has birth rate at most λ and death rate at least μ, whatever its type and however it switches, then the total count N_t can be coupled with a*

linear birth–death process Z_t with per-cell rates λ, μ and the same initial count so that $N_t \le Z_t$ for all t.

Proof. Match each actual cell to a dominating cell. A matched dominating cell proposes births at rate λ; each proposal is accepted as an actual birth with probability $b_{\text{type}}/\lambda \le 1$, and the two new cells are matched. A death clock of rate μ removes a matched pair together; the actual cell carries an additional death clock of rate $d_{\text{type}} - \mu \ge 0$, after which its partner continues as an unmatched dominating cell. Unmatched dominating cells evolve with rates λ, μ. Switching only changes the rates attached to the matched actual cell. Every actual cell then has its own birth and death intensities, Z is a linear birth–death process, and matching gives $N_t \le Z_t$ pathwise. Zero rates are handled by omitting their clocks; finite actual rates and the linear birth bound give nonexplosion. This is the standard coupling construction [17]. □

Lemma 6.4 (Linear birth–death law)**.** *For one initial cell, $\lambda \ne \mu$, $E = e^{-(\lambda-\mu)T}$ and*

$$a = \frac{\mu(1-E)}{\lambda - \mu E}, \qquad q = \frac{\lambda(1-E)}{\lambda - \mu E},$$

one has $\mathbb{P}(Z_T = 0) = a$, $\mathbb{P}(Z_T = j) = (1-a)(1-q)q^{j-1}$ for $j \ge 1$, and

$$\mathbb{P}(Z_T \le k) = F_{\lambda,\mu,T}(k) = 1 - (1-a)q^k. \tag{26}$$

For $\mu = 0$ this is the Yule law $1 - (1 - e^{-\lambda T})^k$. In the critical case $\lambda = \mu$, $\mathbb{P}(Z_T > k) = (\lambda T)^k/(1+\lambda T)^{k+1}$.

This is Kendall's classical result [3, 13]; the generating function g solves $\partial_t g = \lambda g^2 - (\lambda + \mu)g + \mu$ with $g(0, z) = z$. Appendix C records the monotonicity in E used below. Under Contract 2.2 the two lemmas give $\mathbb{P}(N_T \le k) \ge F_{\lambda,\mu,T}(k)$ for every single-founder law, without the finite-state loss of Theorem 6.1: at $\lambda = 501/5000$, $\mu = 1/20$, $T = 7$ the scalar value is $F_{\lambda,\mu,7}(4) \approx 0.96642437695$ and the recorded bound is about 0.95545795258 (numerical evaluations of a proved expression, not replacements for the compiled 0.952037).

Proposition 6.5 (A larger birth cap with cutoff four)**.** *If $0 \le b_S, b_R \le 0.105$, $d_S, d_R \ge 0.05$ are finite, switching is arbitrary and finite, $T = 7$, $\gamma \le 0.001$ and $\eta \le 0.01$, then*

$$\mathbb{P}(D + A \le 4) > 0.95034241414 > 0.95. \tag{27}$$

Proof. Set $\lambda = 21/200$, $\mu = 1/20$, $x = (\lambda - \mu) \cdot 7 = 77/200$ and $E_- = \sum_{j=0}^{7} (-x)^j/j!$. The odd Taylor truncation is a lower bound, $E_- \le e^{-x} = E$, by the sign of the remainder. By Appendix C, $F_{\lambda,\mu,7}(4)$ is increasing in E whenever $4\lambda \ge 5\mu$, which holds here, so

$$F_- = 1 - \frac{\lambda - \mu}{\lambda - \mu E_-}\Big[\frac{\lambda(1-E_-)}{\lambda - \mu E_-}\Big]^4 \le F_{\lambda,\mu,7}(4)$$

is a rational lower bound on the scalar probability. Exact arithmetic gives $\frac{999}{1000}F_- - \frac{1}{100} > 0.95034241414$. Lemmas 6.3 and 6.4 give $\mathbb{P}(N_7 \le 4) \ge F_-$ for every single-founder law, and the preparation and observation accounting of Theorem 6.1 completes the proof. □

Proposition 6.6 (One extra count removes the death floor)**.** *If $0 \le b_S, b_R \le 0.1$, all death and switching rates are arbitrary finite nonnegative numbers, $T = 7$, $\gamma \le 0.001$ and $\eta \le 0.01$, then*

$$\mathbb{P}(D + A \le 5) > 0.95670013878 > 0.95. \tag{28}$$

Proof. Lemma 6.3 with $\mu = 0$ dominates N by a Yule process of rate 0.1, whose distribution function is $1 - (1 - e^{-0.7})^k$ and increases with $e^{-0.7}$. The odd truncation $r = \sum_{j=0}^{7} (-0.7)^j/j! = 1191801547/2400000000 \le e^{-0.7}$ gives $\mathbb{P}(N_7 \le 5) \ge 1 - (1-r)^5$, and exact arithmetic gives $\frac{999}{1000}[1 - (1-r)^5] - \frac{1}{100} > 0.95670013878$. Accounting is as before. □

Table 4: Sufficient source conditions at $T = 7$ under the founder budget $\gamma \leq 0.001$ and false-object budget $\eta \leq 0.01$.

Per-cell rate conditions	Cutoff	Lower bound	Evidence
Box (21)	4	0.952037	Lean-checked (`observed_four`)
Births ≤ 0.1002, deaths ≥ 0.05	4	0.952037	Conventional, Prop. 6.2; same certificate
Births ≤ 0.105, deaths ≥ 0.05	4	> 0.95034241414	Conventional, Prop. 6.5; rational certificate
Births ≤ 0.1, deaths arbitrary	5	> 0.95670013878	Conventional, Prop. 6.6; rational certificate

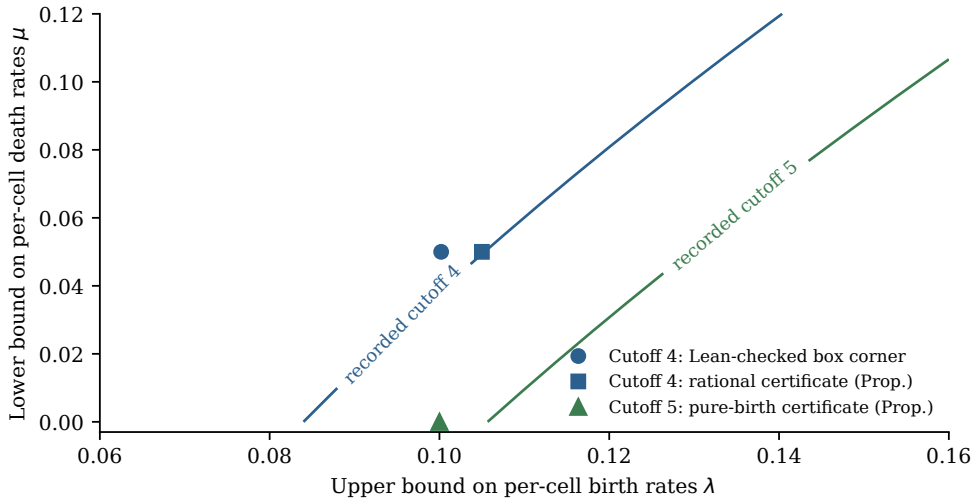

Figure 4: Sufficient demographic conditions in the plane of birth cap λ and death floor μ. The curves are numerical level sets of $0.999\, F_{\lambda,\mu,7}(k) - 0.01 = 0.95$ for recorded cutoffs $k = 4$ and $k = 5$ under full scalar domination (display only); the region above and to the left of a curve is sufficient for that cutoff. The three marked points are the proved operating points of Table 4. The Lean-checked corner lies strictly inside the cutoff-four region because the finite-state certificate discards histories that the full scalar law retains.

With the same pure-birth route, cutoff four gives only $0.9248\ldots$, so the extra count is needed when nothing is assumed about deaths. Table 4 and Figure 4 summarize the operating conditions. None of the rows establishes that an assay satisfies its conditions, and no minimality claim is made for the new cutoffs. The result admits a concrete tradeoff: a wider count region removes an entire demographic lower bound.

7 Robust prediction under inherited-state rate uncertainty

This section uses Contract 2.3, in which the switching rates are small but uncertain and the founder type is random, so that a single-founder count law genuinely depends on the inherited state. Rates are per day.

7.1 The class and the theorem

The reference vector and the uncertainty set are

$$\begin{gathered} a_0 = (b_S, d_S, q_S;\, b_R, d_R, q_R) = (0,\, 0.3,\, 0.001;\;\, 0.1,\, 0,\, 0.00001), \\ b_S + |d_S - 0.3| + |q_S - 0.001| \le 0.0005, \qquad |b_R - 0.1| + d_R + |q_R - 0.00001| \le 0.0005, \end{gathered} \tag{29}$$

that is, two per-type ℓ^1 balls of radius 0.0005 intersected with the nonnegative orthant (not a six-coordinate rectangle). The horizon is $T_0 = \log 2/0.10001 \approx 6.93078$ days. Conditional on a single founder, its resistant probability satisfies $p \in [p_-, p_+]$ with $p_\pm = 5/24 \pm 1/1000$. With a shared static environment U, these bounds hold for every conditional one-founder law, and the mixing law is the distribution of U among single-founder preparations. Preparation and observation are as in Contract 2.2: multiple-founder probability at most γ with an arbitrary multiple-founder law, arbitrary deletion, and $\mathbb{P}(A > 0) \le \eta$.

Theorem 7.1 (Robust endpoint three)**.** *Under Contract 2.3, put*

$$\begin{gathered} a = \frac{599}{603}, \qquad c = \frac{9950}{10051}, \qquad r_3 = \frac{1993}{4000}\Big(1 + \frac{c}{2} + \frac{c^2}{4}\Big), \\ L_3 = (1 - p_+)a + p_+ r_3 = \frac{176701212731773153}{182749885209000000}. \end{gathered} \tag{30}$$

Then, also under conditional shared-environment mixing, $\mathbb{P}(Y \le 3) \ge \max\{0, (1-\gamma)L_3 - \eta\}$. *For* $\gamma = 1/1000$ *and* $\eta = 1/100$*,*

$$\mathbb{P}(Y \le 3) \ge B_3 := \frac{6470259728405606661}{6768514267000000000} = 0.955935006291\ldots > \frac{19}{20}. \tag{31}$$

At the reference rates with one founder, $p = 5/24$ *and perfect observation, every fixed endpoint below three has coverage strictly below* $19/20$*; thus three is the smallest fixed endpoint valid over the class.*

The source theorem, its shared-environment extension, the accounting and the reference minimality are compiled as `robust_endpoint_three_budget`, `robust_endpoint_three_mixed_budget`, `sharp_constant` and `reference_minimum_endpoint` in `LowFounderPrediction/`. The class is supplied, not estimated from the lineage data of Section 8. The theorem concerns a fixed time and endpoint, not a trajectory bound or an optional-stopping guarantee.

7.2 Proof

Finite successful histories. The class (29) implies

$$\begin{aligned} b_S + d_S + q_S &\le Q_S := \tfrac{603}{2000}, \qquad & d_S &\ge d_* := \tfrac{599}{2000}, \\ b_R + d_R + q_R &\le Q_R := \tfrac{10051}{100000}, \qquad & b_R &\ge b_* := \tfrac{199}{2000}. \end{aligned} \tag{32}$$

For $k \ge 1$ retain the states $\mathcal{K}_k = \{0, S, R, 2R, \ldots, kR\}$ and send every exit to a cemetery ∂ with success value zero; every retained state has count at most k, and the killed chain follows the original until first exit, so its success probability is at most $\mathbb{P}(N_t \le k)$ even when an omitted trajectory could later return. The physical-time comparison needs nonexplosion: with $V(s, r) = 1 + s + r$ the generator satisfies $\mathcal{G}V \le (b_S + b_R)V$, and stopping on finite sublevel sets, Gronwall's inequality and Markov's inequality show that infinitely many jumps in finite time have probability zero [21]. The compiled chronological bridge transfers the killed comparison to the actual source measure.

Backward subsolution. Set $E(t) = e^{-Q_R t}$, $z(t) = c(1 - E(t))$ with $c = b_*/Q_R$, and for $1 \le n \le k$

$$v_n(t) = \sum_{j=n}^{k} \binom{j-1}{n-1} E(t)^n c^{j-n} (1 - E(t))^{j-n}, \qquad v_{k+1} \equiv 0. \tag{33}$$

These functions are nonnegative and, using $(j-n)\binom{j-1}{n-1} = n\binom{j-1}{n}$, satisfy $v_n' = nb_* v_{n+1} - nQ_R v_n$ with $v_n(0) = 1$. Put $v_0 = 1$, $v_S = a = d_*/Q_S$ and $v_\partial = 0$. At S the killed generator applied to v is $d_S v_0 + q_S v_1 + b_S \cdot 0 - (b_S + d_S + q_S)a \ge d_* - Q_S a = 0 = v_S'$, because the birth destination SS is not retained and $v_1 \ge 0$. At nR it is at least $nb_R v_{n+1} - n(b_R + d_R + q_R)v_n \ge nb_* v_{n+1} - nQ_R v_n = v_n'$, because retained death and switch destinations have nonnegative value and omitted destinations have value zero. At 0 and ∂ both sides vanish. At time zero v lies below the success indicator h of the retained states. With G the finite killed generator, whose exponential preserves coordinatewise nonnegativity, variation of constants gives

$$e^{tG} h - v(t) = e^{tG}(h - v(0)) + \int_0^t e^{(t-u)G} \{Gv(u) - v'(u)\} \, \mathrm{d}u \ge 0.$$

Hence $\mathbb{P}_S(N_t \le k) \ge a$ and $\mathbb{P}_R(N_t \le k) \ge v_1(t) = E(t) \sum_{j=0}^{k-1} z(t)^j$. The $k = 3$ subsolution and the source bridge are compiled; the argument for general k is conventional.

Exact rationalization. Since $T_0 \le 7$ (from $\log 2 \le 7/10$) and $Q_R = 0.10001 + 1/2000$,

$$E(T_0) = \tfrac{1}{2} e^{-T_0/2000}, \qquad \tfrac{1993}{4000} \le E(T_0) \le \tfrac{1}{2}, \qquad 1 - E(T_0) \ge \tfrac{1}{2}, \tag{34}$$

using $e^{-x} \ge 1 - x$; no numerical exponential enters. Substituting the lower bounds for E and $z \ge c/2$ in v_1 gives the rational r_3 of (30). Since $r_3 \le a$, the worst founder mixture over $p \in [p_-, p_+]$ is at p_+, and every conditional one-founder law has success probability at least L_3; integrating those laws over the environment preserves the bound without any independence between founder type and environment.

Accounting. Empty preparations succeed with probability one. Discarding nonnegative multiple-founder contributions,

$$\mathbb{P}(N_{T_0} \le 3) \ge e_0 + wL_3 = (1 - m)L_3 + e_0(1 - L_3) \ge (1 - \gamma)L_3, \tag{35}$$

and $\{N_{T_0} \le 3\} \cap \{A = 0\} \subseteq \{D + A \le 3\}$ gives $\mathbb{P}(Y \le 3) \ge \mathbb{P}(N_{T_0} \le 3) - \mathbb{P}(A > 0)$. Arbitrary dependence permits subtraction of η but not multiplication by $1 - \eta$: false objects concentrated on otherwise successful outcomes remove their entire mass. This proves (31).

Minimality. At the reference source start with one resistant cell and retain only pure resistant birth histories without switching that end with $n \in \{3, \dots, 7\}$ cells at T_0. With total per-cell event rate $b_R + q_R = 0.10001$ and $e^{-(b_R + q_R)T_0} = 1/2$ exactly, a history with $n - 1$ events, all births, has probability $\frac{1}{2}(c_0/2)^{n-1}$ with $c_0 = b_R/(b_R + q_R) = 10000/10001$. Each such history has $N_{T_0} \ge 3$, so

$$\mathbb{P}(N_{T_0} \ge 3) \ge \frac{5}{48} \sum_{j=2}^{6} \Big(\frac{c_0}{2}\Big)^j > \frac{1}{20}, \tag{36}$$

by rational arithmetic; the complement shows that endpoint two, and every smaller endpoint, undercovers at an admissible source. This establishes uniform minimality over the class, not optimality of three for every individual source. □

7.3 Endpoint four and a confidence allowance

Applying the same proof with $k = 4$ gives

$$r_4 = \frac{1993}{4000}\sum_{j=0}^{3}\Big(\frac{c}{2}\Big)^j = \frac{941626977073717}{1015378162651000},$$
$$B_4 = (1-\tfrac{1}{1000})L_4 - \tfrac{1}{100} = \frac{16473049504746783406959}{17007584224404250000000} = 0.968570802731\ldots, \tag{37}$$

with $L_4 = (1-p_+)a + p_+ r_4$ and $r_4 \le a$. This is a complete conventional proof; endpoint four is not source-connected in Lean. Suppose a confidence event C has probability at least $1-\delta$ and that for almost every training data set in C the *conditional future* law satisfies the full contract with coverage at least B. Then $\mathbb{P}(Y \le k) \ge \mathbb{E}[\mathbf{1}_C \mathbb{P}(Y \le k \mid \text{training})] \ge (1-\delta)B$. With $\delta = 0.01$,

$$\tfrac{99}{100}B_4 = 0.958885094703\ldots > 0.95, \qquad \text{whereas} \qquad \tfrac{99}{100}B_3 < 0.95. \tag{38}$$

This confidence multiplication requires the conditional-law premise; an estimated rate vector lying in a box does not establish it, and it is not the arbitrary-dependence menu theorem of Section 4, which needs no such premise because the ordered endpoints make one error harmless. The accounting algebra is compiled as `conditional_coverage_accounting`; the validity of a statistically estimated confidence region remains a separate obligation.

Example 7.2 (All positive rates). The rates $(0.0001, 0.2999, 0.0011;\ 0.1001, 0.0001, 0.00002)$ per day have ℓ^1 distances 0.0003 and 0.00021 from a_0, both within 0.0005, so Theorem 7.1 gives B_3 for them with arbitrary empty preparations. At $\eta = 0.02$ the endpoint-three certificate falls below 0.95; endpoint four restores the stated 1% inference allowance at $\eta = 0.01$. The zero-rate reference isolates the dependence mechanism; positive-rate examples show mathematical robustness but remain hypothetical operating points.

The short history family can be conservative when switching produces successful mixed-type descendants that are sent to the cemetery; increasing the endpoint alone cannot remove that loss, since the retained-history envelope has a ceiling in k. Successful-history loss, finite-state omission, observation budgets and statistical confidence losses are separate quantities and are kept separate above.

8 Empirical scope and prospective use

The mathematical examples are synthetic. A separate frozen comparison used recorded lineage families released with the lineage-tree hidden Markov study of Mohammadi et al. [19]; observed-phase count models were fitted to the annotations, not a reproduction of that study's model, and with a different target. One training run (50 families), one calibration run (50) and one test run (75) were used; predictions were frozen before test access; all families, including extinct ones, were retained. The target was the recorded descendant count of a family at a common frame. Rechecking the test workbook identified one absorbing extinction (a root dividing at frame 46 with both daughters recorded dead at frame 54), which corrects the outcome count to 51 known and 24 unresolved among 75 without refitting or a new holdout.

If c of $n-u$ resolved outcomes are covered and u are unresolved, the realized coverage is identified only within $[c/n, (c+u)/n]$; both extremes are attainable. With $24/75 > 0.05$ unresolved, even all-success on the resolved set cannot bring the realized lower endpoint to 0.95. These data establish neither biological 95% coverage nor a comparative winner, and they neither refute the model nor the source study. Run-level calibration by the group-maximum rank rule of distribution-free prediction [1, 15, 25] returns an infinite endpoint with one calibration

Table 5: Frozen, uncalibrated endpoints after the extinction correction: 51 resolved and 24 unresolved outcomes in one independent test run. Intervals are logical identification bounds on the realized fraction, not confidence intervals. Mean endpoint uses all 75 predictions.

Frozen model	Covered of 51 known	Identification interval	Mean raw endpoint
Exponential phases	51	[51/75, 75/75]	7.773
Erlang-2 phases	50	[50/75, 74/75]	6.387
Mean-composition scalar	51	[51/75, 75/75]	9.547
Empirical training baseline	49	[49/75, 73/75]	6.000

run, and nineteen calibration runs is merely the first potentially finite case for that rule. Constant predictors cannot retain an efficiency advantage after common residual calibration, since $a + (Y - a)_{(r)} = Y_{(r)}$; covariate-aware predictors can differ but still need independent evidence.

Prospective validation requires untouched independent runs, identical selection, a common horizon, explicit extinctions and censoring, founder records and detector controls. Fifty related families cannot replace fifty independent runs. Even for independent Bernoulli units with all successes, the one-sided exact 95% lower confidence limit is $0.05^{1/n}$, so at least 59 units are needed to exceed 0.95. Table 6 states, for each design step, what the results of this article resolve.

9 Verification, reproducibility and limitations

9.1 Evidence map

Table 7 lists every claim used above with its evidence level. The Lean development uses Lean 4 (toolchain `v4.30.0`) with mathlib, compiled with warnings treated as errors and with transitive source dependencies materialized; the exported declarations were probed for their elaborated types. Kernel checking establishes consequences of the encoded source process and hypotheses. It does not establish biological adequacy of the model, the validity of an estimated rate region, or clinical utility. The Lean statements name measures on an implemented chronological jump-trajectory space and connect the finite killed comparisons to the actual source measure through a compiled bridge, so no premise assumes the desired distribution function.

9.2 Reproducibility

The accompanying script `check_paper.py` replays every exact constant in this article with rational and integer arithmetic (the distribution functions and mass functions from the convolution of single-founder laws, the binomial certificates by an exact integer recurrence, the sixteen-row scalar certificate, the odd Taylor lower bounds, the linear birth–death rational bounds, the class constants L_3, B_3, r_4, B_4 and the minimality witness), verifies that every printed decimal is a downward rounding of its exact value, checks the closed-form linear birth–death law against a numerical solution of the forward equations, and regenerates all four figures. It runs in well under a minute and exits nonzero on the first failed check. The manuscript build consumes only retained figures and generated table rows; it does not run a stochastic search or refit the empirical comparison.

9.3 Limitations and claims to avoid

The menu theorems fix the source to two preparations at a normalized horizon and assume independent, perfectly counted calibration wells; a recorded well identifier is not proof of

Table 6: A design checklist derived from the theorems. Each row names the quantity a researcher records and the question it settles.

Step	What it resolves
Define the event: true versus recorded count, horizon, endpoint, conditioning on preparation	Prevents a recorded-count guarantee being read as latent eradication (Section 2).
Define the sampling unit: well and shared batch or run; related cells grouped	Determines whether the independence premise of calibration is credible (Sections 4, 8).
Establish the preparation: founder count and its uncertainty, empty and multiple-founder frequencies	Determines which source law is predicted and the budget γ (Sections 3, 6).
Specify observation: detection mechanism, constant versus variable efficiency, false objects	Selects the applicable theorem and statistic and the budget η (Section 5).
Obtain a baseline: fixed seven for the menu, or a scalar envelope for a justified demographic class	Supplies a conservative prediction before any smaller-cutoff calibration (Sections 4.4, 6).
Add a discriminating measurement: full two-founder counts, or the detector-aware statistics	Addresses the dependence that changes the cutoff (Propositions 4.4, 4.5, 5.4).
Compare decisions at matched guarantees: coverage, selection frequencies, endpoint size	Prevents a narrower uncalibrated region from being declared superior (Section 4.4).
Validate prospectively on untouched units with an ascertainable common endpoint	Tests whether the source and observation assumptions support actual use (Section 8).

independence, and wells on one plate may share batch effects. The demographic envelopes assume the stated one-sided rate bounds; no assay has been shown to satisfy them. The inherited-state class is supplied, not estimated. Constant common detection is a model; type-dependent, shared or population-dependent detection needs new observation models. No count exceedance in these models demonstrates genetic resistance, persister cells, treatment failure or patient risk, and no dosing or maintenance recommendation follows. A larger cutoff reduces false exceedances under a null model while reducing sensitivity to a genuine alternative; that tradeoff requires an alternative distribution and a decision or power analysis that this article does not supply. Optimal sample sizes, sequential procedures and general detector models are open and are not needed for the claims made here.

Table 7: Evidence level of each claim. Lean names are declarations in the accompanying development (directories `MemoryPrediction/` and `LowFounderPrediction/`); "conventional" means a written proof in this article with exact arithmetic replayed by the checking script.

Claim	Evidence
Single-founder observation equality (Thm. 3.1)	Lean: `chronological_single_observation_equal`; general family-space identity in `SharedEnvironment.lean`. No genealogical-tree implementation.
Cutoffs six and seven (Thm. 3.1)	Lean: `chronological_minimal_endpoints`, `source_cdf_formula`; general-k formula (2) conventional.
Many founders (Prop. 3.3); continuous family (Prop. 3.4)	Conventional.
Nested endpoints (Lemma 4.1)	Lean: `nested_endpoint_lower`, `upper_endpoint_loss`.
4000-well coverage 245/256 (Thm. 4.2(i))	Lean: `classification_error`, `adaptive_prediction_sharp`.
1280- and 1024-well coverage (Thm. 4.2(ii))	Conventional binomial law plus exact integer certificates (Appendix A).
640- and 576-well rule (Prop. 4.4); ρ confidence rule (Prop. 4.5)	Conventional; exact integer certificates (12).
Fixed and randomized baselines; independence formulas (13)	Conventional algebra from exact source values.
Sign reversal and bounded repair (Props. 5.1, 5.2)	Conventional; exact rational examples.
Complete erasure	Lean: `deleted_count_cannot_identify_latent_cutoff`.
Unknown-d identification and feasible-set rule (Props. 5.3, 5.4)	Conventional.
Rate-box bound 0.952037 (Thm. 6.1)	Lean: `observed_four`, `source_scalar_lower`, `finite_scalar_certificate`. No minimality claim.
One-sided conditions, broad caps, pure-birth (Props. 6.2–6.6)	Conventional; rational certificates (Appendix C).
Endpoint three, mixing, minimality (Thm. 7.1)	Lean: `robust_endpoint_three_budget`, `robust_endpoint_three_mixed_budget`, `sharp_constant`, `reference_minimum_endpoint`, `preparation_budget`.
Endpoint four and confidence allowance (37)–(38)	Conventional source proof; accounting algebra Lean: `conditional_coverage_accounting`.
Empirical comparison (Section 8)	Frozen predictions and annotation audit; identification bounds only.
Biological applicability; clinical performance	Not established in this work.

A Exact binomial certificates

For a Bernoulli probability a/b define the integers $a_0 = (b-a)^n$ and $a_{j+1} = a(n-j)a_j/((b-a)(j+1))$; induction gives $a_j = \binom{n}{j}a^j(b-a)^{n-j}$, so every division is exact and $\sum_{j\le k} a_j/b^n$ is the binomial distribution function at k. The five certificates used in Section 4 are the following integer inequalities.

$$160\sum_{j\le 607}\binom{1024}{j}5^j 3^{1024-j} \le 3\cdot 8^{1024},$$

$$256\sum_{j\le 759}\binom{1280}{j}5^j 3^{1280-j} \le 3\cdot 8^{1280},$$

$$256\sum_{j\le 2374}\binom{4000}{j}5^j 3^{4000-j} \le 3\cdot 8^{4000},$$

$$160 \sum_{j \ge 149} \binom{576}{j} 7^j 25^{576-j} \le 3 \cdot 32^{576},$$
$$256 \sum_{j \ge 165} \binom{640}{j} 7^j 25^{640-j} \le 3 \cdot 32^{640}.$$

The first row bounds e_W for the low-count rule (8) at $n = 1024, 1280, 4000$; the second bounds e_W for the counts-3-or-4 rule (11) at $n = 576, 640$. The thresholds are $\lceil 19n/32 \rceil - 1 = 607, 759, 2374$ and $\lceil 33n/128 \rceil = 149, 165$. The script evaluates these with the recurrence above and no floating-point arithmetic; the I errors in Table 2 are computed the same way. These are exact computations supporting conventional theorems, not statements trusted by the Lean kernel; the compiled 4000-well theorem uses the Hoeffding route instead.

B The finite scalar certificate

Let Q be the seven-state generator of the scalar killed chain in Theorem 6.1 (live states $0, \dots, 5$, cemetery six) and $P = I + Q$. Nonnegativity of the Poisson weights gives

$$v(7,1) = e^{-7} \sum_{j \ge 0} \frac{7^j}{j!} (P^j h)(1) \ge e^{-7} \sum_{j=0}^{15} \frac{7^j}{j!} (P^j h)(1).$$

Set $M = 10^6$, $a_{0,i} = M \mathbf{1}\{i \le 4\}$, $a_{j,6} = 0$, and for $0 \le i \le 5$

$$a_{j+1,i} = \Big\lfloor \frac{(5000 - 751i) a_{j,i} + 501i\, a_{j,i+1} + 250i\, a_{j,\max(i-1,0)}}{5000} \Big\rfloor. \tag{39}$$

All coefficients are nonnegative, so induction gives $a_{j,i}/M \le (P^j h)(i)$. The sixteen values at the initial state are

j	0	1	2	3	4	5	6	7
$a_{j,1}$	1000000	1000000	1000000	1000000	997580	992141	983746	972743
j	8	9	10	11	12	13	14	15
$a_{j,1}$	959583	944728	928611	911617	894077	876268	858418	840711

Since $e < 271829/100000$ and the seventh power of that rational is below 1097, $e^{-7} \ge 1/1097$, and exact rational arithmetic yields

$$\frac{1}{1097 \cdot 10^6} \sum_{j=0}^{15} \frac{7^j}{j!} a_{j,1} = \frac{3134312565151423153}{3252877056000000000} = 0.9635508847\ldots > \frac{963}{1000}.$$

Only the rational lower bound 963/1000 is used. Lean checks the initial rows, each step inequality, nonnegativity, the exponential bound and the final comparison (`ScalarCertificate.lean`). The finite lower comparison is related to finite-state projection [20], but the source-to-scalar inequality is proved separately here.

C The linear birth–death law and its monotonicity

With $E = e^{-(\lambda-\mu)T}$, $\lambda \ne \mu$, the generating function of Z_T from one cell is $g(z) = (a + (1 - a - q)z)/(1 - qz)$ with a, q as in Lemma 6.4; expanding in z gives the stated masses and (26). For $\mu \to 0$, $a \to 0$ and $q \to 1 - E$, the Yule law. Regard the tail $t(E) = (1 - a)q^k$ as a function of $E \in (0, 1)$ at fixed λ, μ; for $k = 4$,

$$\frac{\mathrm{d}}{\mathrm{d}E} \log t(E) = \frac{5\mu - 4\lambda - \mu E}{(1 - E)(\lambda - \mu E)},$$

a closed form the script verifies symbolically. The denominator is positive on $(0, 1)$ because $\lambda > \mu E$, and the numerator is negative whenever $4\lambda \geq 5\mu$. Hence $F_{\lambda,\mu,T}(4)$ increases with E, and any rational $E_- \leq E$ gives a rational lower bound on F. In Proposition 6.5, E_- is the degree-seven Taylor polynomial of e^{-x} at $x = 77/200$, whose signed remainder makes it a lower bound. The same route at the corner $\lambda = 501/5000$, $\mu = 1/20$ of box (21) gives a recorded bound above 0.95545, confirming that the finite-state certificate of Appendix B is conservative.

Preparation statement. This manuscript and its accompanying artifacts were prepared with AI assistance. No collaborator authorship, institutional affiliation, external peer review or publication acceptance is asserted. Author and affiliation fields are intentionally blank in this version.

References

[1] Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: a gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023. doi: 10.1561/2200000101.

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

[3] Norman T. J. Bailey. *The Elements of Stochastic Processes with Applications to the Natural Sciences*. Wiley, New York, 1964.

[4] Nathalie Q. Balaban, Jack Merrin, Remy Chait, Lukasz Kowalik, and Stanislas Leibler. Bacterial persistence as a phenotypic switch. *Science*, 305(5690):1622–1625, 2004. doi: 10.1126/science.1099390.

[5] Alexander P. Browning, David J. Warne, Kevin Burrage, Ruth E. Baker, and Matthew J. Simpson. Identifiability analysis for stochastic differential equation models in systems biology. *Journal of the Royal Society Interface*, 17(173):20200652, 2020. doi: 10.1098/rsif.2020.0652.

[6] Alexander P. Browning, Rebecca M. Crossley, Chiara Villa, Philip K. Maini, Adrianne L. Jenner, Tyler Cassidy, and Sara Hamis. Identifiability of phenotypic adaptation from low-cell-count experiments and a stochastic model. *PLOS Computational Biology*, 21(6):e1013202, 2025. doi: 10.1371/journal.pcbi.1013202.

[7] 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.

[8] Companion manuscript. Prediction coverage in inherited-state branching models: an exact mean-closure counterexample, a repaired count threshold, and the effect of founder number, 2026. Manuscript, September 2026; author fields left blank in this version.

[9] 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.

[10] Theodore E. Harris. *The Theory of Branching Processes*. Springer, Berlin, 1963.

[11] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

[12] Arun Iyer, Arnob Alva, Adrián E. Granada, and Shaon Chakrabarti. Inheritable cell-states shape drug-persister correlations and population dynamics in cancer cells. *PLOS Computational Biology*, 21(9):e1013446, 2025. doi: 10.1371/journal.pcbi.1013446.

[13] David G. Kendall. On the generalized "birth-and-death" process. *Annals of Mathematical Statistics*, 19(1):1–15, 1948. doi: 10.1214/aoms/1177730285.

[14] Marek Kimmel and David E. Axelrod. *Branching Processes in Biology*. Springer, New York, 2nd edition, 2015. doi: 10.1007/978-1-4939-1559-0.

[15] Yonghoon Lee, Rina Foygel Barber, and Rebecca Willett. Distribution-free inference with hierarchical data. arXiv:2306.06342, 2023. URL https://arxiv.org/abs/2306.06342.

[16] E. L. Lehmann and Joseph P. Romano. *Testing Statistical Hypotheses*. Springer, New York, third edition, 2005. doi: 10.1007/0-387-27605-X.

[17] Torgny Lindvall. *Lectures on the Coupling Method*. Wiley, New York, 1992. Dover reprint, 2002.

[18] Salvador E. Luria and Max Delbrück. Mutations of bacteria from virus sensitivity to virus resistance. *Genetics*, 28(6):491–511, 1943. doi: 10.1093/genetics/28.6.491.

[19] Farnaz Mohammadi, Shakthi Visagan, Sean M. Gross, Luka Karginov, J. C. Lagarde, Laura M. Heiser, and Aaron S. Meyer. A lineage tree-based hidden Markov model quantifies cellular heterogeneity and plasticity. *Communications Biology*, 5:1258, 2022. doi: 10.1038/s42003-022-04208-9.

[20] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124(4):044104, 2006. doi: 10.1063/1.2145882.

[21] James R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, 1997. doi: 10.1017/CBO9780511810633.

[22] 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.

[23] 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.

[24] 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.

[25] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, New York, 2005. doi: 10.1007/b106715.

[26] G. Udny Yule. A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F.R.S. *Philosophical Transactions of the Royal Society of London. Series B*, 213: 21–87, 1925. doi: 10.1098/rstb.1925.0002.