When single-lineage measurements fail to determine treatment order: sister dependence, branching extinction, and decision-focused assays

Abstract

Measurements sufficient to describe a retained cell lineage need not determine which intervention minimises the survival of the whole family. We construct a finite-type branching source in which complete retained-branch histories have identical laws across a family of sister couplings, while two interventions applied for exactly the same durations have opposite preferred orders. An exact finite-duration formula identifies the decision boundary and shows that it is strictly positive, not at zero covariance. Two new ingredients then control recurring descendant division: the extinction probability is monotone in the descendant division rate, and a sensitivity reserve replaces the crude occupation bound. Together they certify the reversal at $h = \log \frac{8}{7}$ for descendant division rates up to $\varepsilon = 1/30$ instead of $1/500$, and they certify it for the *entire* range $0 \le \varepsilon \le 1$ at pulses of length $\log \frac{50}{49}$ rather than 10^{-6} — four orders of magnitude of pulse length — so that the phenomenon survives descendants dividing exactly as fast as the founder. A sister-agreement measurement repairs the missing information: forty independent division records give 95% correct choice on a separated constant-treatment class, now throughout $\varepsilon \le 1/4$, and exact binomial inversion gives full-class rules with an explicit unresolved output. We distinguish calibration needed for a shifted decision threshold from calibration-free sign information, price correlated read errors as an additive covariance reserve, and express residual uncertainty in excess-survival-risk units. The source is synthetic: its rate ratios, instantaneous readout and independence assumptions are stated explicitly, and the conclusion is a statement about an observation protocol, not about a clinically effective regimen.

Contents

1 Introduction

1.1 Synopsis

A cell's descendants form a family, but a measurement may follow only one branch of that family. Two sources can produce exactly the same observed branch while coupling the states of sister cells differently. We show that this hidden dependence can reverse the preferred order of two specified treatment phases, so that no amount of data from the admitted protocol resolves the comparison. A paired sister measurement recovers the information the decision actually needs, and exact binomial calculations quantify when it supports a declaration. The examples are synthetic and demonstrate an information requirement, not a clinically effective regimen.

1.2 Inference about a cell and a decision about a family

Inferring proliferation and phenotype transitions, and choosing an intervention, are related but different tasks. Branching inference from cell numbers and fractions can exploit information well beyond mean trajectories [5], and optimal-dosing analysis studies interventions under a specified source and response law [6]. Here we ask a prior question: does an explicitly declared observation experiment determine a particular intervention comparison, when the source is not fully known?

The distinction matters at division. A retained daughter can have a completely known state distribution while its dependence with the omitted sister is unknown. Family extinction, however, uses *products* of daughter continuation probabilities, so it is sensitive to exactly that dependence. A schedule comparison can therefore turn on information erased by the observation protocol. This is not an assertion that all cell counts or all lineage trees are uninformative. We prove equality of one specific complete pruned-record law, including its timestamps, event labels, death and censoring records, and we say precisely which protocols are excluded.

1.3 What this paper establishes

The paper assembles a single source-specific package.

1. **An exact observation equivalence** (Theorem 1). Two bounded-rate finite-type sources with common founder law, life-event rates, observation channels and one-daughter transition marginals, but different sister couplings, produce identical complete retained-branch records on any bounded horizon, under every shared policy measurable from the retained history.
2. **An exact finite-duration treatment-order boundary** (Proposition 3). For the constructed source the difference of extinction probabilities between the chronological schedules AB and BA factors completely; the boundary $c_{\mathrm{p}}(e^{-h})$ is strictly positive and strictly increasing in the pulse length, with $c_{\mathrm{p}}(e^{-h}) = 7h/8 + O(h^2)$, and it leaves the feasible covariance range at a single point.
3. **Monotonicity in the descendant division rate** (Lemma 5). Family extinction is nonincreasing in the descendant division rate ε, under a supersolution condition that Lemma 6 verifies for the source at hand and, at every pulse length where the reversal exists, for an arbitrary descendant offspring law. This is new here and it halves every perturbation reserve at once, because the two schedules can only move in the same direction.
4. **A sensitivity reserve** (Proposition 9). A differential-inequality bound replaces the coupling bound "probability that an extra division occurs" by a bound that also charges how much such a division can matter. It is between eight and thirty-one times smaller than the occupation reserve it replaces, at no cost in generality: it holds for every $\varepsilon \geq 0$.
5. **A much larger certified range for the reversal** (Theorem 11, Theorem 14, Table 1). At $h = \log \frac{8}{7}$ the opposite preferences at $c = \mp 1/5$ now persist through $\varepsilon = 1/30$, and at $c = \mp 1/4$ through $\varepsilon = 1/16$. At the shorter pulse $h = \log \frac{50}{49}$ the reversal holds for *every* $\varepsilon \in [0, 1]$, that is, with descendants dividing exactly as fast as the founder.
6. **A uniform decision band** (Corollary 13). For every pulse length, every ε and every c, the preferred order is determined outside an explicit band of half-width proportional to ε and, for small h, to about $11h$. The previously available statement of this shape had half-width $16000\,h$ and applied only at $\varepsilon \leq 1$.
7. **Finite-sample decision rules for the repaired measurement** (Theorem 16, Proposition 18), with an unresolved output, exact binomial validity over the whole admitted source class, and calculated resolution probabilities.
8. **Calibration and error-model accounting** (§7), including a calibration-free sign rule that is now usable at finite pulse length, and an additive reserve (Proposition 22) that prices correlated

Table 1: What the new certificates change. "Previous" refers to the occupation-coupling certificate; every entry is for the same source, the same schedules and the same exact arithmetic. Factors are rounded down.

Statement	Previous	This paper	Factor
Reversal at $c = \mp\frac{1}{5}$, $h = \log\frac{8}{7}$	$\varepsilon \le 1/500$	$\varepsilon \le 1/30$	16.8
Reversal at $c = \mp\frac{1}{4}$, $h = \log\frac{8}{7}$	—	$\varepsilon \le 1/16$	—
Reversal throughout $0 \le \varepsilon \le 1$	$h \le 10^{-6}$	$h = \log\frac{50}{49} \approx .0202$	$> 2 \times 10^4$
Half-width of the undecided band in c	$16000\,h$	$\approx 11\,h\,\varepsilon$	$> 10^3$
Separated 40-division class	$\varepsilon \le .001$	$\varepsilon \le 1/4$	250
Full-class pulse rule at .95 resolution	$m = 1600$, $\varepsilon \le .001$	$m = 800$, $\varepsilon \le .01$	2 and 10
Matched direct-endpoint budget	7 814 272 roots	5 377 920 roots	1.45

Table 2: The four decision tasks, their thresholds and their certified descendant-division ranges. "Low c" and "high c" refer to the hidden founder covariance. $h = T = \log\frac{8}{7}$ unless stated.

Task	Threshold on c and which side wins	Certified ε	Where
Pulses AB vs BA	$c_{\mathrm{p}} \approx .13706$; low $c \Rightarrow AB$	$\le 1/30$ $(c = \mp\frac{1}{5})$	Thm 11
		$\le 1/16$ $(c = \mp\frac{1}{4})$	Thm 11
Short pulses, $h = \log\frac{50}{49}$	$c_{\mathrm{p}} \approx .01813$; low $c \Rightarrow AB$	≤ 1 $(c = \mp\frac{1}{4})$	Thm 14
Constant A vs B, separated	$c_* \approx .08432$; high $c \Rightarrow A$	$\le 1/4$	Thm 16
Either, full class with abstention	own threshold	$\le 1/100$	Prop 18

read errors instead of assuming them away.

Items 3–6 and 8 are new relative to the earlier treatment of this source; items 1, 2 and 7 are retained, with the numerical content recomputed. Table 1 summarises the quantitative change.

1.4 Sign conventions and a map of the decision tasks

Throughout, F_a is the probability that the family is *extinct* by the deadline under schedule a, and $R_a = 1 - F_a$ is the survival risk that the decision minimises. We write

$$D := F_{AB} - F_{BA} = -(R_{AB} - R_{BA}).$$

A positive D favours AB; a positive $R_{AB} - R_{BA}$ favours BA. The same convention is used for constant actions, with $D_C := F_A - F_B$, so a positive D_C favours A. All reported "gaps" are in units of absolute survival risk. The hidden coordinate is the founder sister covariance c; for binary state indicators with marginal $1/2$ the correlation is $4c$, so $c = \pm 1/5$ means correlation ± 0.8 and $c = \pm 1/4$ means correlation ± 1. These are strong sister associations, not perturbations of independence.

Four decision tasks appear, and they must not borrow one another's parameter ranges. Table 2 is the map.

1.5 Related work, and what is not claimed

Decision-focused experimental design, prediction identifiability, partial identification and minimax regret are established subjects [4, 10, 14]. Information lower bounds for identification from repeated sampling are standard [8, 15], and we use them as background rather than as new inequalities. Multi-type branching processes and their backward equations are classical [1, 7, 9]; the comparison arguments in §5 use the standard theory of quasimonotone differential inequalities [12, 16, 17].

Our contribution is a source-specific connection: an exact intervention-order boundary hidden from retained lineages, sharp control of the perturbation that transfers it to recurring division, and a paired measurement with quantitative decision guarantees. We make no exhaustive priority claim for the combined statement, and we do not claim to have established a clinically effective treatment, a generally optimal control policy, a universal failure of lineage tracing, or an optimal laboratory sample budget. Section 9 states the boundaries in detail.

2 The source, the endpoint and the admitted record

2.1 The branching source

There are three labelled types, P, S, T, and initially one P. The founder divides at rate one, without death or switching, and is replaced by an ordered pair of daughters with law

$$K_c = \begin{pmatrix} 1/4 + c & 1/4 - c \\ 1/4 - c & 1/4 + c \end{pmatrix}, \qquad -\tfrac{1}{4} \le c \le \tfrac{1}{4}. \tag{1}$$

The row and column index the two daughter states, not mother types. An S or T cell divides at rate $\varepsilon \ge 0$, using the same joint law, and dies at the action-dependent rates

$$(d_S^A, d_T^A) = (2, 4), \qquad (d_S^B, d_T^B) = (3, 3). \tag{2}$$

Different cells have independent event clocks and independent kernel draws conditional on their realised genealogy. Their future subtrees are independent *conditional on both daughter states*; those states themselves need not be independent. Replacing the joint offspring probabilities by products of marginals would delete the mechanism under study.

The source (1) has sister dependence but no marginal mother-to-daughter phenotype memory: every daughter marginal is $(1/2, 1/2)$. That separation is deliberate, and §8 proves that the conclusions survive a refinement with genuine marginal persistence. The response law (2), the preparation and the remaining source assumptions are known; c is the hidden coordinate, and we sometimes allow ε to remain unknown inside a certified interval. This is not extrapolation to an unmeasured action whose response is unrestricted.

2.2 Schedules and the endpoint

For a deterministic schedule a and horizon T the survival risk is

$$R_a(c, \varepsilon) = \mathbb{P}_P^a(N_T > 0), \qquad F_a = 1 - R_a. \tag{3}$$

The principal comparison is chronological AB against BA, each phase of length h, so $T = 2h$. These two schedules use exactly the same interventions for exactly the same durations and differ only in order; there is no exposure asymmetry to explain the difference away. Constant A against constant B at a common horizon supplies a second, easier decision task. Their equal sums of death hazards, $2 + 4 = 3 + 3$, are a mathematical actuator normalisation and are not an assertion that the two actions have equal dose, toxicity or resource cost.

2.3 The admitted observation protocol

One observation protocol follows a fresh founder until it divides or until an independent deadline of rate $\lambda > 0$. It records the elapsed time, which event occurred and, at a division, the state or marker

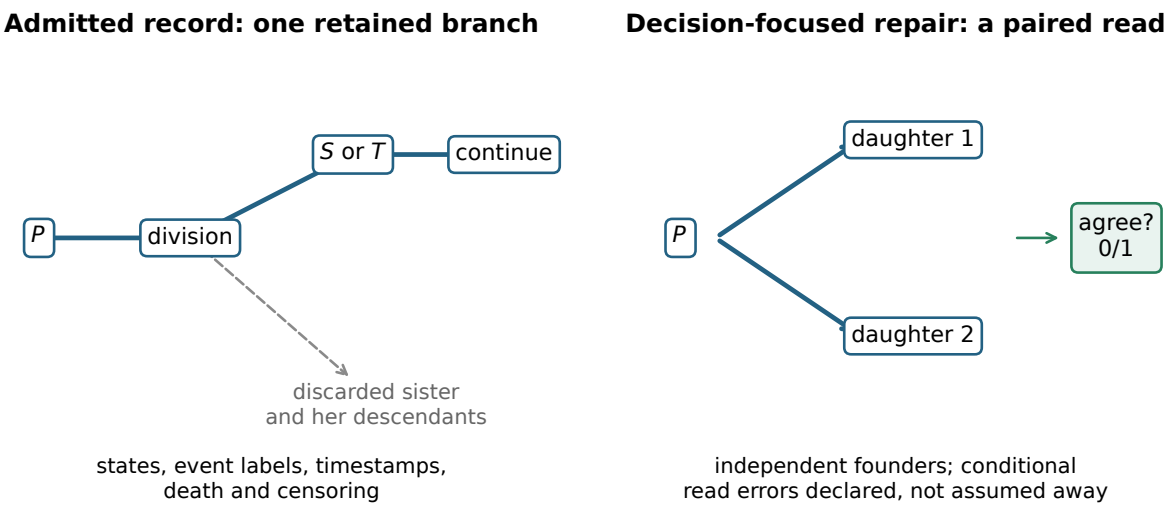

Figure 1: The admitted observation and its repair. In the admitted record the unretained sister and all of her descendants are absent, even though the record is otherwise complete: states, event labels, timestamps, death and censoring are all retained, for as long as the horizon allows. A paired read at the founder's division supplies one family-level feature — whether the two daughters agree — and that single bit is what the schedule comparison needs. No selection on subsequent survival is permitted in either protocol.

of a uniformly and independently chosen daughter. The record then ends. Its densities are

$$\lambda e^{-(1+\lambda)t} \quad \text{(deadline)}, \qquad \tfrac{1}{2} e^{-(1+\lambda)t} \quad \text{(division, either retained state)}. \tag{4}$$

A fixed emission channel acts on the state marginal. All of these densities are independent of c and of ε. Independent founders, not repeated observations of one lineage, are the independent sample units.

The richer protocol treated in Theorem 1 continues along the retained branch indefinitely: at every division the observer retains one daughter by a fresh uniform choice independent of the states, discards the sister and her entire subtree, and continues. Figure 1 shows this record and the paired repair studied in §6.

2.4 Why the backward equation is the source probability

Bounded per-cell division rates give domination by a Yule process of rate $\max(1, \varepsilon)$, so the chronological source is nonexplosive: the independent Yule holding-time expectations have divergent harmonic sum and summable centred variances, hence their sum diverges almost surely. The dominating population has finite integrated cell time on every bounded horizon, which also excludes accumulation of death events.

Define the offspring polynomial

$$G_c(s,t) = (\tfrac{1}{4} + c)(s^2 + t^2) + (\tfrac{1}{2} - 2c)\, st = \tfrac{1}{4}(s+t)^2 + c\,(s-t)^2. \tag{5}$$

Conditioning on the first event of a cell yields, for a constant action a, the field

$$\Phi_a^\varepsilon(p,s,t) = \begin{pmatrix} G_c(s,t) - p \\ d_S^a(1-s) + \varepsilon\,(G_c(s,t) - s) \\ d_T^a(1-t) + \varepsilon\,(G_c(s,t) - t) \end{pmatrix}. \tag{6}$$

The product structure in G_c is legitimate only after conditioning on both daughter states. The first-event integral equation on finite trees, followed by bounded convergence under nonexplosion, gives $v' = \Phi_a^\varepsilon(v)$, $v(0) = 0$, where the independent variable u is *time remaining to the deadline*. The field is polynomial, hence locally Lipschitz, and the unit cube is invariant because the field points inward on every face; uniqueness therefore identifies the solution with the actual extinction probabilities. For chronological phases, conditioning at the phase boundaries expresses the extinction vector as the flow composition $F_{a_1}^{t_1} \circ \cdots \circ F_{a_k}^{t_k}(0)$, in which the phases are evaluated in reverse chronological order. No population-size truncation is involved.

We write $A(u) = e^{-R_S(u)}$ and $B(u) = e^{-R_T(u)}$, where $R_i(u)$ is the death hazard of a type-i cell integrated over the last u units of the schedule. At $\varepsilon = 0$ the daughter coordinates decouple and

$$s^0(u) = 1 - A(u), \qquad t^0(u) = 1 - B(u). \tag{7}$$

3 Identical retained lineages, different family decisions

Theorem 1 (Retained-branch equivalence)**.** *Consider two bounded-rate finite-type branching sources with the same founder law, the same life-event rates, the same observation channels and the same one-daughter transition marginals under each action, but possibly different sister couplings. At every observed division retain one daughter by a fresh uniform choice independent of the states, discard the sister and her subtree, and continue only along the retained branch. Record every observed state, life event and time, including death and censoring. Then the complete records on a bounded horizon have identical laws. The conclusion remains true for shared action and stopping policies measurable from the retained history, and for shared external randomisation.*

Proof. Given the retained history, the next event clock and the event type have the same conditional law in both sources, because they depend only on the retained state and the currently chosen action. At a division the retained state is distributed according to the common daughter marginal, and the readout channel is common. Induction over events proves equality of the law of every finite-event record, including the terminal stopping or death flag. Equal histories induce equal actions and equal stopping decisions under a shared policy, and common external randomisation can be coupled identically.

Along a single branch the total event intensity is bounded by a constant, since the action menu is finite. The number of events by a bounded horizon is therefore dominated by a Poisson count, so truncating the record after n events has a tail probability tending to zero uniformly in both sources. Passing to the limit proves equality of the complete history laws. □

Remark 2. The theorem excludes: records of hidden family outcomes, records of both daughter states at any division, survivor-biased sampling, and any rule that chooses a daughter using information about the omitted sister. It establishes much more than equality of means, and it is stable under post-processing: because the two record laws are equal, *no* measurable statistic, reconstruction algorithm or learned model applied to those records can distinguish the sources. Collecting more time points on the pruned branch, or more independent pruned branches, cannot help either.

The algebraic reason for the phenomenon is the identity

$$G_{c+u}(s,t) - G_c(s,t) = u\,(s-t)^2, \tag{8}$$

immediate from the second form in (5). A change of c at fixed marginals is invisible to a retained branch, but it changes a family outcome whenever the two daughter continuation probabilities differ. The quantity the decision depends on is a product of continuation risks, not a marginal

state frequency. A schedule comparison then asks whether the two interventions weight the hidden direction $(s-t)^2$ differently.

4 An exact finite-duration treatment-order boundary

Set $x = e^{-h}$. For a daughter whose death rate is r in the first phase and s in the second, put

$$B_{r,s}(x) = \frac{x^{s+1} - x^{r+s}}{r-1} + \frac{x^2 - x^{s+1}}{s-1}, \qquad r, s \neq 1. \tag{9}$$

Indeed a daughter born at time $t < h$ survives to $2h$ with probability $e^{-r(h-t)-sh}$, while one born at $h \le t < 2h$ survives with probability $e^{-s(2h-t)}$; multiplying by the founder division density e^{-t} and integrating separately over the two intervals gives (9).

At $\varepsilon = 0$ the two daughter subtrees are independent given their states, so we may expand the joint extinction product conditional on the founder's division time. For first-phase rates (r_1, r_2) and second-phase rates (s_1, s_2),

$$\begin{aligned} F = 1 - x^2 - B_{r_1,s_1} - B_{r_2,s_2} \\ + (\tfrac{1}{4} + c)\big(B_{2r_1,2s_1} + B_{2r_2,2s_2}\big) + (\tfrac{1}{2} - 2c)\, B_{r_1+r_2,\, s_1+s_2}, \end{aligned} \tag{10}$$

the founder that never divides surviving and contributing no extinction mass. Substituting $(2,4;3,3)$ and $(3,3;2,4)$ and collecting powers of x gives the complete factorisation

$$D := F_{AB} - F_{BA} = \frac{x^2(x-1)^3}{210}\big\{4c\,\mathcal{A}(x) + (x-1)\mathcal{Q}(x)\big\}, \tag{11}$$

where

$$\begin{aligned} \mathcal{A}(x) &= (x^4 + x^3 + x^2 + x + 1)(3x^5 + 6x^4 + 9x^3 + 12x^2 + 8x + 4), \\ \mathcal{Q}(x) &= 3x^8 + 12x^7 + 30x^6 + 60x^5 + 98x^4 + 137x^3 + 170x^2 + 159x + 66. \end{aligned}$$

Equivalently, writing

$$s_0(x) = \frac{4\mathcal{A}(x)\, x^2(1-x)^3}{210} > 0, \tag{12}$$

we have the affine form $D = s_0(x)\,(c_{\mathrm{p}}(x) - c)$, with the boundary defined next.

Proposition 3 (Exact phase boundary)**.** *For $0 < x < 1$ and $\varepsilon = 0$, the schedule AB has strictly smaller survival risk if and only if $c < c_{\mathrm{p}}(x)$, where*

$$c_{\mathrm{p}}(x) = \frac{(1-x)\mathcal{Q}(x)}{4\mathcal{A}(x)}. \tag{13}$$

This threshold is strictly positive on $(0,1)$ and strictly increasing in the pulse length h. There is exactly one $x_0 \in (.80120, .80121)$ with $c_{\mathrm{p}}(x_0) = 1/4$. For $x < x_0$ the schedule AB is strictly preferred throughout the feasible range $-1/4 \le c \le 1/4$; for $x > x_0$ that range contains both preferences. Moreover $c_{\mathrm{p}}(e^{-h}) = \frac{7}{8}h + O(h^2)$.

Proof. $\mathcal{A}$ and $\mathcal{Q}$ have positive coefficients and the prefactor $x^2(x-1)^3/210$ in (11) is negative on $(0,1)$; hence $D > 0$, that is $R_{AB} < R_{BA}$, exactly when $4c\mathcal{A} + (x-1)\mathcal{Q} < 0$, i.e. below the threshold (13). Direct differentiation gives $c_{\mathrm{p}}'(x) = -P(x)/(4\mathcal{A}(x)^2)$ with

$$\begin{aligned} P(x) = 735x^{10} + 4410x^9 + 11655x^8 + 20580x^7 + 28560x^6 \\ + 31290x^5 + 26740x^4 + 17850x^3 + 9030x^2 + 3080x + 420, \end{aligned} \tag{14}$$

all of whose coefficients are positive; so c_{p} decreases strictly in x and therefore increases strictly in h. Its endpoint limits are $33/8$ at $x = 0$ and 0 at $x = 1$, so by continuity and strict monotonicity there is exactly one x_0 with $c_{\mathrm{p}}(x_0) = 1/4$; rational substitution at .80120 and .80121 isolates it, and monotonicity — not a floating-point root count — proves uniqueness. Finally $\mathcal{A}(1) = 210$ and $\mathcal{Q}(1) = 735$, so $c_{\mathrm{p}}(e^{-h}) = (735/840)h + O(h^2) = \frac{7}{8}h + O(h^2)$. □

Remark 4. The finite-duration decision threshold is *not* at $c = 0$. This is the single most consequential correction to the naive picture: the leading-order calculation of §5.5 sees only the cubic term $-4ch^3$ and therefore places the boundary at zero covariance, whereas at $h = \log\frac{8}{7}$ the boundary sits at $c_{\mathrm{p}} = 8967219299/65423564404 \approx .1370640591$, more than half way to the feasible ceiling. Consequently a measurement that determines only the *sign* of c does not determine the preferred order at finite pulse length; see §7. Proposition 3 is also a positive identification result of independent interest: for $x < x_0$ the hidden parameter remains completely unknown and yet its entire feasible range gives the same decision. What matters is whether the decision boundary meets the observationally equivalent set, not whether every parameter is recovered.

At $x = 7/8$, that is $h = \log\frac{8}{7} \approx .133531$, the exact values are

$$\begin{aligned} c_{\mathrm{p}} &= \tfrac{8967219299}{65423564404} \approx .1370640591, & s_0 &= \tfrac{114491237707}{32985348833280} \approx 3.47097 \times 10^{-3}, \\ D(-\tfrac{1}{4}) &= \tfrac{5635}{4194304} \approx 1.34349 \times 10^{-3}, & D(\tfrac{1}{4}) &= -\tfrac{25860351307}{65970697666560} \approx -3.91998 \times 10^{-4}, \\ D(-\tfrac{1}{5}) &= \tfrac{771817626293}{659706976665600} \approx 1.16994 \times 10^{-3}, & D(\tfrac{1}{5}) &= -\tfrac{48037425121}{219902325555200} \approx -2.18449 \times 10^{-4}. \end{aligned} \tag{15}$$

The contrast has opposite signs at $c = -1/5$ and $c = +1/5$, and again at $c = \mp 1/4$: negative covariance favours AB and positive covariance favours BA. Figure 2 displays the boundary and the contrast.

5 Recurring descendant division

The reference calculation of §4 is exact only at $\varepsilon = 0$. This section transfers it to $\varepsilon > 0$, and it is where the present treatment differs most from what was previously available. We prove two structural facts — a monotonicity and a sensitivity bound — and only then evaluate constants.

Throughout this section a is a schedule with phases of constant death rates, U is its total horizon, and R_S, R_T are the integrated death hazards over the last u units, as in §2.4.

5.1 Extinction is monotone in the descendant division rate

We allow the descendants to use an offspring law G that may differ from the founder's G_c; this costs nothing and is what makes §8 work. Thus G is any polynomial $w_{SS}s^2 + w_{ST}st + w_{TT}t^2$ with nonnegative coefficients summing to one, and (6) is read with G_c in the first coordinate and G in the other two. Everything below is driven by one inequality on the *reference* trajectory.

Lemma 5 (Monotonicity)**.** *Assume $d_S \le d_T$ in every phase, and assume the supersolution condition*

$$G(s^0(u), t^0(u)) \;\le\; s^0(u) \qquad \text{for all } u \in [0, U]. \tag{16}$$

Then for every $\varepsilon \ge 0$ the extinction vectors satisfy $v^\varepsilon(u) \le v^0(u)$ componentwise on $[0, U]$. In particular $F_a^\varepsilon \le F_a^0$: recurring descendant division can only reduce *family extinction.*

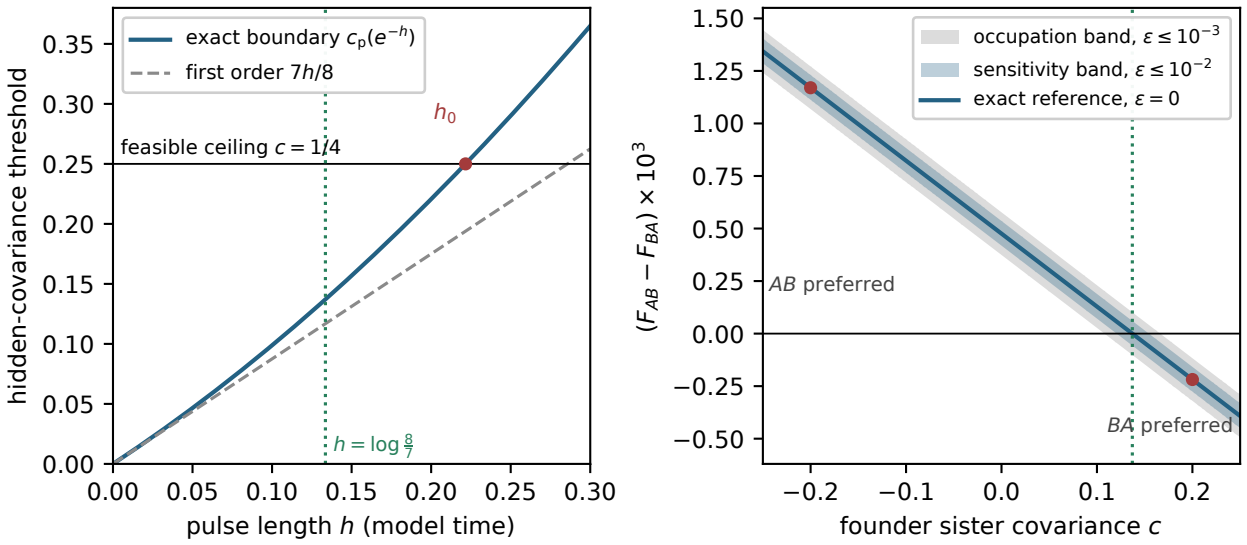

Figure 2: Left: the exact analytic threshold $c_{\mathrm{p}}(e^{-h})$ and its first-order approximation $7h/8$, with the feasible ceiling $c = 1/4$ and the unique crossing $h_0 = -\log x_0$. Below h_0 the schedule AB wins for every feasible covariance; above it both preferences occur, and the gap between the exact curve and $7h/8$ is what a sign-only measurement cannot see. Right: the exact reference contrast at $h = \log\frac{8}{7}$, with the two witnesses marked, enclosed by the proved sensitivity band at $\varepsilon \le 10^{-2}$ (Proposition 9) and, for comparison, by the occupation band at the ten times smaller $\varepsilon \le 10^{-3}$. Both bands are analytic enclosures, not solver-tolerance estimates. Negative contrast favours BA.

Proof. Since $d_S \le d_T$ we have $s^0 \le t^0$, so (16) says $G(s^0, t^0) \le \min(s^0, t^0)$ and both ε-terms of (6) are ≤ 0 at v^0. Hence $(v^0)' = \Phi_a^0(v^0) \ge \Phi_a^\varepsilon(v^0)$: the reference trajectory is a supersolution of the ε-system.

On $[0,1]^3$ each component of Φ_a^ε is nondecreasing in the other two coordinates, because $\partial_s G_c = 2(\frac{1}{4} + c)s + (\frac{1}{2} - 2c)t \ge 0$ and $\partial_t G_c \ge 0$ there (both coefficients $\frac{1}{4} + c$ and $\frac{1}{2} - 2c$ are nonnegative for $|c| \le \frac{1}{4}$), because the same holds for G, and because $\varepsilon \ge 0$. The system is therefore quasimonotone.

On the first phase v^ε solves $v' = \Phi_a^\varepsilon(v)$ with $v(0) = 0$ and v^0 is a supersolution of the same system with the same initial value; both remain in the invariant cube $[0,1]^3$. The comparison theorem for quasimonotone systems [16, §10] (see also [12, Ch. 3], [17, §III.12]) gives $v^\varepsilon \le v^0$ on that phase. At a phase boundary the ordering is inherited as initial data, (16) holds on the next phase as well, and the argument repeats. □

Condition (16) says that dividing now is never better for extinction than remaining a single cell of the more robust type. Two criteria cover every use we make of it, one specific to the founder's kernel and uniform in the pulse length, the other valid for an arbitrary descendant kernel on a restricted range of pulse lengths.

Lemma 6 (Criteria for (16))**.**

(a) If $G = G_c$ with $|c| \le \frac{1}{4}$ and if $d_S \le d_T \le 2d_S$ in every phase, then (16) holds, for every pulse length.

(b) For an arbitrary *descendant kernel G, (16) holds as soon as $(t^0(u))^2 \le s^0(u)$ for all u. For the two actions (2) and each of the schedules AB, BA, A, B with common phase length h, this in*

turn holds if and only if

$$2x^2 - x^6 \;\geq\; 1, \qquad \textit{that is} \qquad x \;\geq\; x_\flat := \left(\tfrac{\sqrt{5}-1}{2}\right)^{1/2} = .7861513\ldots \tag{17}$$

Since $x_\flat < x_0$, criterion (b) is available at every *pulse length at which both preferences occur in the feasible covariance range.*

Proof. (a) By the second form in (5) and $|c| \leq \frac{1}{4}$, $G_c(s,t) = \frac{1}{4}(s+t)^2 + c(s-t)^2 \leq \frac{1}{2}(s^2+t^2)$. The rate hypothesis integrates to $R_S \leq R_T \leq 2R_S$, so with A, B as in (7) we get $0 < B \leq A \leq 1$ and $A^2 = e^{-2R_S} \leq e^{-R_T} = B$; hence $A^2 + B^2 \leq B + B = 2B$ and

$$\tfrac{1}{2}((1-A)^2 + (1-B)^2) = 1 - (A+B) + \tfrac{1}{2}(A^2+B^2) \leq 1 - A = s^0.$$

(b) Every monomial of G is a product of two of s^0, t^0, each at most $\max(s^0,t^0) = t^0$, and the coefficients sum to one, so $G(s^0,t^0) \leq (t^0)^2$. The condition $(t^0)^2 \leq s^0$ reads $(1-B)^2 \leq 1-A$, i.e.

$$A \;\leq\; B\,(2-B). \tag{18}$$

We check (18) on each phase, writing $x = e^{-h}$. For a phase with $d_S = d_T$ one has $A = B$ and (18) reads $A \leq A(2-A)$, i.e. $A \leq 1$, which always holds; this settles the schedule B, the first backward phase of AB and the second backward phase of BA except for the accumulated prefactors, treated next. For the first backward phase of BA, and for the constant schedule A, the rates are $(2,4)$ from the deadline, so $A = e^{-2u}$ and $B = A^2$ with $A \in [x^2, 1]$; (18) becomes $2A - A^3 \geq 1$. The cubic $A^3 - 2A + 1 = (A-1)(A^2+A-1)$ shows $2A - A^3 \geq 1$ exactly for $A \geq \frac{\sqrt{5}-1}{2}$, and the smallest value of A on the phase is x^2, so the condition is $x^2 \geq \frac{\sqrt{5}-1}{2}$, i.e. (17). For the second backward phase of BA we have $A = x^2 z$, $B = x^4 z$ with $z = e^{-3(u-h)} \in [x^3, 1]$, and (18) becomes $2x^2 - x^6 z \geq 1$, whose worst case $z = 1$ is again (17). For the second backward phase of AB we have $A = x^3 w^2$, $B = x^3 w^4$ with $w = e^{-(u-h)} \in [x, 1]$, and (18) becomes $f(w) := 2w^2 - x^3 w^6 \geq 1$; since $f'(w) = 2w\,(2 - 3x^3 w^4)$ vanishes at most once on $(0,1]$ and f is there first increasing, the minimum of f over $[x,1]$ is attained at an endpoint, and $f(1) = 2 - x^3 \geq 1$ while $f(x) = 2x^2 - x^9 \geq 2x^2 - x^6 \geq 1$ by (17). Conversely the case $A = e^{-2u}$ above shows that (17) is necessary. Finally $x_\flat^2 = \frac{\sqrt{5}-1}{2} < .61804 < .64192 < x_0^2$, so $x_\flat < x_0$. □

Remark 7. Criterion (a) applies to our source with $(d_S^A, d_T^A) = (2,4)$ saturating $d_T \leq 2d_S$ exactly, and it is uniform in h; criterion (b) is sharp in x but indifferent to the descendant kernel. Both are sufficient conditions for (16), which is what the analysis actually uses.

5.2 A sensitivity reserve

Lemma 5 fixes the sign of the perturbation. The next proposition bounds its size, and does so more sharply than the obvious coupling argument.

Definition 8 (Two reserves)**.** For $i \in \{S, T\}$ put

$$L_i(u) = e^{-R_i(u)} \int_0^u e^{R_i(\tau)}\, d\tau, \qquad \sigma_i(u) = e^{-R_i(u)} \int_0^u (e^{R_i(\tau)} - 1) d\tau, \tag{19}$$

and define the *occupation reserve* and the *sensitivity reserve*

$$V_a = \int_0^U e^{-(U-u)} \left[L_S(u) + L_T(u)\right] du, \qquad W_a = \int_0^U e^{-(U-u)}\, t^0(u) \left[\sigma_S(u) + \sigma_T(u)\right] du. \tag{20}$$

Here $L_i(u)$ is the expected time a type-i daughter present at remaining time u stays alive, so V_a is the expected integrated live-daughter time of the reference process — the quantity a naive coupling produces. In σ_i that live time is weighted by $s_i^0(\tau) = 1 - e^{-R_i(\tau)}$, the probability that the daughter's own subtree would have died anyway, and in W_a the result is weighted again by $t^0(u) \le 1$. Both weights are at most one, so $W_a \le V_a$; the point of the proposition is that the weights are in fact small.

Proposition 9 (Sensitivity reserve)**.** *Assume $d_S \le d_T$ in every phase together with the supersolution condition (16). Then for every $c \in [-\frac{1}{4}, \frac{1}{4}]$ and every $\varepsilon \ge 0$,*

$$0 \;\le\; F_a^0 - F_a^\varepsilon \;\le\; \varepsilon\, W_a \;\le\; \varepsilon\, V_a. \tag{21}$$

Consequently, for two schedules a, b and $D^\varepsilon = F_a^\varepsilon - F_b^\varepsilon$,

$$D^0 - \varepsilon W_a \;\le\; D^\varepsilon \;\le\; D^0 + \varepsilon W_b. \tag{22}$$

Proof. Let $\Delta = v^0 - v^\varepsilon$, which is nonnegative by Lemma 5. Subtracting the two systems componentwise gives the exact identities

$$\Delta_s' = -d_S \Delta_s + \varepsilon\big(s^\varepsilon - G(s^\varepsilon, t^\varepsilon)\big), \tag{23}$$
$$\Delta_t' = -d_T \Delta_t + \varepsilon\big(t^\varepsilon - G(s^\varepsilon, t^\varepsilon)\big),$$
$$\Delta_p' = G_c(s^0, t^0) - G_c(s^\varepsilon, t^\varepsilon) - \Delta_p. \tag{24}$$

Since $G \ge 0$ on the cube and $s^\varepsilon \le s^0$, the forcing in (23) is at most $\varepsilon\, s^0$. Multiplying by the integrating factor e^{R_S} and using $\Delta_s(0) = 0$,

$$\Delta_s(u) \;\le\; \varepsilon \int_0^u e^{-(R_S(u) - R_S(\tau))} s^0(\tau)\, d\tau \;=\; \varepsilon\, \sigma_S(u),$$

the last equality because $e^{R_S(\tau)} s^0(\tau) = e^{R_S(\tau)} - 1$. Identically $\Delta_t(u) \le \varepsilon\, \sigma_T(u)$.

For (24), the segment joining $(s^\varepsilon, t^\varepsilon)$ to (s^0, t^0) lies in $[0, s^0] \times [0, t^0]$, so by the mean value theorem there is a point $(\tilde s, \tilde t)$ on it with $G_c(s^0, t^0) - G_c(s^\varepsilon, t^\varepsilon) = \partial_s G_c(\tilde s, \tilde t)\Delta_s + \partial_t G_c(\tilde s, \tilde t)\Delta_t$. Because $\tilde s \le s^0 \le t^0$ and $\tilde t \le t^0$, and because $2(\frac{1}{4} + c) + (\frac{1}{2} - 2c) = 1$,

$$\partial_s G_c(\tilde s, \tilde t) = 2(\tfrac{1}{4} + c)\tilde s + (\tfrac{1}{2} - 2c)\tilde t \le t^0,$$

and the same bound holds for $\partial_t G_c$. Hence $\Delta_p' \le t^0(\Delta_s + \Delta_t) - \Delta_p$, and with $\Delta_p(0) = 0$,

$$F_a^0 - F_a^\varepsilon = \Delta_p(U) \le \int_0^U e^{-(U-u)} t^0(u)\big(\Delta_s + \Delta_t\big)(u)\, du \le \varepsilon W_a.$$

The comparison $W_a \le V_a$ follows from $t^0 \le 1$ and $\sigma_i \le L_i$. For (22), subtract the two one-sided bounds and use that both deficits are nonnegative: $D^\varepsilon = D^0 - (F_a^0 - F_a^\varepsilon) + (F_b^0 - F_b^\varepsilon)$ with the first deficit in $[0, \varepsilon W_a]$ and the second in $[0, \varepsilon W_b]$. □

Remark 10. Two distinct improvements are combined in (22). First, Lemma 5 makes the band *one-sided in each direction*: without knowing the sign, one would have to write $|D^\varepsilon - D^0| \le \varepsilon(W_a + W_b)$, losing a factor close to two because $W_{AB} \approx W_{BA}$. Second, W_a itself is far below V_a. The two gains multiply, and neither costs any generality: (21) is valid for *all* $\varepsilon \ge 0$, not only for small ε.

Both reserves are elementary integrals of exponentials with integer rates, so for rational $x = e^{-h}$ each evaluates to an element of $\mathbb{Q} + \mathbb{Q}\, h$. Appendix B gives the closed forms; Table 3 gives the values at $x = 7/8$.

Table 3: The two reserves at $x = 7/8$. The occupation values reproduce the previously published constants exactly, which validates the evaluation; the sensitivity values are the ones used from here on. Here $h = \log\frac{8}{7}$, and the decimal column is an enclosure computed from a rational bracket for h of width below 10^{-30}.

Schedule	Occupation V_a	Sensitivity W_a	V_a/W_a
AB	$\frac{957367}{18874368} \approx .0507231$	$-\frac{6624078645151}{54975581388800} + \frac{10392084486973}{10995116277760}h \approx .00571652$	8.87
BA	$\frac{2552741}{50331648} \approx .0507184$	$-\frac{520588799010593}{3958241859993600} + \frac{6795050968489}{6597069766656}h \approx .00601852$	8.43
A	$\approx .0149943$	$-\frac{11540424829}{105696460800} + \frac{206906567}{251658240}h \approx .000601292$	24.9
B	$\approx .0149740$	$-\frac{79873}{1228800} + \frac{321391}{655360}h \approx .000483482$	31.0

5.3 The reversal under recurring division

Theorem 11 (Exposure-matched decision reversal)**.** *Take the source of §2 with $h = \log\frac{8}{7}$ and compare the chronological schedules AB and BA.*

1. *At $c = -1/5$ the schedule AB has strictly smaller survival risk, and at $c = +1/5$ the schedule BA does, uniformly for $0 \le \varepsilon \le 1/30$. The absolute risk gap is at least 1.78×10^{-5} throughout that range.*
2. *At $c = -1/4$ and $c = +1/4$ the same opposite preferences hold uniformly for $0 \le \varepsilon \le 1/16$, with absolute risk gap at least 1.58×10^{-5}.*

The exact suprema of the admissible ranges are

$$\varepsilon_\star(\tfrac{1}{5}) = \min\Big(\frac{D(-\frac{1}{5})}{W_{AB}},\ \frac{-D(\frac{1}{5})}{W_{BA}}\Big) \approx .0362961, \qquad \varepsilon_\star(\tfrac{1}{4}) \approx .0651318.$$

For each fixed ε the complete retained-branch laws of the two sources are identical, under every shared admissible retained-history policy.

Proof. At $\varepsilon = 0$ the exact evaluation (15) of the factored contrast (11) gives $D(-1/5) > 0 > D(1/5)$ and $D(-1/4) > 0 > D(1/4)$. Here the descendant kernel is the founder's, and $d_S \le d_T \le 2d_S$ holds for both actions, so Lemma 6(a) gives (16) and Proposition 9 applies with $a = AB$, $b = BA$, yielding (22). At $c = -1/5$, positivity of D^ε therefore follows from $D(-1/5) > \varepsilon W_{AB}$, and at $c = +1/5$ negativity follows from $-D(1/5) > \varepsilon W_{BA}$; the binding constraint is the second. Substituting the values of Table 3 and clearing denominators in exact rational arithmetic proves both claims at $\varepsilon = 1/30$, with the stated gap, and similarly at $\varepsilon = 1/16$ for $c = \mp 1/4$. Theorem 1 supplies the observation equality. □

Remark 12. The previously available certificate for the same statement used $|D^\varepsilon - D^0| \le \varepsilon(V_{AB} + V_{BA})$, which certifies only $\varepsilon \le 1/500$ and is exhausted near $\varepsilon \approx .0021534$. Its exhaustion was, correctly, never interpreted as evidence of a true sign crossing; the present theorem shows that it was indeed an artefact of the bound, by a factor of 16.8. Direct numerical integration of the backward system, which is a diagnostic and not a proof, shows no sign change anywhere in $0 \le \varepsilon \le 1$ at $h = \log\frac{8}{7}$.

The affine form $D^0 = s_0(c_\mathrm{p} - c)$ converts Proposition 9 directly into a decision rule that does not depend on choosing witnesses.

Corollary 13 (Uniform decision band)**.** *Put $\gamma_-(x) = W_{AB}/s_0(x)$ and $\gamma_+(x) = W_{BA}/s_0(x)$. For every pulse length, every $\varepsilon \ge 0$ and every $c \in [-\frac{1}{4}, \frac{1}{4}]$,*

$$c < c_\mathrm{p}(x) - \varepsilon\,\gamma_-(x) \implies AB \text{ strictly preferred}, \qquad c > c_\mathrm{p}(x) + \varepsilon\,\gamma_+(x) \implies BA \text{ strictly preferred}.$$

At $x = 7/8$, $\gamma_- \approx 1.64695$ *and* $\gamma_+ \approx 1.73396$. *As* h *decreases,* $\gamma_\pm/h$ *decreases towards about* 10.94*: the computed values of* γ_+/h *are* 12.99, 11.21, 11.07, 10.95, 10.94 *at* $h = \log\frac{8}{7}, \log\frac{50}{49}, \log\frac{100}{99}, \log\frac{1000}{999}, \log\frac{10^4}{10^4-1}$.

This should be compared with the legacy cubic estimate of §5.5, which yields a band of half-width $16000\,h$, valid only for $\varepsilon \le 1$ and centred at $c = 0$ rather than at c_{p}. The band above is centred correctly, shrinks linearly in ε, and has half-width about $11h\varepsilon$ — better by more than three orders of magnitude at $\varepsilon = 1$, and by correspondingly more for small ε.

5.4 Short pulses: the whole range of descendant division rates

The certified range $\varepsilon_\star$ grows without bound as the pulse shortens, because the contrast is of order h^3 while the reserve is of order h^4. There is therefore a finite pulse length at which the reversal is certified for every descendant division rate up to and including the founder's own rate.

Theorem 14 (Short pulses)**.** *Take* $h = \log\frac{50}{49} \approx .0202027$*, so that* $c_{\mathrm{p}} \approx .0181305$*. Then at* $c = -1/4$ *the schedule* AB *has strictly smaller survival risk, and at* $c = +1/4$ *the schedule* BA *does, uniformly for* every $\varepsilon \in [0, 1]$*. The certified supremum is* $\varepsilon_\star \approx 1.0238$*.*

Proof. Identical to Theorem 11, with $x = 49/50$. Exact evaluation gives $s_0 \approx 2.81790 \times 10^{-5}$, $W_{AB} \approx 5.90140 \times 10^{-6}$ and $W_{BA} \approx 6.38199 \times 10^{-6}$, whence $s_0(c_{\mathrm{p}} + \frac{1}{4}) - W_{AB} > 0$ and $s_0(\frac{1}{4} - c_{\mathrm{p}}) - W_{BA} > 0$; the smaller of the two residual gaps is about 1.519×10^{-7}. □

Remark 15. This is the statement that the previously available small-pulse bound (§5.5) was reaching for, and it improves it in pulse length by more than four orders of magnitude: the cubic remainder certifies $c = \pm 1/5$ at $\varepsilon \le 1$ only for $h \le 1/80000$, whereas $\log\frac{50}{49} \approx 2.02 \times 10^{-2}$. Under the illustrative scale of §8.2 — a mean founder waiting time of 24 hours — Theorem 14 concerns pulses of about 29 minutes each, with descendants dividing exactly as fast as the founder. That is a qualitatively different regime from a pulse of 86 milliseconds, and it removes the most serious restriction of the earlier treatment.

5.5 The legacy small-pulse remainder

For completeness we record the uniform estimate that the above supersedes, since it is proved by a different technique — a semigroup remainder rather than a differential inequality — and its symmetric form is occasionally convenient. For $0 \le \varepsilon \le 1$ the source satisfies

$$\left|F_{AB} - F_{BA} + 4ch^3\right| \le 64000\,h^4, \tag{25}$$

so $c > 16000h$ favours BA and $c < -16000h$ favours AB; at $c = \pm 1/5$ this requires $h \le 1/80000$. Appendix C proves the bound. Its cubic term comes from a founder division followed by two daughter deaths; daughter divisions first enter at fourth order, which is why the estimate is uniform in ε but so lossy in h. The shift of the true boundary away from zero is consistent with the fourth-order term $D(0) = \frac{7}{2}h^4 + O(h^5)$.

6 Measuring the hidden coordinate

Suppose a founder division yields a paired read of both daughters, with independent symmetric marker flips of known probability η. Coding the states as ± 1 shows that the agreement indicator has success probability

$$p(c) = \tfrac{1}{2} + 2kc, \qquad k = (1 - 2\eta)^2, \tag{26}$$

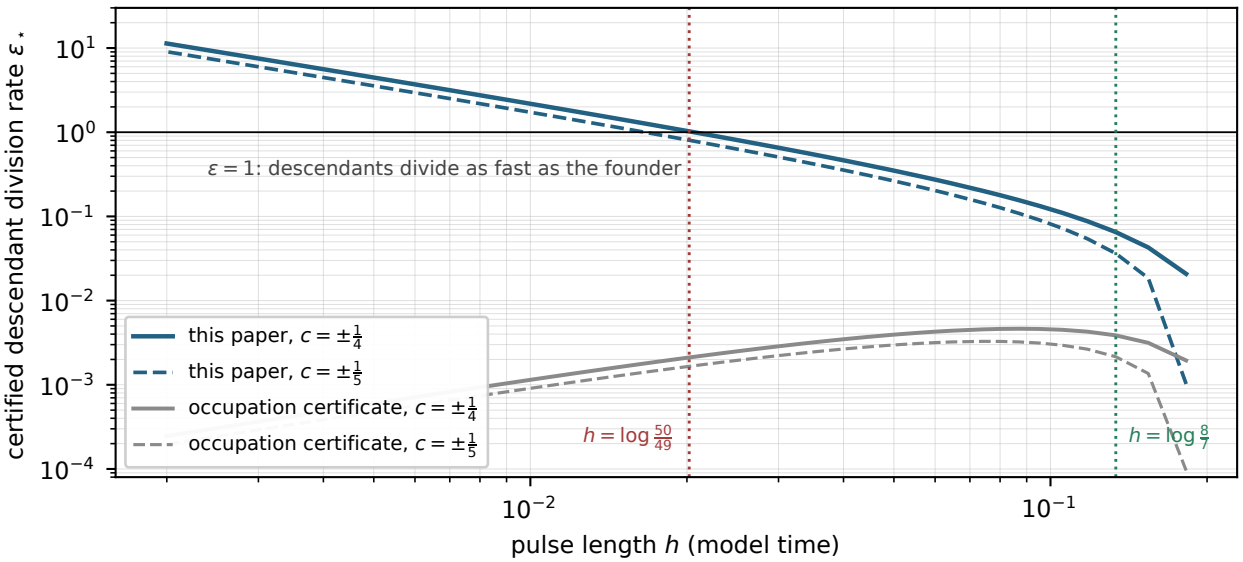

Figure 3: The certified descendant division rate $\varepsilon_\star$ as a function of pulse length, for witnesses at $c = \pm\frac{1}{4}$ and $c = \pm\frac{1}{5}$. Solid and dashed blue: the certificate of Proposition 9. Grey: the occupation certificate it replaces. Both families terminate on the right where c_{p} leaves the feasible range and the reversal genuinely ceases to exist (Proposition 3); the separation between blue and grey, by contrast, is pure slack in the older bound. The horizontal line at $\varepsilon = 1$ marks descendants dividing as fast as the founder; the blue curves cross it at pulse lengths near $\log\frac{50}{49}$, which is Theorem 14.

since each conditional marker mean is attenuated by $1 - 2\eta$ and conditional independence multiplies the two factors. Thus m independent accepted founder divisions give $X \sim \mathrm{Bin}(m, p(c))$. At $\eta = 1/10$ this is $p(c) = \frac{1}{2} + 1.28\,c$. We do not treat correlated read errors as independent calibration noise; see Proposition 22.

This single bit is exactly the family-level feature that Theorem 1 removes. It does not reconstruct the hidden subtree; it measures the one functional of the joint daughter law that the decision depends on.

6.1 Constant actions: a forty-division classification theorem

At $T = \log\frac{8}{7}$ and $\varepsilon = 0$, first-event integration under constant actions gives

$$F_A = a + sc, \quad F_B = b, \qquad a = \tfrac{1684141}{335544320}, \quad s = \tfrac{428807}{251658240}, \quad b = \tfrac{6767}{1310720}, \tag{27}$$

obtained by integrating $G_c(1 - e^{-2(T-t)}, 1 - e^{-4(T-t)})e^{-t}$ over $0 < t < T$, each exponential term contributing $(x - x^r)/(r - 1)$ with $x = 7/8$ and constant term $1 - x$. The contrast is $D_C = F_A - F_B = s\,(c - c_*)$ with

$$c_* = \frac{b - a}{s} = \frac{144633}{1715228} \approx .0843229, \tag{28}$$

so for constant actions *high* c favours A — the opposite orientation to the pulse comparison, and a reason to keep the two thresholds strictly separate.

Theorem 16 (Specified separated decision task)**.** *Assume $\eta = 1/10$, $0 \le \varepsilon \le 1/4$ and $|c - c_*| \ge 1/10$. From 40 independent founder-division agreement bits, choose A if $X \ge 25$ and B otherwise. Then the probability of choosing the better constant treatment exceeds .95 throughout this class.*

Proof. By Proposition 9 applied to the two one-phase schedules (Lemma 6(a) again supplies (16), since $2 \le 4 \le 2 \cdot 2$ and $3 \le 3 \le 2 \cdot 3$), $D_C^0 - \varepsilon W_A \le D_C^\varepsilon \le D_C^0 + \varepsilon W_B$. On the separated class $|D_C^0| \ge s/10 = 428807/2516582400 \approx 1.70382 \times 10^{-4}$, while $\frac{1}{4}\max(W_A, W_B) = \frac{1}{4}W_A \approx 1.50323 \times 10^{-4} < s/10$. Hence the sign of D_C^ε agrees with the sign of $c - c_*$ throughout the class, and classifying c relative to c_* determines the true recurring-source ranking.

For the classification itself, the extreme Bernoulli probabilities nearest the decision boundary are $p_L = \frac{1}{2} + 2k(c_* - \frac{1}{10}) = \frac{51449691}{107201750}$ and $p_U = \frac{1}{2} + 2k(c_* + \frac{1}{10}) = \frac{78893339}{107201750}$. Coupling Bernoulli trials through common uniform variables shows that binomial upper tails increase with p and lower tails decrease, so it suffices to check

$$\sum_{j=25}^{40} \binom{40}{j} p_L^j (1-p_L)^{40-j} < \tfrac{1}{20}, \qquad \sum_{j=0}^{24} \binom{40}{j} p_U^j (1-p_U)^{40-j} < \tfrac{1}{20}.$$

Exact rational summation proves both; the two error probabilities are approximately .0464141 and .0422845. □

Remark 17. The admissible range $\varepsilon \le 1/4$ is 250 times the range previously available for this theorem, and the exact supremum is $s/(10W_A) \approx .2834$. The decision rule and its two error probabilities are unchanged; only the robustness statement improves, because the reserve shrank. This theorem *assumes* separation: it does not certify that assumption from the data, does not provide a confidence interval, does not permit abstention, does not apply near the tie, and does not use the pulse threshold.

6.2 Full-class guarantees with an unresolved output

Without a separation assumption, the honest output has three values. An exact equal-tailed binomial confidence interval [3] for p inverts through (26) to a set of admissible c; mapping that set through the affine reference contrast and adding the coupling reserves to the two ends gives an interval for the true contrast, and a declaration is made only when that interval has one sign. Shared nuisance parameters must remain shared when the contrast is formed.

Equivalently, and this is the implementation we certify: for a threshold t, a reference slope magnitude $s_0 > 0$ and reserves $\rho_-, \rho_+ \ge 0$, define

$$p_- = \tfrac{1}{2} + 2k(t - \rho_-/s_0), \qquad p_+ = \tfrac{1}{2} + 2k(t + \rho_+/s_0). \tag{29}$$

When both lie in $[0,1]$, declare the low-c action if $\mathbb{P}_{p_-}(\mathrm{Bin}(m, p_-) \le X) \le \alpha/2$; declare the high-$c$ action if $\mathbb{P}_{p_+}(\mathrm{Bin}(m, p_+) \ge X) \le \alpha/2$; otherwise return *unresolved*. If an endpoint falls outside the parameter range the corresponding decision is vacuous or forced and can be read off the whole contrast range directly; in every example below both endpoints lie strictly inside $[0,1]$.

Proposition 18. *The above test makes an incorrect declaration with probability at most α throughout the admitted source class, without any separation assumption. For known $\eta = 1/10$, $\alpha = .05$ and $0 \le \varepsilon \le 1/100$:*

1. Constant actions, *$t = c_*$, $s_0 = s$, $\rho_- = \varepsilon W_B$, $\rho_+ = \varepsilon W_A$, $m = 200$: declare B at $X \le 106$ and A at $X \ge 137$. The resolution probability at $c = .2$ is $\approx .9909068$.*
2. Pulses *at $h = \log\frac{8}{7}$, $t = c_{\mathrm{p}}$, s_0 as in (12), $\rho_- = \varepsilon W_{AB}$, $\rho_+ = \varepsilon W_{BA}$, $m = 1600$: declare AB at $X \le 1009$ and BA at $X \ge 1153$. The resolution probability at $c = .2$ is $\approx .9994762$.*
3. *With the same pulse rule and $m = 800$: declare AB at $X \le 496$ and BA at $X \ge 584$, with resolution $\approx .9590464$ at $c = .2$.*

All three resolution probabilities exceed .95 by exact rational comparison, and at $c = -.2$ each exceeds .999999.

Proof. A low-c declaration can be wrong only if the low-c action is not in fact better, which by (22) requires the true c to be at least $t - \rho_-/s_0$ and hence $p \geq p_-$; the probability of the low-c acceptance region is then at most its value at p_-, which is at most $\alpha/2$ by monotonicity of the binomial lower tail. Symmetrically for a wrong high-c declaration, and the union bound gives α. The acceptance regions are integer tails, so their probability at any specified rational p is an exact finite sum of rationals; the cutoffs and resolution probabilities above are computed that way, with the displayed decimals descriptive and the comparisons against .95 exact integer inequalities. □

Remark 19. Three separate things improved at once here, and they should not be conflated. The reserve shrank, which narrows the unresolved band; the admissible ε grew tenfold; and the resulting budget fell. Concretely, the previously published rule at $\varepsilon \leq .001$ and $m = 1600$ declared AB at $X \leq 982$ and BA at $X \geq 1177$, with resolution .9721808 at $c = .2$. The new rule at ten times the ε range resolves more often with the same m, and item 3 shows that half the sample now suffices to clear the same .95 resolution target.

Remark 20. False-declaration probability is controlled *unconditionally*. This is not a claim that the conditional probability of error given a declaration is at most .05. "Resolution" means producing either declaration; correct resolution is a different event, although wrong declarations are bounded. Near a robust tie an unresolved output is expected and appropriate.

6.3 Effort, and a matched direct-endpoint benchmark

A founder divides before its exponential deadline with probability $y = 1/(1+\lambda)$, so collecting until m divisions yields i.i.d. marks and a negative-binomial expected root count m/y. The expected observed founder time per root is $1/(1+\lambda)$, so the expected total founder time is m; these are expectations, not fixed-budget guarantees. Imposing a cap C, with its flag retained, changes the yield to $y_C = \big(1 - e^{-(1+\lambda)C}\big)/(1+\lambda)$ while leaving the conditional bit law unchanged, because time and daughter draw are independent. With no timeout every immortal founder eventually divides, so exactly m roots suffice and the expected total founder time is still m: a timeout limits the length of an individual record but is not root-optimal in this idealised source. Calibration reads and the paired-read itself are additional costs, charged separately.

Direct endpoint randomisation needs no sister assay at all, and a fair comparison must therefore fix the same task. We use the separated class of Theorem 16, the same .05 error, known nuisance bounds and no abstention. Let W be the difference of extinction indicators under A and under B from two independent experimental roots. By Lemma 5, $F_A^{\varepsilon} \leq F_A^0$ and $F_B^{\varepsilon} \leq F_B^0$, so throughout the class $F_A, F_B \leq \max(a + s/4,\, b) < 3/500$ with no reserve needed at all; hence $\mathrm{Var}(W) \leq .012$ and $|W - \mathbb{E}W| \leq 2$. With n independent pairs Bernstein's inequality [2, §2.8] gives

$$\mathbb{P}(\text{wrong empirical sign}) \leq \exp\Big\{-\frac{ng^2}{2\big(.012 + 2g/3\big)}\Big\}, \qquad g = \tfrac{s}{10} - \varepsilon W_A. \tag{30}$$

At $\varepsilon \leq 1/100$ this is $g \approx 1.6438 \times 10^{-4}$ and the right-hand side falls below .05 at $n = 2\,688\,960$ per arm, certified using $\log 20 < 3$; that is $5\,377\,920$ roots in total, against $7\,814\,272$ under the previous reserve. It remains a sufficient, not a necessary, cost: no optimal sample-complexity separation is proved, and Table 4 does not assert an optimal experimental advantage. Direct outcomes retain the important advantage of not requiring the mechanistic and marker assumptions that the paired protocol needs.

Table 4: Different guarantees and their costs. Pair-record rows use known $\eta = .1$; expected root counts use $\lambda = 1$ without a cap. All observations are independent roots or independently accepted founder divisions. The ε column is the certified descendant-division range for that row.

Task	Guarantee and assumption	ε	Divisions	Roots
Constant classification	.95 correct; assumes $\lvert c - c_* \rvert \geq .1$	$\leq 1/4$	40	80 exp.
Constant abstention	full-class error $\leq .05$; resolution .9909 at $c = .2$	$\leq 1/100$	200	400 exp.
Pulse abstention	own threshold; error $\leq .05$; resolution .9995 at $c = .2$	$\leq 1/100$	1600	3200 exp.
Pulse abstention, lean	as above; resolution .9590 at $c = .2$	$\leq 1/100$	800	1600 exp.
Direct constant endpoint	same separated task; Bernstein sufficient bound	$\leq 1/100$	—	5 377 920

7 Calibration, correlated errors and residual ambiguity

7.1 What a sign needs, and what a threshold needs

Unknown but symmetric marker error does not always prevent a decision. If $k > 0$ then (26) gives

$$\operatorname{sign}(p - \tfrac{1}{2}) = \operatorname{sign}(c), \qquad |c| \geq \tfrac{1}{2}|p - \tfrac{1}{2}|, \tag{31}$$

both statements free of any knowledge of η. Combining (31) with Corollary 13 gives a genuinely calibration-free rule at *finite* pulse length, in two asymmetric halves.

Proposition 21 (Calibration-free declarations)**.** *Fix x and ε with $c_{\mathrm{p}}(x) - \varepsilon\gamma_-(x) > 0$, and let I be a $(1-\alpha)$ confidence interval for p built without knowledge of η.*
1. If $I \subset [0, \frac{1}{2})$ then AB is strictly preferred, with error probability at most α.
2. If $I \subset (\frac{1}{2} + 2(c_{\mathrm{p}}(x) + \varepsilon\gamma_+(x)), 1]$ then BA is strictly preferred, with error probability at most α.
At $x = 7/8$ and $\varepsilon \leq 1/100$ the hypothesis of the proposition holds, since $c_{\mathrm{p}} - \varepsilon\gamma_- \approx .12059 > 0$, and the second threshold is $\approx .80879$.

Proof. On the coverage event, in case 1 we get $p < \frac{1}{2}$, hence $c < 0$ by (31), hence $c < c_{\mathrm{p}} - \varepsilon\gamma_-$ by hypothesis, and Corollary 13 applies. In case 2 we get $p - \frac{1}{2} > 2(c_{\mathrm{p}} + \varepsilon\gamma_+)$, hence $|c| \geq \frac{1}{2}|p - \frac{1}{2}| > c_{\mathrm{p}} + \varepsilon\gamma_+$ with $c > 0$, and Corollary 13 applies again. □

The asymmetry is intrinsic and worth stating plainly. Certifying AB needs only the sign of c, because the threshold is on the other side of zero; certifying BA needs $|c|$ to exceed a positive threshold, and a calibration-free lower bound on $|c|$ loses a factor of two, so it demands a strong channel. At $\eta = .1$ one has $p(1/4) = .82 > .80879$, so case 2 is feasible, but only for nearly extremal covariance. There is no uniform sample bound as $k \to 0$: at pure noise no signal remains.

The previously available calibration-free rule required $I \subset (\frac{1}{2} + 32000h, 1]$, which at $h = \log\frac{8}{7}$ demands $p > 4273$ and is vacuous; it was usable only in the regime $h \leq 10^{-6}$. The rule above is usable at the pulse lengths the rest of the paper works with.

Calibration is genuinely needed to locate c against a *shifted* threshold. For example $(c, \eta) = (.05, 0)$ and $(.2, .25)$ both give $p = .6$, yet lie on opposite sides of c_*; no amount of agreement data resolves that pair. A calibration interval must then be projected jointly with the bit confidence set, and the union-bound coverage event for the two measurements is sufficient. Estimating the channel does not turn it into a known constant.

7.2 Asymmetric and correlated read channels

For a common asymmetric binary channel with $\mathbb{E}[Y_i \mid Z_i] = a + bZ_i$, $\mathbb{P}(Z_i = 1) = \frac{1}{2}$, $\mathrm{Cov}(Z_1, Z_2) = c$ and conditionally independent read errors,

$$\mathrm{Cov}(Y_1, Y_2) = a^2 + ab + b^2(\tfrac{1}{4} + c) - (a + \tfrac{b}{2})^2 = b^2 c. \tag{32}$$

The centred observed covariance therefore preserves the sign of c whenever $b \neq 0$. This requires the marginal read frequencies as well as the paired reads: an agreement bit alone does not identify the unknown centre.

Correlated read errors are a more serious threat, and the right response is to price them rather than to assume them away.

Proposition 22 (Additive reserve for correlated read errors)**.** *Suppose $\mathbb{E}[Y_i \mid Z_1, Z_2] = a + bZ_i$, so that the conditional mean of one read does not itself acquire a dependence on the sister state, and suppose the two read errors have conditional covariance $q(Z_1, Z_2)$. Then*

$$\mathrm{Cov}(Y_1, Y_2) = b^2 c + \mathbb{E}\big[q(Z_1, Z_2)\big]. \tag{33}$$

Consequently, if $|\mathbb{E}q| \leq \delta_{\mathrm{err}}$ is known, then a confidence interval for the observed covariance lying entirely above δ_{err} certifies $c > 0$, and one lying entirely below $-\delta_{\mathrm{err}}$ certifies $c < 0$.

Proof. By the law of total covariance, $\mathrm{Cov}(Y_1, Y_2) = \mathrm{Cov}\big(\mathbb{E}[Y_1 \mid Z], \mathbb{E}[Y_2 \mid Z]\big) + \mathbb{E}\big[\mathrm{Cov}(Y_1, Y_2 \mid Z)\big]$. The first term is $\mathrm{Cov}(a + bZ_1, a + bZ_2) = b^2 c$ and the second is $\mathbb{E}[q]$. □

The value of (33) is not that it rescues the assay — it does not — but that it shows precisely what must be measured: a bound on the *average* error covariance, not a claim that errors are independent. Note that a certified sign of c combines with Proposition 21 case 1 to certify AB; certifying BA still needs a magnitude.

7.3 Irreducible ambiguity, and decisions under it

On the two sources of Theorem 11 the retained-record laws are identical, so any possibly randomised binary selector has worst-case error at least $1/2$, for any number of independent records: if E denotes the event of choosing one of the actions, the two error probabilities sum to $1 - P_1(E) + P_2(E) = 1$. With unequal data laws the same two-point argument gives $\max \mathbb{P}(\text{error}) \geq (1 - \mathrm{TV}(P_1, P_2))/2$, and multiplying by a common gap lower bound converts this into an excess-risk lower bound. For the paired assay the one-root information is $y\ \mathrm{kl}(p_1 \| p_2)$ and independent-root Kullback–Leibler divergences add. These are standard testing tools [8, 15] specialised to this record law, not new information inequalities.

When the remaining gap makes a ranking costly, report the uncertainty in risk units. If an estimated contrast has absolute error at most e, then a wrong sign costs at most e in excess risk; $2e$ is the correct bound when the two risks were estimated separately each with error e, and the two conventions must be kept distinct. On a contrast interval $[L, U]$, choosing the first action has worst regret $\max(U, 0)$ and the second $\max(-L, 0)$; for $L < 0 < U$, randomising with probability $-L/(U - L)$ on the first action balances the two endpoint losses and minimises the worst regret at

$$\Gamma = \frac{(-L)\,U}{U - L} \tag{34}$$

[10, 14]. A declared tolerance δ can therefore justify an action even when the ranking is unresolved, and coverage failure adds at most α to the unconditional expected regret.

8 A controlled refinement, and what the numbers mean

8.1 Mother-dependent phenotype marginals

The source (1) deliberately separates sister coupling from marginal phenotype persistence, so that the theorem cannot be explained away as an artefact of inheritance. To include genuine persistence, keep the founder law (1) but give a type-i descendant the pair law

$$K_i = \begin{pmatrix} p_i^2 + \kappa_i & p_i(1-p_i) - \kappa_i \\ p_i(1-p_i) - \kappa_i & (1-p_i)^2 + \kappa_i \end{pmatrix}, \qquad p_S = .6,\ p_T = .4,\ \kappa_i \in \{-.1, .1\}, \tag{35}$$

whose entries are all nonnegative (they are, in the four cases, $\frac{13}{50}, \frac{17}{50}, \frac{17}{50}, \frac{3}{50}$ and permutations). Now a daughter matches her mother's type with probability .6. The founder covariance c and the descendant covariances κ_i are separate parameters and must be named separately.

Corollary 23. *Theorems 11 and 14, with the same ε ranges and the same gaps, remain valid when descendants use (35). More generally they remain valid when descendants use* any *type-dependent offspring law on $\{S, T\}$ whatsoever. Sources sharing the founder's daughter marginals and the event rates, but differing in their founder or descendant sister covariances, remain observationally equivalent under the retained-branch protocol.*

Proof. The descendant offspring law enters the backward system (6) only through the terms $\varepsilon\,(G - \cdot)$ in the two daughter coordinates, and Lemma 5 and Proposition 9 were stated for an arbitrary such G. The only hypothesis that refers to G is (16), and Lemma 6(b) supplies it for every kernel once $x \geq x_\flat$. Both theorems use $x = 7/8$ and $x = 49/50$, which exceed $x_\flat$, and indeed every pulse length at which the reversal exists does. The reference contrast at $\varepsilon = 0$ and the reserves W_a are computed from the $\varepsilon = 0$ trajectory alone and are therefore untouched, so the certificates transfer verbatim. The source remains nonexplosive. Finally Theorem 1 uses only the common daughter marginals, which (35) shares across κ_i, giving the observation claim. □

This is a robustness theorem, not a calibration of real inherited phenotypes. Its content is that sister dependence and mother-to-daughter memory are distinct model features and that only the former is at stake here: the descendant mechanism may be arbitrary, and the reversal is unaffected.

8.2 Physical scale

Let the physical founder division hazard be β. Model time is βt_{phys}; the descendant division hazard is $\varepsilon\beta$, the death hazards are the model values times β, and a pulse lasts h/β. Rescaling time changes all dimensional rates together and therefore cannot repair their ratios. Note also that an exponential mean division waiting time is not a population doubling time, and that the reciprocal descendant hazard is a latent division-clock scale, not the realised waiting time of a cell that can die first.

The earlier restriction to $\varepsilon \leq .001$ was a genuine weakness: a population whose founders and daughters have comparable division rates was simply outside the certified regime. Theorem 14 removes that restriction entirely, at the cost of a shorter pulse, and Theorem 11 relaxes it by more than an order of magnitude at the original pulse length. What remains restrictive is the specification of the response law (2) and the instantaneous nonperturbing readout, not the demographic ratio.

8.3 Effect sizes and multiple founders

At $c = .2$ and $\varepsilon = 0$,

$$F_{AB} = \tfrac{40043373837487}{1319413953331200} \approx .030349364, \qquad F_{BA} \approx .030567813.$$

Table 5: Model quantities under the illustrative mean founder waiting time $1/\beta = 24$ hours. The last four rows are the point of §5: the certified regime has moved from a latent descendant division clock of a thousand days to one of a few weeks, and — at a half-hour pulse — to no restriction at all.

Quantity	Physical reading	Status
$h = \log \frac{8}{7}$	3.20 h per pulse, 6.41 h in total	worked example
$h = \log \frac{50}{49}$	29.1 min per pulse, 58.2 min in total	Theorem 14
$h = 10^{-6}$	86.4 ms per pulse	legacy bound (25)
$\varepsilon = .001$	descendant division clock 24 000 h = 1000 d	previously certified
$\varepsilon = 1/30$	descendant division clock 720 h = 30 d	Theorem 11
$\varepsilon = 1/16$	descendant division clock 384 h = 16 d	Theorem 11
$\varepsilon = 1$	descendants divide as fast as the founder	Theorem 14

The absolute survival-risk improvement is $\approx 2.18449 \times 10^{-4}$, that is .0218 percentage points, while the relative increase in one-founder extinction is about .72%. These are different effect-size descriptions and should not be interchanged. Both schedules leave survival near 97%, so the example establishes a ranking mechanism, not effective eradication.

For n independent P founders, extinction is F_a^n and

$$|F_{AB}^n - F_{BA}^n| \le n \max(F_{AB}, F_{BA})^{n-1} |F_{AB} - F_{BA}|. \tag{36}$$

The ranking is preserved for every n, but the magnitude collapses: at $n = 2$ the extinction difference is $\approx 1.331 \times 10^{-5}$ and at $n = 10$ it is $\approx 4.930 \times 10^{-17}$, the individual extinction probabilities there being $\approx 6.630 \times 10^{-16}$ and $\approx 7.123 \times 10^{-16}$. Random preparations require averaging the full extinction product rather than substituting a mean count, and a shared environment requires conditional analysis before any family factorisation.

9 Scope and interpretation

9.1 What the obstruction is specific to

The obstruction is specific to a pruned observation and a family-level endpoint. Exact full-state molecular conservation removes it: if a mother contains (a, r) molecules and the observed daughter contains (i, j), the sister must contain $(a - i, r - j)$ under pure partitioning, so the joint law is determined by the one-daughter law. Our phenotype covariance family must not be treated as a free molecular covariance under those stronger constraints.

Founder immortality before division, exponential clocks, known intervention responses, instantaneous nonperturbing reads, conditionally independent marker errors and independent families are all substantive assumptions. Persistent marginals are covered by Corollary 23; arbitrary reporter delay, unobserved age structure, unknown responses and shared environments are not. Correlated reads are covered only to the extent that Proposition 22 is supplied with a defensible δ_{err}.

The endpoint (3) uses zero terminal extinction data. Eventual extinction after withdrawal is governed by the extinction vector of the withdrawal source and may erase these finite-time differences entirely. Clinical cure, relapse, detection and resistance establishment are different events. Nor does comparing this finite menu identify a globally optimal feedback policy.

Table 6: Protocol correspondence checks. A nonuniform but state-independent retention rule would admit an analogous theorem provided the induced daughter marginals still agree; a physically selected branch requires its own proof.

Candidate data situation	Required check
Microfluidic devices that follow one branch [18]	Is retention independent of state, or do pole age, position or survival select the branch? Include the relevant state variables and derive the actual kernel.
Destructive single-cell sampling	This is generally a snapshot protocol, not a complete retained history. Establish its own observation law; do not apply Theorem 1 verbatim.
Subsampled retrospective lineage reconstruction	Does the record retain shared ancestry, sister pairs or family sizes that reveal dependence? Derive the sampling law.
Full time-lapse lineage trees [11, 13]	Both sisters may be observed, so the exact pruning obstruction can vanish. Investigate the actual resolution and error model — this is the regime in which the repair of §6 is natural.

9.2 Before invoking Theorem 1 for a real protocol

Theorem 1 is a statement about one precisely specified observation rule. A platform may resemble it superficially and fail its hypotheses. Table 6 lists the checks we consider mandatory. None of the entries is an assertion that a named platform does or does not satisfy them; establishing that is a separate, empirical task that this paper does not perform.

9.3 The narrow conclusion

The mathematical conclusion is narrow and we state it without ornament: a specified missing family-level measurement can reverse a specified schedule comparison, this remains true when descendants divide at any rate up to the founder's own, and the uncertainty relevant to that comparison can be bounded without recovering every parameter. Observing how an individual cell behaves and choosing an intervention for its family are different inference problems, and the right measurement is dictated by the continuation products that appear in the decision. A small paired observation can resolve a comparison that complete retained-branch histories cannot resolve at all.

A Exact certificates

The boundary. The polynomial appearing in the derivative proof of Proposition 3 is (14); all eleven coefficients are positive. Expanding at $x = 1$ gives $\mathcal{A}(1) = 210$ and $\mathcal{Q}(1) = 735$. The equation $c_{\mathrm{p}} = 1/4$ is, after clearing denominators,

$$-6x^9 - 18x^8 - 36x^7 - 60x^6 - 76x^5 - 78x^4 - 66x^3 - 13x^2 + 81x + 62 = 0,$$

and rational substitution at $x = .80120$ and $x = .80121$ brackets its unique root in $(0, 1)$; monotonicity of c_{p}, not a floating-point root count, proves the uniqueness.

Reserves and margins. At $h = \log \frac{8}{7}$ the pulse threshold is $8967219299/65423564404$ and the affine slope magnitude is $114491237707/32985348833280$. The reserve polynomials of Table 3 are

evaluated using a rational bracket for $\log\frac{8}{7}$ of width below 10^{-30}, obtained from the alternating series for $\log(1+y)$ at $y=1/7$; for the coarser statements the bound $\log\frac{8}{7} < 134657/1008420$ suffices. The residual gaps quoted in Theorem 11 are

$$\min\Big(D(-\tfrac{1}{5}) - \tfrac{1}{30}W_{AB},\ -D(\tfrac{1}{5}) - \tfrac{1}{30}W_{BA}\Big) \approx 1.78315\times10^{-5},$$

and the corresponding quantity at $\varepsilon = 1/16$, $c = \mp 1/4$ is $\approx 1.58398\times10^{-5}$. For the record, the previously published occupation certificate at $\varepsilon = 1/500$ left the gap $154033465997/9895604649984000 > 1.55\times10^{-5}$, and at $\varepsilon = .001$ the gap $1157858798221/9895604649984000 \approx 1.17007\times10^{-4}$.

Binomial arithmetic. The exact binomial implementation represents $p = a/d$ and the jth term by the integer $\binom{m}{j}a^j(d-a)^{m-j}$ over the common denominator d^m; prefix and suffix sums then give every tail comparison without floating-point rounding, and the same routine evaluates the resolution probabilities of Proposition 18. The displayed decimals are descriptive; the comparisons against .95 and $\alpha/2$ are exact integer inequalities.

B The sensitivity reserve in closed form

The quantities of Definition 8 are elementary. On a phase with constant death rate r over $[0,u]$ measured from the deadline,

$$L(u) = \frac{1-e^{-ru}}{r}, \qquad \sigma(u) = \frac{e^{ru}-1-ru}{r}\,e^{-ru} = \frac{1-e^{-ru}-rue^{-ru}}{r}.$$

For a two-phase schedule one splits the inner integral in (19) at the phase boundary, so that σ_i is again a combination of 1, u and exponentials e^{-ku} with integer k; the outer integral in (20) then has an integrand of the form $\sum_k(\alpha_k+\beta_k u)e^{-ku}$, and each term integrates in closed form. With $x = e^{-h}$ rational the result lies in $\mathbb{Q}+\mathbb{Q}\,h$, the linear term in h arising exactly from the constant ($k=0$) contributions. The occupation reserve V_a has no h term, since its integrand contains no polynomial factor; this is why V_a is a plain rational in Table 3.

Writing out the two-phase occupation reserve for one daughter type with phase hazards (r,s) gives the compact expression

$$V(r,s) = \frac{(1-x)-(1-x^r)/r}{r-1} + \frac{x-x^r}{r-1}\cdot\frac{1-x^s}{s} + x\,\frac{(1-x)-(1-x^s)/s}{s-1},$$

whose three terms integrate the occupation before the switch, the survivors present at the switch, and the daughters born after it; summing over the two daughter types reproduces the V_a column of Table 3. The corresponding expression for W_a is longer but is produced by the same mechanical procedure, and the accompanying script emits it as an exact element of $\mathbb{Q}+\mathbb{Q}\,h$.

Two sanity checks are worth recording. First, the V_a column reproduces exactly the constants obtained previously by a completely different route — a first-added-event coupling with the bound "probability that an extra division occurs" — which confirms that V_a is the reserve that argument produces and that W_a is a strict refinement of it rather than a different quantity. Second, direct high-accuracy numerical integration of the backward system at $x = 7/8$ gives $\lim_{\varepsilon\downarrow0}(F_a^0 - F_a^\varepsilon)/\varepsilon \approx .0035$ for both schedules, against $W_a \approx .0057$ and $V_a \approx .0507$: the sensitivity reserve is within a factor of 1.7 of the true linear response, while the occupation reserve is off by a factor of 14.

C The legacy small-pulse remainder

We prove (25). Write $K(u,v) = \sum_{jk} K_c(j,k) u_j v_k$. Daughter extinction starts as $d_a t + O(t^2)$. Reverse propagation for chronological AB begins with B; integrating the founder's offspring term over the next backward phase gives the cubic coefficient $K(d_A, d_A)/3 + 4K(d_B, d_B)/3 + K(d_A, d_B)$. Interchanging A and B leaves the difference $K(d_B, d_B) - K(d_A, d_A) = -4c$, and the remaining terms first contribute at order four.

To bound the remainder, use the coefficient ℓ^1 norm on polynomials in (p, s, t). For $\varepsilon \le 1$ every component of (6) has norm at most 10, so if $L_a f = \nabla f \cdot \Phi_a^\varepsilon$ and $\deg f \le d$ then $\|L_a f\|_1 \le 10d\|f\|_1$. The first word $L_a p = G_c - p$ has norm 2, so every four-letter word has norm at most $2 \cdot 20 \cdot 30 \cdot 40 = 48000$. Let $T_a(t) f = f \circ F_a^t$; invariance of the cube makes this a sup-norm contraction, and $T_a(t) f = f + \int_0^t T_a(u) L_a f \, du$. Expanding $T_B(h) T_A(h) p$ to total degree three leaves five remainder terms with coefficients $h^4/(i!\,(4-i)!)$, $i = 0, \dots, 4$, multiplying words of length four evaluated on the cube; their coefficient sum is $2^4 h^4/4!$. Each schedule remainder is therefore at most $32000h^4$, and subtracting the two gives (25).

This bound proves a genuinely uniform statement in ε, which is why it was worth having, but it is extremely lossy in h and it is centred at $c = 0$ rather than at the true boundary c_{p}. Both defects are removed by Corollary 13 and Theorem 14, which are uniform in ε as well.

D Reproducibility and formal scope

What is proved how. The source construction, the identification of the backward system with the extinction probabilities, the retained-history equality of Theorem 1, the analytic integration of §4, the comparison and differential-inequality arguments of §5, the semigroup remainder of Appendix C and the concentration and testing arguments of §6 and §7 are conventional mathematical proofs, given in full in this paper. Every numerical constant that enters a proof step — the factored contrast (11), the positivity of the coefficients of P, the bracket for x_0, the four reserve polynomials, the residual gaps, the binomial tails and resolution probabilities, the Bernstein budget, and the feasibility of (35) — is recomputed in exact rational or symbolic arithmetic by the accompanying script, which performs 82 assertions and reports them individually. Where an irrational quantity is unavoidable, namely $h = \log \frac{8}{7}$ and $\log \frac{50}{49}$, it is enclosed in a rational interval of certified width and the enclosure is propagated. No floating-point value supports any claim in this paper; floating-point numbers appear only as descriptive decimals.

What is not claimed. We make no claim of machine-checked verification of the stochastic content. Any formalisation accompanying earlier versions of this material covered finite algebraic identities and rational margins only; a description such as "formally verified stochastic treatment theorem" would overstate it, and we do not use one. The numerical ODE integrations reported in Appendix B are diagnostics that motivated and cross-check the analytic bounds; no proof depends on them.

Reproduction. The submission contains `main.tex`, `refs.bib`, the three vector figures, and `scripts/` holding the exact-arithmetic engine, the certificate script and the figure generator. The certificate script prints a labelled dump of every constant quoted above, followed by its assertion log. It has no dependency beyond a standard Python distribution with `sympy`; the figure script additionally uses `numpy` and `matplotlib`, and the numerical cross-check in Appendix B uses `scipy`.

Provenance. A shorter and unpublished working note covering part of this material, together with an accompanying supplement, circulated privately in September 2026; the present paper supersedes both, and no argument here depends on either. The exact witnesses at $c = \pm 1/5$ and $h = \log \frac{8}{7}$, the factored contrast (11), the retained-branch equivalence and the forty-division rule originate in that note; the monotonicity lemma, the sensitivity reserve, the enlarged ε ranges, the uniform decision band, the short-pulse theorem, the calibration-free rule at finite pulse length and the correlated-error reserve are new here. No clinical parameters were estimated and no new data were collected.

References

[1] Krishna B. Athreya and Peter E. Ney. *Branching Processes*, volume 196 of *Die Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 1972. doi: 10.1007/978-3-642-65371-1. Reprinted by Dover Publications, 2004.

[2] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, Oxford, 2013. ISBN 978-0-19-953525-5. doi: 10.1093/acprof:oso/9780199535255.001.0001.

[3] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404.

[4] Robert J. Flassig, Iryna Migal, Esther van der Zalm, Liisa Rihko-Struckmann, and Kai Sundmacher. Rational selection of experimental readout and intervention sites for reducing uncertainties in computational model predictions. *BMC Bioinformatics*, 16(1):13, 2015. doi: 10.1186/s12859-014-0436-5.

[5] Einar Bjarki Gunnarsson, Jasmine Foo, and Kevin Leder. Statistical inference of the rates of cell proliferation and phenotypic switching in cancer. *Journal of Theoretical Biology*, 568: 111497, 2023. doi: 10.1016/j.jtbi.2023.111497.

[6] Einar Bjarki Gunnarsson, Benedikt Vilji Magnússon, and Jasmine Foo. Optimal dosing of anti-cancer treatment under drug-induced plasticity. *npj Systems Biology and Applications*, 11 (1):98, 2025. doi: 10.1038/s41540-025-00571-5. Preprint: arXiv:2412.16391.

[7] Theodore E. Harris. *The Theory of Branching Processes*. Springer, Berlin, Heidelberg, 1963. doi: 10.1007/978-3-642-51866-9.

[8] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42, 2016. URL https://jmlr.org/papers/v17/kaufman16a.html.

[9] Marek Kimmel and David E. Axelrod. *Branching Processes in Biology*, volume 19 of *Interdisciplinary Applied Mathematics*. Springer, New York, NY, second edition, 2015. doi: 10.1007/978-1-4939-1559-0.

[10] Charles F. Manski. Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246, 2004. doi: 10.1111/j.1468-0262.2004.00530.x.

[11] Oded Sandler, Sivan Pearl Mizrahi, Noga Weiss, Oded Agam, Itamar Simon, and Nathalie Q. Balaban. Lineage correlations of single cell division time as a probe of cell-cycle dynamics. *Nature*, 519(7544):468–471, 2015. doi: 10.1038/nature14318.

[12] Hal L. Smith. *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*, volume 41 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1995. doi: 10.1090/surv/041.

[13] Sabrina L. Spencer, Suzanne Gaudet, John G. Albeck, John M. Burke, and Peter K. Sorger. Non-genetic origins of cell-to-cell variability in TRAIL-induced apoptosis. *Nature*, 459(7245): 428–432, 2009. doi: 10.1038/nature08012.

[14] Jörg Stoye. Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics*, 166(1):138–156, 2012. doi: 10.1016/j.jeconom.2011.06.012.

[15] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, NY, 2009. ISBN 978-0-387-79051-0. doi: 10.1007/b13794.

[16] Wolfgang Walter. *Differential and Integral Inequalities*, volume 55 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, Berlin, Heidelberg, 1970. doi: 10.1007/978-3-642-86405-6.

[17] Wolfgang Walter. *Ordinary Differential Equations*, volume 182 of *Graduate Texts in Mathematics*. Springer, New York, NY, 1998. doi: 10.1007/978-1-4612-0601-9.

[18] Ping Wang, Lydia Robert, James Pelletier, Wei Lien Dang, François Taddei, Andrew Wright, and Suckjoon Jun. Robust growth of Escherichia coli. *Current Biology*, 20(12):1099–1103, 2010. doi: 10.1016/j.cub.2010.04.045.