Inherited-family dependence can reverse an equal-exposure treatment decision

21 September 2026

Abstract

We exhibit two seven-type branching models that agree on every type-resolved population mean, under every common deterministic schedule, and nevertheless prefer opposite orders for two equal-exposure interventions. The models differ only in the joint law of sister birth states: one partitions the mother's epigenetic marks complementarily between daughters, the other draws both daughters independently from the very same one-daughter marginal. All division, death, switching and mutation rates, the one-mutant-daughter convention and both exposures are identical. The objective is eventual total extinction under a declared common continuation, so the comparison is a probability, not a mean. Exact outward arithmetic places the four one-founder extinction probabilities in rational intervals of width 3×10^{-15} and proves opposite order margins exceeding 1.6×10^{-6}; an analytic curvature bound together with four additional enclosed flows certifies the reversal throughout an interval of killing rates of width 1.6×10^{-5}. We identify the mechanism: the two laws share their first-order response to rare acquisition, and their second-order responses differ by the transported sister covariance. We prove this correction has a definite sign, so the independent approximation is systematically optimistic about eradication, and we give an exact finite-amplitude response identity valid at the witness parameters. The operational cost of the approximation is the loss 1.7013×10^{-6} in extinction probability incurred by following its recommendation when the complementary law holds; a decision-preservation criterion says when a schedule ranking survives the approximation, and a scalar survival floor bounds what any predictable feedback course can achieve. Two continuation perturbations are quantified separately: residual acquisition up to 3×10^{-5} and finite clearing of duration 152, the latter at a common exposure cost 15.2 times the course itself. Arguments are conventional probability and ODE proofs, numerical claims are exact interval arithmetic under a written soundness proof, and a thin Lean root checks selected algebra but not the stochastic theorem. This is a synthetic counterexample with a small near-tie effect, not a calibrated treatment recommendation.

1 Introduction

1.1 A model-reduction question

Reversible tolerance can sustain a cell population long enough for durable resistance to arise, and branching models already connect phenotypic switching, permanent resistance and drug sequencing [1, 4]. Dose-response and induction laws shape optimal dosing [2], and the timing of a second strike has been analysed directly in terms of evolutionary rescue [6, 5]. Neither treatment sequencing nor the general observation that offspring dependence can affect extinction is new here.

Our question is narrower and is about model reduction. When a stochastic population model is simplified, a common and reasonable defence is that the simplification preserves the observable mean dynamics. Suppose it does so exactly: fix the molecular dynamics, all demographic rates,

the acquisition convention, every one-daughter marginal, the two actions and their exposures, and change only the *dependence between the two sister states produced at a division*. Every type-resolved mean trajectory is then unchanged, under every common schedule. Can the better treatment order for an exact extinction objective nevertheless change?

1.2 What we prove

It can. We construct a literal seven-type source (Section 2) with six molecular types carrying activating and repressive marks and one durable resistant type M. Under the *complementary* law J each old mark of the mother is allocated to one daughter or the other by a fair coin, so the pair conserves the mother's marks. Under the *independent* law I both daughters are drawn independently from exactly the marginal that J induces. The two laws have identical marginals and, as a consequence of the one-mutant-daughter convention, identical complete mean operators (Proposition 2.1).

Two actions A and B of ten time units each are compared in both chronological orders, with equal deadline and equal integrated rate increments. The objective is eventual total extinction under a declared post-course continuation which we prove clears every finite molecular population almost surely and creates no new M (Proposition 2.3). Theorem 4.1 states the counterexample: from any number of identical AA founders, J strictly prefers BA and I strictly prefers AB, with one-founder margins above 1.6×10^{-6}, and both preferences persist throughout a certified interval of the killing rate in B.

Section 3 explains why. Expanding in the acquisition scale ε, the two laws share the first-order coefficient exactly; the second-order coefficients differ by the transported covariance between sisters (equation (14)). We prove in Theorem 3.3 that this covariance is nonpositive coordinatewise, so the independent approximation over-estimates extinction probability at leading order — it is systematically optimistic about eradication. That sign alone does not decide a schedule comparison, because a reversal depends on the *difference* of two such errors; Proposition 3.4 supplies an exact finite-amplitude identity for that difference, and the schedule-gap identity (19) connects it to the certified numbers.

Section 5 gives the operational content: the direct cost of acting on the approximation, a criterion under which a schedule ranking is protected from it, the founder count at which the advantage is largest, and a minimax reading of the information limitation. Section 6 proves a survival floor valid for *every* admitted predictable feedback course, so the exhibited schedules are compared against a genuine feasibility limit. Section 7 quantifies how much the declared continuation can be perturbed before the conclusion is lost, and states the exposure cost of the finite version honestly.

1.3 What is not claimed

We do not optimise over schedules, identify the maximal reversal set, or refute any model-specific timing theorem proved under different assumptions. The effect is small and sits near a tie by construction: that is the point, since a decision procedure is exactly where a small modelling error becomes consequential. All rates are synthetic mathematical assumptions, not calibrated drug concentrations, doses or clinical cure probabilities, and nothing here is a treatment recommendation. The acquisition probabilities are stipulated coarse-grained per-division quantities and are not per-base mutation rates.

1.4 Evidence conventions

Claims carry one of four labels, used consistently and summarised in Appendix C. **C** is a conventional written proof. **E** is a computation in exact rational or outward dyadic interval arithmetic, rigorous under a conventional soundness proof given in Appendix B; the arithmetic is exact, but the integrator itself is not machine-verified. **K** is a strictly compiled Lean module; these cover selected algebra only and no Lean theorem asserts the probability-to-ODE bridge or the complete result. **N** is a floating-point diagnostic, used for illustration and never as an input to a certified sign.

2 The literal source and the extinction endpoint

2.1 Molecular types, marks and rates

A molecular cell carries two sites, each unmarked, activating or repressive. Its type is the pair (a, r) of activating and repressive mark counts with $a + r \le 2$; there are six such types, and we write $u = 2 - a - r$ for the number of unmarked sites.

Type	UU	UR	RR	AU	AR	AA
Activating marks a	0	0	0	1	1	2
Repressive marks r	0	1	2	0	1	0

Throughout, r denotes a repressive mark count. The extinction probability of the resistant family is written ρ, never r. The nonzero molecular jump rates are

$$(a, r) \to (a + 1, r) : u(.01 + a/2), \qquad (a, r) \to (a, r + 1) : u(.01 + r/2), \tag{1}$$

$$(a, r) \to (a - 1, r) : a(e + r/4), \qquad (a, r) \to (a, r - 1) : r(e + a/4). \tag{2}$$

Writing and erasing are each autocatalytic in their own mark and antagonised by the opposite mark. Vanishing prefactors suppress invalid boundary transitions. Let Q_e be the resulting conservative generator, with diagonal minus the outgoing sum. Division has constant rate $b = .1$; molecular death is $d_i = d_i^0 + c$, where $d_i^0 = .01$ for the protected types $a > r$, namely AU and AA, and $d_i^0 = .3$ otherwise. The erasure level e and the additional death c are the two action variables.

The seventh type M is durable resistance. It divides into two M cells at rate $\beta = .1$, dies at rate $\delta = .02$, never reverts, and is unaffected by either action. Its stipulated rates do not depend on the mark counts of the cell it came from. Its extinction probability from one cell is the minimal first-event fixed point

$$\rho = \delta/\beta = 1/5. \tag{3}$$

Note that RR means "two repressive marks" and is a molecular type; it is not the resistant type M.

2.2 Two sister laws with the same marginals

Under the complementary law J, each of the mother's a activating and r repressive marks is assigned independently to one of the two daughters by a fair coin:

$$(Y, Z) = \big((I, J_r),\, (a - I, r - J_r)\big), \qquad I \sim \mathrm{Bin}(a, 1/2), \quad J_r \sim \mathrm{Bin}(r, 1/2), \tag{4}$$

independently of each other. Unmarked sites refill each daughter to two sites. The pair conserves the mother's old marks exactly. For molecular test functions f, g write

$$D_{D,i}(f, g) = \mathbb{E}_{D,i}[f(Y)g(Z)], \qquad (\Lambda f)_i = \mathbb{E}_i[f(Y)].$$

The comparison law I draws both daughters independently from that same one-daughter marginal, so

$$D_{I,i}(f,g) = (\Lambda f)_i(\Lambda g)_i. \tag{5}$$

Both laws are exchangeable and both induce the same Λ, which is the only sense in which they agree at the level of a single daughter. Law I is an approximation: it generally does not conserve the mother's old marks across the pair. It is exactly the reduction obtained by recording the empirical one-daughter distribution and then assuming sisters are independent.

2.3 Division-associated acquisition of durable resistance

At each molecular division, with probability $\mu_i = \varepsilon\kappa_i$, exactly one of the two sampled daughters, chosen uniformly, is replaced by an M cell; otherwise both daughters retain their molecular types. We set $\kappa_i = .4$ for the protected types AU, AA and $\kappa_i = .1$ otherwise, and admit $0 \le \varepsilon \le 5/2$, which is exactly the range making every μ_i a probability. The mother is removed and exactly two daughters remain. Exchangeability gives the offspring probability generating function

$$G_{D,i}(x,y) = (1-\mu_i)D_{D,i}(x,x) + \mu_i\, y\, (\Lambda x)_i. \tag{6}$$

This convention is part of the model statement. Mutating each daughter independently, or adding a third cell, changes (6) and is a different source; we do not study those variants. We call the event *division-associated acquisition of durable resistance*, and emphasise that it is a stipulated coarse-grained probability per division, not a per-base DNA mutation probability.

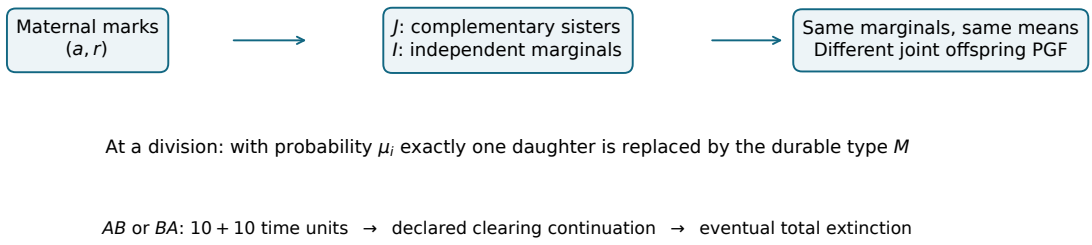

Figure 1: The comparison changes only sister dependence. Both laws use the same daughter marginals, the same demographic rates and the same one-mutant-daughter convention. The declared continuation is part of the extinction endpoint.

2.4 Backward field and the shared mean operator

The backward field on the invariant cube $[0,1]^7$ is

$$\Phi_{D,i}(x,y) = (Q_e x)_i + d_i(1-x_i) + b\{G_{D,i}(x,y) - x_i\}, \tag{7}$$

$$\Phi_M(y) = \delta(1-y) + \beta(y^2 - y). \tag{8}$$

By (3), $\Phi_M(\rho) = 0$: the mutant coordinate is stationary at ρ, for every action, so from the terminal value ρ it stays there and the molecular block decouples into a six-dimensional autonomous system with $y \equiv \rho$. We use this repeatedly.

Proposition 2.1 (Matched means; C, with E arithmetic)**.** *For molecular test weights f and mutant weight g, the mean operator of the augmented process has molecular coordinates*

$$(\mathcal{A}_\varepsilon(f,g))_i = (Q_e f)_i + b\{(2-\mu_i)(\Lambda f)_i + \mu_i g - f_i\} - d_i f_i \tag{9}$$

and mutant coordinate $(\beta-\delta)g$. It does not involve the joint daughter law. Consequently J and I have identical type-resolved mean trajectories under every common deterministic schedule, including the continuation. At $\varepsilon = 0$ the molecular block is $A = Q_e + b(2\Lambda - \mathrm{Id}) - \mathrm{diag}(d)$. Differentiating (7) at the all-one vector gives $D\Phi(\mathbf{1}) = \mathcal{A}_\varepsilon$, with the positive *mean-operator sign.*

Proof. Differentiate (6) at $(x,y) = (\mathbf{1},1)$ in the direction (f,g). Bilinearity and symmetry of D_D with $D_D(\mathbf{1},\mathbf{1}) = 1$ and $D_D(\mathbf{1},f) = \Lambda f$ give $(1-\mu_i)\cdot 2(\Lambda f)_i$ from the first term, and $\mu_i(g + (\Lambda f)_i)$ from the second; the sum is $(2-\mu_i)(\Lambda f)_i + \mu_i g$, which involves D_D only through Λ. Adding the transport, death and removal terms of (7) gives (9). Both laws induce the same Λ, so the two mean operators coincide. Forward mean row vectors obey $m' = m\mathcal{A}_\varepsilon$ and backward rewards are columns; the derivative at $\mathbf{1}$ carries a positive sign, and inserting a negative one corrupts every sensitivity equation below. The exact rational identity of the two 7×7 mean matrices is checked symbolically (E). □

Proposition 2.1 says only that the two laws agree law-to-law for a fixed common schedule. It does *not* say AB and BA have the same means, and it does not automatically extend to unconditional means under population-dependent feedback, where the realised action and the population can be correlated differently under the two laws.

2.5 Construction and the probability bridge

Proposition 2.2 (Construction and probability bridge; C)**.** *For bounded action ranges and a finite initial population the process is nonexplosive. For deterministic piecewise constant actions, the backward flow of (7)–(8) from a terminal payoff $g \in [0,1]^7$ equals the actual expected product of g over the terminal population. Chronological phases A then B give $q = F_A^{10}(F_B^{10}(g))$, where $F_\bullet^t$ is the time-t flow of the corresponding phase field. Fixed founders z give probability $\prod_i q_i^{z_i}$; a random preparation requires averaging that product over its actual law.*

Proof. Construct bounded per-cell marked Poisson proposals on a countable binary genealogical tree, with thinning for type changes, death and division. A Yule process of per-cell rate $b = .1$ dominates the population size. Its finite mean on compact intervals bounds the expected integrated population and hence the expected total proposal count, so localisation excludes explosion; this also covers predictable population feedback, used in Section 6. Under deterministic rates, future descendant clocks are independent conditional on daughter types. First-event conditioning gives (7): death contributes the empty product 1, a type change contributes the new continuation value, and a division contributes the joint offspring monomial, averaged over the joint birth law — this is the unique place where J and I differ. The probability of two events descended from one founder in time h is $O(h^2)$ uniformly. Bounded convergence and nonexplosion justify the integral identity and its differentiation. The polynomial field is Lipschitz on the cube, so uniqueness identifies its solution with the expectation. At a zero coordinate the field is nonnegative and at a one coordinate it is nonpositive, proving invariance of the cube. Conditioning at phase boundaries gives chronological composition. Independent founders multiply only under the stated deterministic schedule. □

2.6 The declared continuation and the endpoint

The objective is eventual total extinction, so the regime after the finite course is part of the statement. After the deadline T we set $e = .3$, $c = 0$ and $\mu = 0$, keeping the M rates fixed.

Proposition 2.3 (Clearing continuation and terminal payoff; C, with E arithmetic)**.** *Under this continuation every finite molecular population becomes extinct almost surely, and the continuation extinction vector is*

$$g = (1, 1, 1, 1, 1, 1, \rho) = (1, 1, 1, 1, 1, 1, 1/5). \tag{10}$$

Consequently eventual total survival is exactly survival of an M family seeded by time T.

Proof. Let A_0 be the mutation-free molecular mean matrix of the clearing source. For the positive weight $w = (14, 11, 10, 84, 43, 107)$, exact rational substitution gives

$$A_0 w \le -.09\, w, \qquad -.09\, w - A_0 w = (.2, .2, .2, 1.56, .78, 1.49) \quad \text{(E)}. \tag{11}$$

By the stopped generator identity and Gronwall, the expected weighted molecular population decays at rate at least .09. Its minimum positive weight is 10, so by Markov's inequality the probability of being nonempty at time t is at most $e^{-.09t}\langle w, z\rangle/10$ and tends to zero; hence almost sure molecular extinction from any finite configuration. There is no acquisition after T, so no new M appears, and M is autonomous with extinction probability ρ per initial cell by (3); nonexplosion identifies the limit over genealogical depths with extinction. Finiteness of the population at T and conditional independence of the autonomous M families then give (10) and the stated equivalence. □

Four natural events must be kept apart, since an M cell can appear and later die. Table 1 lists them with the terminal payoff computing each, together with floating diagnostics at the witness parameters which show they are genuinely different numbers.

Table 1: Distinct endpoints and their backward constructions. Probabilities are complementary-law AB diagnostics from one AA founder, evidence **N**; the certified table is Table 2.

Event	Backward construction	Diagnostic
Total survival past the deadline T	terminal 0, complement	.14318
Some M present at T	terminal $(\mathbf{1}, 0)$, complement	.03292
Some M ever appeared by T	killed molecular field, mutant term deleted	.03785
Eventual total survival (our objective)	terminal $(\mathbf{1}, \rho)$, complement	.03042

3 The dependence mechanism

The certified reversal of Section 4 does not rest on this section: the four enclosed flows prove it outright. But without a mechanism the result reads as an unexplained numerical coincidence, so we record what actually drives it. Everything here is a conventional proof, exact algebra, or a finite exact check.

Throughout this section the terminal molecular payoff is **1** and the mutant coordinate is frozen at its exact fixed point $y \equiv \rho$, which is the situation of (10).

3.1 The first-order response is common

Proposition 3.1 (Rare-acquisition expansion; C)**.** *On a fixed finite deterministic course, the molecular backward solution admits the expansion* $x_D = \mathbf{1} - \varepsilon h + \varepsilon^2 k_D + O(\varepsilon^3)$*, in which*

$$h' = Ah + b\kappa(1-\rho), \qquad h(0) = 0, \tag{12}$$

$$k_D' = Ak_D + bD_D(h,h) + b\kappa(2-\rho)\Lambda h, \qquad k_D(0) = 0, \tag{13}$$

with $A = Q_e + b(2\Lambda - \mathrm{Id}) - \mathrm{diag}(d)$ *and all products with* κ *taken coordinatewise. The first coefficient* h *is the same for both daughter laws.*

Proof. At $\varepsilon = 0$ we have $\mu \equiv 0$ and $x \equiv \mathbf{1}$ solves (7) exactly, so the expansion starts at $\mathbf{1}$; smooth dependence of a polynomial ODE on a parameter over a compact finite-horizon domain justifies it. Substituting $x = \mathbf{1} - \varepsilon h + \varepsilon^2 k + O(\varepsilon^3)$ into D_D and using bilinearity, symmetry and $D_D(\mathbf{1}, f) = \Lambda f$ gives

$$D_D(x,x) = 1 - 2\varepsilon\Lambda h + \varepsilon^2\{2\Lambda k + D_D(h,h)\} + O(\varepsilon^3),$$

while $(\Lambda x)_i = 1 - \varepsilon(\Lambda h)_i + O(\varepsilon^2)$ and $\mu_i = \varepsilon\kappa_i$. Collecting the $O(\varepsilon)$ and $O(\varepsilon^2)$ coefficients of (7) with $y = \rho$ gives (12) and (13); the detailed collection is in Appendix A. The ε-coefficient involves D_D only through Λ, which is common, so h is common. □

3.2 The second order is transported sister covariance

Both laws are exchangeable with the same marginal, so $\mathbb{E}_{J,i}[h(Y)] = \mathbb{E}_{J,i}[h(Z)] = (\Lambda h)_i$ and, by (5),

$$D_{I,i}(h,h) - D_{J,i}(h,h) = (\Lambda h)_i^2 - \mathbb{E}_{J,i}[h(Y)h(Z)] = -\operatorname{Cov}_{J,i}(h(Y), h(Z)).$$

Subtracting the two copies of (13) — the acquisition term $b\kappa(2-\rho)\Lambda h$ is common and cancels — gives the transport equation for $\Delta = k_I - k_J$:

$$\boxed{\Delta' = A\Delta - b\,\operatorname{Cov}_J(h(Y), h(Z)), \qquad \Delta(0) = 0,} \tag{14}$$

coordinatewise, and hence, with $U_A(t,s)$ the mean propagator of the phase-dependent A,

$$\Delta(t) = -\int_0^t U_A(t,s)\, b\, \operatorname{Cov}_{J,s}(h(Y), h(Z))\, ds. \tag{15}$$

Sensitivity states and propagators carry through phase boundaries without reset. These are finite-horizon expansions; they assert nothing uniformly in time or in founder number.

3.3 The correction has a definite sign

Order the molecular types by

$$(a,r) \preceq (a',r') \quad :\Longleftrightarrow \quad a \le a' \text{ and } r \ge r', \tag{16}$$

so that "more activating, fewer repressive" is higher. This is a partial order with least element RR and greatest element AA, and exactly six covering pairs: $RR \lessdot UR$, $UR \lessdot UU$, $UR \lessdot AR$, $UU \lessdot AU$, $AR \lessdot AU$ and $AU \lessdot AA$. Let $\mathcal{C}$ be the closed convex cone of vectors that are nonnegative and $\preceq$-increasing.

Lemma 3.2 (h stays in the monotone cone; C, with a finite E check)**.** *At the witness parameters,* $h(t) \in \mathcal{C}$ *for every* $t \in [0, 20]$ *and for both chronological orders.*

Proof. The forcing $b\kappa(1-\rho)$ of (12) is nonnegative, and κ is $\preceq$-increasing: a violation would need $\kappa_i = .4$ and $\kappa_j = .1$ with $i \preceq j$, hence $i \in \{AU, AA\}$, but every j above $AU = (1,0)$ has $a \geq 1$ and $r \leq 0$ and so lies in $\{AU, AA\}$, and the only j above $AA = (2,0)$ is AA (E). Since $h(0) = 0 \in \mathcal{C}$, it suffices that $\mathcal{C}$ is forward invariant for the affine field $f \mapsto Af + b\kappa(1-\rho)$ in each phase. The cone is the polyhedron cut out by $f_{RR} \geq 0$ together with $f_j - f_i \geq 0$ over the six covering pairs, so by Nagumo's subtangentiality criterion invariance holds as soon as each active constraint has nonnegative outward rate.

For the constraint $f_{RR} \geq 0$: if $f \in \mathcal{C}$ with $f_{RR} = 0$ then $f \geq 0$ because RR is least, and A is Metzler (E), so $(Af)_{RR} = A_{RR,RR} \cdot 0 + \sum_{k \neq RR} A_{RR,k} f_k \geq 0$; adding the nonnegative forcing keeps it nonnegative.

For a covering constraint $i \lessdot j$: if $f \in \mathcal{C}$ with $f_i = f_j$ we need $(Af)_j - (Af)_i + b(1-\rho)(\kappa_j - \kappa_i) \geq 0$, and the last term is nonnegative, so it suffices that $(Af)_j \geq (Af)_i$. The slice $\mathcal{C} \cap \{f_i = f_j\}$ is again a cone, and the layer-cake decomposition $f = \int_0^\infty \mathbf{1}_{\{f>t\}}\, dt$ writes any of its elements as a nonnegative combination of indicators of up-sets, each of which satisfies the same equality $\mathbf{1}_{\{f>t\}}(i) = \mathbf{1}_{\{f>t\}}(j)$ because $f_i = f_j$. The poset has exactly seven nonempty up-sets, so the requirement reduces to a finite list of inequalities. Exact rational evaluation confirms all of them: 32 inequalities for phase A and 32 for phase B, all nonnegative (E). Both phases therefore leave $\mathcal{C}$ invariant, and the two phases compose because the cone is the same. The explicit inequalities are in Appendix A. □

Theorem 3.3 (The independent approximation is optimistic at leading order; C)**.** *Under the hypothesis of Lemma 3.2,* $\mathrm{Cov}_{J,i}(h(Y), h(Z)) \leq 0$ *for every molecular type i and every time, and therefore $\Delta = k_I - k_J \geq 0$ coordinatewise. Ignoring sister dependence over-estimates the eventual extinction probability at leading order in ε, equivalently it under-estimates the probability of rescue.*

Proof. Fix a mother of type (a, r) and let $\xi_1, \ldots, \xi_a$ and $\zeta_1, \ldots, \zeta_r$ be the independent fair allocation bits of (4), so that $Y = (\sum_m \xi_m, \sum_m \zeta_m)$ and $Z = (a - \sum_m \xi_m,\ r - \sum_m \zeta_m)$. Introduce the independent bit vector

$$W = (\xi_1, \ldots, \xi_a,\ 1-\zeta_1, \ldots, 1-\zeta_r) \in \{0,1\}^{a+r},$$

whose coordinates are again independent and fair. Increasing a coordinate of W either increases the activating count of Y or decreases its repressive count, so Y is $\preceq$-increasing in W; simultaneously it decreases the activating count of Z or increases its repressive count, so Z is $\preceq$-decreasing in W. By Lemma 3.2, h is $\preceq$-increasing, so $h(Y)$ is a coordinatewise increasing and $h(Z)$ a coordinatewise decreasing function of the independent vector W. Harris' correlation inequality for product measures [3] gives $\mathrm{Cov}(h(Y), h(Z)) \leq 0$.

The forcing $-b\,\mathrm{Cov}_J(h(Y), h(Z))$ in (14) is therefore nonnegative. Each phase matrix A is Metzler, so its propagator $U_A(t,s)$ is entrywise nonnegative, and (15) with $\Delta(0) = 0$ gives $\Delta \geq 0$. Since $x_D = \mathbf{1} - \varepsilon h + \varepsilon^2 k_D + O(\varepsilon^3)$ with h common, the first term of $x_I - x_J$ that does not cancel is $\varepsilon^2 \Delta$, and it is nonnegative; by Propositions 2.2 and 2.3 x is the eventual extinction probability. □

Two qualifications matter. First, Theorem 3.3 is a statement about the fixed source above, with constant division rate and the stated monotone death and acquisition profiles; state-dependent division would break Lemma 3.2, whose finite check is specific to these two phase matrices. Second, and decisively for our purpose, a nonnegative dependence error *for each schedule separately* does not decide which schedule has the larger error. The reversal depends on that difference, which the next two results address.

3.4 An exact finite-amplitude response identity

The expansion above explains where the effect comes from but cannot certify it. At the witness value $\varepsilon = 1/10$ the best remainder bound we obtained from a generic third-derivative estimate exceeds 600, seven orders of magnitude larger than the 10^{-6} decision (Appendix A). That is a failed sufficient bound, not a false expansion. The following identity replaces it and is exact at every ε.

Proposition 3.4 (Exact dependence response; C)**.** *For a common terminal molecular payoff let $z = x_I - x_J$ and put*

$$B(t)f = Q_e f - (d+b)f + b(1-\mu)D_I(x_I + x_J, f) + b\mu\rho\Lambda f, \tag{17}$$

$$C(t) = b(1-\mu)\{D_I(x_J, x_J) - D_J(x_J, x_J)\} = -b(1-\mu)\,\mathrm{Cov}_J(x_J(Y), x_J(Z)). \tag{18}$$

Then $z' = B(t)z + C(t)$ with $z(0) = 0$, so $z(t) = \int_0^t U_B(t,s)C(s)\,ds$. The operator $B(t)$ is Metzler on the invariant cube, so U_B is nonnegative, and the response to $|C|$ bounds the absolute dependence error.

Proof. Subtract the two copies of (7) and add and subtract $D_I(x_J, x_J)$. Bilinearity and symmetry of D_I give $D_I(x_I, x_I) - D_I(x_J, x_J) = D_I(x_I + x_J, x_I - x_J)$, which is the third term of (17) applied to z; the transport, death, removal and mutant terms are linear in x and common to both laws, giving the remaining terms of (17). What is left over is exactly (18), and the second form of C follows as in (14). All coefficients of Q_e, D_I and Λ off the diagonal are nonnegative and $x_I + x_J \geq 0$ on the cube, so $B(t)$ is Metzler. □

Proposition 3.4 isolates the whole dependence error: the common acquisition terms enter only through the linear response B, and cancel from the forcing C. Note that (18) is again a sister covariance, now of the full finite-ε solution rather than of the first-order coefficient.

3.5 From the mechanism to the certified reversal

Write $E_s = q_I(s) - q_J(s)$ for the dependence error of schedule s and $G_D = q_D(AB) - q_D(BA)$ for the signed order gap of law D. Then, trivially but usefully,

$$G_J = G_I - (E_{AB} - E_{BA}). \tag{19}$$

This is the bridge between mechanism and certificate. Theorem 3.3 says both E_{AB} and E_{BA} are positive at leading order, and Proposition 3.4 says each is a transported covariance; because the covariance is generated during one pulse and then propagated through the other, the two schedules transport it differently, and their difference is what (19) feeds into the order gap. Section 4 certifies that this difference exceeds the order margin, which is exactly the condition for a reversal. Figure 2(c) displays the leading-order transport for the two schedules; at $\varepsilon = 1/10$ it reproduces each certified error, and their difference, to within about 18%, which is the expected size of the neglected $O(\varepsilon^3)$ term (N).

Neither the covariance sign alone nor any qualitative asymptotic argument proves the finite witness. The enclosed flows do.

4 A certified equal-exposure decision reversal

4.1 Actions and exposure accounting

Let $\varepsilon = 1/10$, so the acquisition probability per division is .04 for the protected types and .01 otherwise. Fix the two actions

$$A : (e, c) = (3/10,\ 0), \qquad B : (e, c) = (1/100,\ 120553/10^6),$$

each applied for ten time units and followed by the continuation of Proposition 2.3. Action A erases marks aggressively without extra killing; action B kills without erasing. Both orders have deadline $T = 20$, intrinsic-erasure increment exposure 2.9 above the baseline .01, and additional molecular-death exposure 1.20553. The M rates are unaffected by either action.

These are equal integrals of model rate increments. They are not equal measured drug masses, equal toxicity burdens, equal pharmacokinetic exposures or equal healthy-cell constraints; the equality is an assumption of the comparison, chosen so that no difference in the total amount of intervention can explain the result.

4.2 The counterexample

Theorem 4.1 (Matched means, opposite decisions; C with E arithmetic)**.** *For every positive integer n of identical AA founders, complementary inheritance J gives strictly larger eventual total-extinction probability under BA than under AB, while independent-marginal inheritance I gives strictly larger probability under AB than under BA. At the central parameters each one-founder order margin exceeds 1.6×10^{-6} in magnitude. Both preferences persist for every*

$$.120545 \le c_B \le .120561, \tag{20}$$

all other quantities held fixed.

Proof. The exact outward arithmetic of Appendix B places the four one-founder extinction probabilities in the rational intervals of Table 2, whence

$$-1.702 \times 10^{-6} < q_J(AB) - q_J(BA) < -1.701 \times 10^{-6}, \tag{21}$$

$$1.770 \times 10^{-6} < q_I(AB) - q_I(BA) < 1.771 \times 10^{-6}. \tag{22}$$

Propositions 2.2 and 2.3 identify these AA coordinates with eventual total extinction from one AA founder. All four numbers lie strictly in $(0, 1)$, so raising them to the power n preserves the strict orders. Proposition 2.1 gives equality of all mean trajectories between the two laws, schedule by schedule. The interval assertion is Lemma 4.2 below. □

4.3 Curvature turns four endpoint flows into a continuum

The interval statement cannot be obtained by sampling. We prove it from two enclosed endpoint gaps, the preserved central certificate and an analytic bound on the second derivative of the gap in the parameter. We write the argument out in full, because the intermediate step — how a secant becomes a uniform derivative bound — is what makes the conclusion valid off the sampled points.

Lemma 4.2 (Curvature and the certified interval; C with E arithmetic)**.** *Let $G_D(c) = q_D(AB; c) - q_D(BA; c)$ as a function of the additional death rate c in pulse B, all else fixed. Then $|G_D''| \le 6912$ uniformly on the nonnegative parameter range, and $G_J < 0 < G_I$ throughout (20).*

Table 2: Rational outward intervals for eventual total extinction from one AA founder, at the central parameters. The exact dyadic intervals in the reproduction files are tighter; each order-gap enclosure has width below 4×10^{-16}.

Law and order	Lower endpoint	Upper endpoint
J, AB	.969579866104740	.969579866104742
J, BA	.969581567445558	.969581567445561
I, AB	.969623237350068	.969623237350071
I, BA	.969621466763777	.969621466763781

Proof. Work with the six molecular coordinates, the mutant coordinate being frozen at ρ. Write the molecular backward field along the schedule as $f(t, x, c)$, let $\chi_B(t)$ indicate the phase carrying the varied parameter, and set $J_x = D_x f$, $u = \partial_c x$, $v = \partial_c^2 x$. Since c enters (7) only through $d_i = d_i^0 + c$ in the term $d_i(1 - x_i)$, we have $\partial_c f = \chi_B(\mathbf{1} - x)$, $\partial_c \partial_x f = -\chi_B \mathrm{Id}$ and $\partial_c^2 f = 0$. Differentiating the flow equation once and twice in c gives

$$u' = J_x u + \chi_B(t)(\mathbf{1} - x), \qquad u(0) = 0, \tag{23}$$

$$v' = J_x v + D_{xx} f[u, u] - 2\chi_B(t) u, \qquad v(0) = 0. \tag{24}$$

The factor two in (24) arises because $\partial_c J_x = D_{xx} f[u, \cdot] - \chi_B \mathrm{Id}$ contributes $-\chi_B u$ in addition to the $-\chi_B u$ obtained by differentiating the forcing of (23). Sensitivities pass through the phase boundary without resetting, and pulse durations are fixed, so there is no moving-boundary derivative.

Three uniform bounds on the cube are needed, all by direct inspection of (1)–(6) (E). The Jacobian J_x is Metzler with molecular row sums at most $b - d_i - b\mu_i(2 - \rho) \le b - d_i \le .09$, so its logarithmic sup norm is at most .09. The state Hessian satisfies $\|D_{xx} f[u, u]\|_\infty \le 2b\|u\|_\infty^2 = .2\|u\|_\infty^2$, since the only quadratic term is $b(1-\mu_i)D_i(x, x)$ and D_i has unit total mass. At the molecular vector $\mathbf{1}$ with mutant coordinate ρ the field equals $-b\mu_i(1 - \rho)$, of sup norm at most $b\varepsilon\kappa_{\max}(1 - \rho) = .0032$; this is *not* evaluated at the full seven-dimensional all-one vector, where the field vanishes.

Variation of constants applied to $\mathbf{1} - x$, then to (23), gives for $0 \le t \le 20$

$$\|\mathbf{1} - x(t)\|_\infty \le .0032\, t e^{.09t}, \qquad \|u(t)\|_\infty \le \int_0^t e^{.09(t-s)} .0032\, s e^{.09s} ds = .0016\, t^2 e^{.09t} < 6, \tag{25}$$

the last step using $e^{1.8} < 9$, which follows from $e < 3$ and hence $e^{1.8} < e^2 < 9$. Feeding (25) into (24),

$$\|v(20)\|_\infty \le \int_0^{20} e^{.09(20-s)} \{.2\|u(s)\|_\infty^2 + 2\|u(s)\|_\infty\} ds \le 20 \cdot 9\{.2 \cdot 36 + 2 \cdot 6\} = 3456. \tag{26}$$

Each G_D is a difference of two such schedule coordinates, so $|G_D''| \le 2 \cdot 3456 = 6912$ uniformly, and the bound holds on both phases and on the whole nonnegative parameter range.

Now the endpoints. Exact arithmetic on four additional flows with identical source equations gives

$$-5.114 \times 10^{-7} < G_J(.120545) < -5.113 \times 10^{-7}, \qquad 5.830 \times 10^{-7} < G_I(.120561) < 5.831 \times 10^{-7} \text{ (E).} \tag{27}$$

For law J form the secant of G_J between .120545 and the central value .120553, a span of 8×10^{-6}; combining (27) with (21) encloses that secant slope in $(-.1487475, -.1487474)$. By the mean value theorem the secant equals $G_J'(\xi)$ at some ξ strictly inside that span, hence inside (20). Every

point c of (20) satisfies $|c - \xi| \le 1.6 \times 10^{-5}$, so the curvature bound gives $|G'_J(c) - G'_J(\xi)| \le 6912 \cdot 1.6 \times 10^{-5} = .110592$ and therefore

$$G'_J(c) < -.1487474 + .110592 = -.03815 \qquad \text{for all } c \text{ in (20).}$$

The same computation for I, using the secant between .120553 and .120561 together with (22), gives $G'_I(c) < -.03785$ on the same interval. Both gaps are therefore strictly decreasing across (20). Since G_J is negative at the left endpoint and decreasing, it is negative throughout; since G_I is positive at the right endpoint and decreasing, it is positive throughout. No sampling and no root bracket are used. □

Remark 4.3. The interval (20) is an *inner* certified interval, not the maximal reversal set: it is bounded by where we placed endpoint flows. Its width 1.6×10^{-5} is 8000 times the width of the elementary Lipschitz neighbourhood (2×10^{-9}) available before this argument, which measures an improvement in the certificate and not an amplification of any biological effect. The margins at the endpoints are correspondingly smaller than the central margins. Determining the maximal set would require either many more endpoint flows or a validated root enclosure; neither is needed here.

4.4 The approximation error against the decision margin

Ignoring sister dependence changes the extinction coordinate by

$$E_{AB} \approx 4.3371245 \times 10^{-5}, \qquad E_{BA} \approx 3.9899318 \times 10^{-5}, \qquad E_{AB} - E_{BA} \approx 3.4719271 \times 10^{-6}, \tag{28}$$

all three enclosed exactly (E). Both absolute errors are small — about four parts in 10^5 of a probability close to .97 — and they have the sign predicted by Theorem 3.3. But their *difference* exceeds the approximate order margin 1.771×10^{-6}, and by (19) that is precisely what flips the sign of the gap. Absolute probability accuracy and decision accuracy are therefore different things. This result isolates decision vulnerability near a tie; it does not claim a large population-level benefit.

5 What the approximation costs a decision maker

5.1 The direct misspecification cost

The cleanest operational quantity involves no hypothetical randomisation over models. Suppose the complementary law J is the true source and a modeller, having validated the means, adopts the independent approximation and follows its recommendation AB. The loss relative to the recommendation BA that J would have given is

$$R_{I\to J} = q_J(BA) - q_J(AB) = m_J \in (1.70134081854 \times 10^{-6},\ 1.70134081866 \times 10^{-6}) \ \text{(E)}. \tag{29}$$

This is simultaneously lost extinction probability and added eventual-survival probability. We write $m_J = q_J(BA) - q_J(AB) > 0$ and $m_I = q_I(AB) - q_I(BA) > 0$ for the two order margins throughout.

5.2 When a schedule ranking is protected

Corollary 5.1 (A replacement for unqualified mean-based decisions; C)**.** *If two schedules admit valid bounds $|q_J(s) - q_I(s)| \le E_s$ and $q_I(A) - q_I(B) > E_A + E_B$, then $q_J(A) > q_J(B)$. Signed bounds on $E_A - E_B$, when available, are sharper and may be used instead. Conversely, no decision rule using only the complete type-resolved mean trajectories can always select the better of our two schedules for both inheritance laws.*

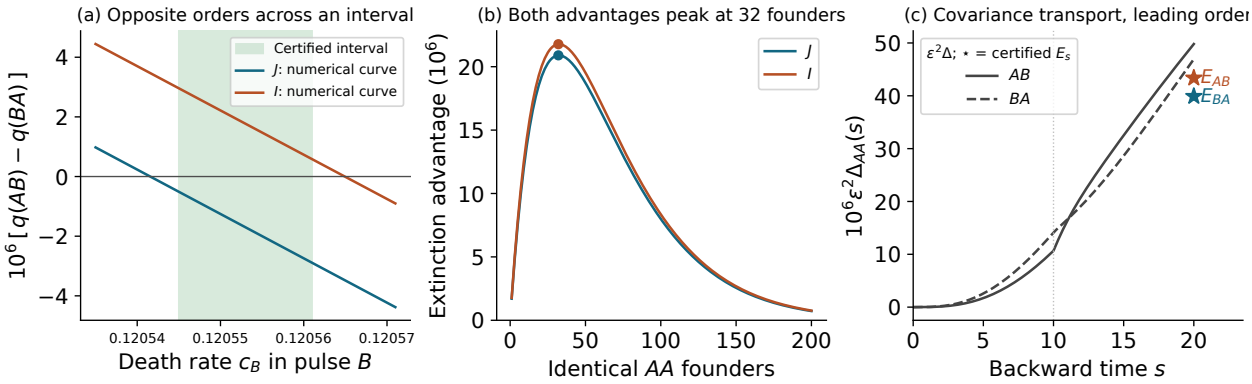

Figure 2: (a) The shaded interval is certified by two endpoint enclosures, the preserved central certificate and the analytic curvature bound of Lemma 4.2; the smooth curves are floating-point diagnostics and their zero crossings, near .1205416 and .1205649, are neither certified root brackets nor a claim about a maximal reversal region. (b) Curves use midpoints of the enclosed one-founder probabilities; the unique maxima at 32 founders are established by exact rational increment signs, not by the plotted samples. (c) Leading-order dependence error $\varepsilon^2\Delta_{AA}$ transported backward from the deadline, from (14), with the phase boundary at $s = 10$; stars mark the certified finite-ε values E_s of (28). The curves cross: the schedules differ in how the covariance forcing generated in one pulse is propagated through the other.

Proof. The first assertion is the triangle inequality. For the second, the mean data are identical for the two models by Proposition 2.1, so a rule using only those data returns the same answer in both; Theorem 4.1 requires opposite answers. □

The worked instance is instructive. The independent-model central margin lies in $(1.770, 1.771) \times 10^{-6}$ and the paired correction $E_{AB} - E_{BA}$ in $(3.471, 3.473) \times 10^{-6}$, so subtracting the signed correction certifies the complementary preference for BA. Had only a symmetric bound of $\pm 4 \times 10^{-6}$ on that correction been available, the corrected margin interval would contain zero and the defensible output would be "order unresolved at this approximation-error level" (E). This concerns model approximation error and is not an empirical confidence interval; empirical uncertainty would have to be modelled separately.

For a fixed heterogeneous preparation, telescoping the product of coordinates gives $|P_J - P_I| \le \sum_i z_i |q_{J,i} - q_{I,i}|$. An ordering for identical AA founders proves no ordering for an arbitrary mixture.

5.3 Founder scaling

Proposition 5.2 (Unique founder maximum; C with E arithmetic)**.** *Fix a law and let $0 < l < h < 1$ be its worse and better one-founder extinction probabilities. The advantage $A_n = h^n - l^n$ is positive, tends to zero, and satisfies $A_n \le n(h-l)h^{n-1}$. Its increment $A_{n+1} - A_n = l^n(1-l) - h^n(1-h)$ changes sign at most once. For both laws the increment is positive at $n = 31$ and negative at $n = 32$, so both advantages attain their unique integer maximum at 32 identical AA founders, with values in $(2.08950953, 2.08950954) \times 10^{-5}$ for J and $(2.17745028, 2.17745029) \times 10^{-5}$ for I.*

Proof. Positivity and the factorisation $A_n = (h-l)\sum_{j<n} h^{n-1-j}l^j \le n(h-l)h^{n-1}$ are immediate, as is $A_n \to 0$. Dividing the increment by h^n gives $(l/h)^n(1-l) - (1-h)$, and $(l/h)^n$ strictly decreases, so the increment changes sign at most once. The two increment signs are evaluated in

exact rational arithmetic from the endpoints of Table 2 (E), which avoids selecting an integer by rounded logarithms. □

A larger founder count therefore does not amplify the effect indefinitely; it changes achieved risk without changing the preferred schedule (Figure 2(b)).

5.4 A minimax reading of the information limitation

The next proposition is a statement about a class of models that the permitted information cannot separate. It is secondary to (29) and we state it as such.

Proposition 5.3 (Regret under mean-indistinguishable information; C with E arithmetic)**.** *A randomised decision rule that sees only information common to the two models, including their complete type-resolved mean trajectories, has worst-case expected extinction regret at least*

$$R_* = \frac{m_J m_I}{m_J + m_I} \in (8.67636512 \times 10^{-7},\ 8.67636514 \times 10^{-7}),$$

and the bound is attained in this two-model problem by choosing AB with probability $m_I/(m_J + m_I)$.

Proof. Such a rule selects AB with the same probability p under both laws. Its regrets are pm_J and $(1-p)m_I$. If $p \geq m_I/(m_J + m_I)$ the first is at least R_*, otherwise the second is; the displayed probability equalises them. The function $(s,t) \mapsto st/(s+t)$ increases in each positive argument, so substituting the endpoints of the certified intervals for m_J and m_I gives the enclosure (E). The randomisation selects an entire deterministic course independently of the population; the claim does not restrict a rule that observes sister pairs or any other information distinguishing the two laws. □

Remark 5.4. The premise is an information restriction, not a physical randomisation of nature. If the conservative law J is already known to hold, the two-world premise does not apply and the relevant loss is the deterministic (29), which is about twice R_*. Halving the number by randomisation does not reduce the cost of the deterministic mistake under J. Nor is p_* a dosing recommendation. Note also that an optimiser working inside model I uses that model's assumed offspring law as well as the mean data; the actual modelling failure is treating validated means as sufficient warrant for that additional assumption.

6 A survival floor for every admitted feedback course

The reversal compares two fixed schedules. It is natural to ask what any controller could achieve, including one that adapts to the observed population. The next theorem answers this in the direction that matters — it is a feasibility obstruction — and it holds without assuming independence between founder outcomes, which fails under feedback.

Theorem 6.1 (Scalar feedback envelope; C with E arithmetic)**.** *On* $[0,T]$ *let actions be bounded and predictable, let division have constant rate* $b > 0$*, and let* $d_i \leq d_{\max}$ *and* $\mu_i \geq \mu_{\min} > 0$*. Retain the normalised binary daughter law and the one-mutant-daughter convention. Let* M *have autonomous rates* $\beta > \delta > 0$ *with* $\rho = \delta/\beta$*, and assume* $d_{\max} \geq b\rho$*. Let the post-course continuation create no new* M*, clear every finite molecular population almost surely, and retain the autonomous* M *rates. Define* q *by*

$$q' = d_{\max}(1-q) + b\{(1-\mu_{\min})q^2 + \mu_{\min}\rho q - q\}, \qquad q(0) = 1. \tag{30}$$

Then from n molecular founders of arbitrary types and m initial mutants every such feedback course satisfies

$$\mathbb{P}(\textit{eventual total survival}) \geq 1 - q(T)^n \rho^m. \tag{31}$$

Proof. The scalar field of (30) is $-b\mu_{\min}(1-\rho) < 0$ at $q = 1$ and $(1-\rho)(d_{\max} - b\rho) \geq 0$ at $q = \rho$, so $[\rho, 1]$ is invariant and $q \in [\rho, 1]$ throughout. Evaluate the actual molecular field at the constant molecular vector with all coordinates equal to q and mutant coordinate ρ. All molecular jumps cancel because Q_e is conservative, and every normalised two-daughter law satisfies $D(q\mathbf{1}, q\mathbf{1}) = q^2$ — this is where the comparison becomes law-independent. The difference between the scalar field and the actual one is

$$(d_{\max} - d_i)(1-q) + b(\mu_i - \mu_{\min})q(q-\rho) \geq 0.$$

For a population z set $V(t,z) = q(T-t)^{\sum_{i\in S} z_i} \rho^{z_M}$. Direct evaluation of each jump gives, for the currently chosen predictable action u,

$$(\partial_t + \mathcal{G}_u)V = V \sum_i \frac{z_i}{q_i}\{\Phi_i^u(q) - q_i'\} \leq 0, \qquad q_M = \rho.$$

This is a conditional generator identity for a population test function and requires no independence between founder outcomes. Stop at bounded population and proposal count, apply the stopped Dynkin inequality, and remove the stopping by nonexplosion (Proposition 2.2) and bounded convergence, using $0 \leq V \leq 1$. At T, $V(T, Z_T) = \rho^{Z_M(T)}$ is the conditional probability of eventual extinction under the declared continuation, since that continuation clears the molecular population and creates no new M. Taking expectations gives (31). □

On the admitted action rectangle $.01 \leq e \leq .3$, $0 \leq c \leq .121$ with $\varepsilon = .1$, we have $b = .1$, $d_{\max} = .421$, $\mu_{\min} = .01$ and $\rho = .2$, and $d_{\max} \geq b\rho$ holds comfortably. Writing $a = 1 - q$ reduces (30) exactly to the Riccati equation

$$a' = .0008 - .3228\, a - .099\, a^2, \qquad a(0) = 0 \quad \text{(E)}, \tag{32}$$

and an exact scalar enclosure gives

$$.0024725782 < a(20) < .0024725784 \quad \text{(E)}. \tag{33}$$

A one-founder eventual-survival target below .00247 is therefore impossible in this actuator class, whatever the feedback policy. The risk actually achieved by complementary BA is about .03041843. The gap between the achieved risk and the universal necessary bound is a factor of about twelve: the floor is a genuine obstruction, but neither global optimality nor near-optimality of the exhibited schedule follows from it. The floor pertains to the stated endpoint and action class; an arbitrary new continuation is not automatically covered.

7 How much continuation perturbation the decision tolerates

The baseline clearing continuation runs indefinitely and is not drug-free: maintaining $e = .3$ above the baseline .01 has infinite integrated erasure increment if the continuation is charged to the exposure budget. The equal exposures of Section 4 compare the two finite courses. The two propositions below quantify two separate departures from the idealised continuation, both at the central course parameters. They are separate scenarios; combining them, altering the M dynamics, or adding population interactions requires a further argument.

Both bounds are one-sided, which is what makes them efficient. Writing $\widehat{q}_s$ for the perturbed eventual extinction probability, in each case $0 \leq q_s - \widehat{q}_s \leq \theta_s$, so only θ_{BA} can erode the J preference m_J and only θ_{AB} can erode the I preference m_I.

7.1 Residual acquisition after the course

Proposition 7.1 (Residual acquisition; C with E arithmetic)**.** *Keep $e = .3$, $c = 0$ after T, but replace zero acquisition by $\mu_i = \eta\kappa_i$. For either inheritance law and either course s,*

$$0 \le q_s(0) - q_s(\eta) \le \eta U_s, \qquad U_{AB} < .050002, \quad U_{BA} < .053172.$$

Both central order preferences persist for every $0 \le \eta \le 3 \times 10^{-5}$.

Proof. Let A_0 be the mutation-free molecular clearing mean matrix and $f = b\kappa$. Matrices act on column rewards, forward mean rows obeying $m' = mA$. The perturbed molecular mean matrix is $A_\eta = A_0 - \operatorname{diag}(b\eta\kappa)\Lambda \le A_0$ entrywise and is Metzler, so semigroup comparison together with (11) preserves almost sure molecular clearing and bounds the molecular occupation means. Define the nonnegative column

$$v = \int_0^\infty e^{tA_0} f \, dt = -A_0^{-1} f,$$

solved exactly in rationals (E). Integrating the acquisition intensity against the dominated molecular mean and conditioning on each seed's autonomous descendants, the expected number of additional surviving resistant families from molecular type i is at most $(1-\rho)\eta v_i$; a union bound converts that expectation into a probability bound. Coupling with zero acquisition shows extinction can only decrease, because baseline molecular descendants already clear and the additional surviving mutant families are the only new source of survival. Conditioning at the deadline and writing A_A, A_B for the finite-course molecular mean matrices gives

$$U_{AB} = (1-\rho)\big[e^{10A_A} e^{10A_B} v\big]_{AA}, \qquad U_{BA} = (1-\rho)\big[e^{10A_B} e^{10A_A} v\big]_{AA},$$

enclosed by validated linear flows as described in Appendix B.2. Since the corrections are one-sided, it suffices that $\eta U_{BA} < m_J$ and $\eta U_{AB} < m_I$. At $\eta = 3\times10^{-5}$ these read $1.59515\times10^{-6} < 1.70134\times10^{-6}$ and $1.50006 \times 10^{-6} < 1.77058 \times 10^{-6}$ (E). □

The binding constraint is the first, and the same bounds fail at $\eta = 3.2\times10^{-5}$, so 3×10^{-5} is close to what this argument can give. It allows post-course acquisition probabilities up to 1.2×10^{-5} per division, still far below the in-course value .04; the improvement is a sharper sufficient radius, not a demonstration of broad biological robustness. Under residual acquisition the molecular continuation extinction coordinates fall below one and become law dependent, so (10) is not reused as their exact value.

7.2 Finite clearing and its exposure cost

Proposition 7.2 (Finite clearing; C with E arithmetic)**.** *Apply the baseline mutation-free clearing source for only H time units after T, and then a common specified source that retains autonomous M, has no immigration, and differs from baseline only through the remaining molecular cells. If $\widetilde{q}_s$ is the resulting eventual extinction probability, then*

$$0 \le q_s - \widetilde{q}_s \le C_s e^{-.09H}, \qquad C_{AB} < 1.430424, \quad C_{BA} < 1.452994,$$

and $H = 152$ preserves both central preferences. The additional common erasure increment exposure is $.29H = 44.08$.

Proof. The contraction weight (11) bounds the probability that any molecular cell remains at $T+H$ by $C_s e^{-.09H}$, with $C_{AB} = [e^{10A_A} e^{10A_B} w]_{AA}/10$ and similarly for BA; the constants are enclosed by validated linear flows (Appendix B.2).

For the sign, condition on the configuration at $T+H$. Under the baseline continuation the conditional eventual extinction probability is exactly ρ^{Z_M}, by Proposition 2.3. Under the alternate source, total extinction implies extinction of the autonomous descendants of the M cells already present, an event of probability exactly ρ^{Z_M} and unaffected by any later molecular behaviour; hence the conditional probability is at most ρ^{Z_M}. Therefore $\widetilde{q}_s \le q_s$. On the event that no molecular cell remains at $T+H$ the two continuations couple exactly, since the alternate source differs from baseline only through remaining molecular cells and admits no immigration, so the difference is supported on the complementary event and is bounded by its probability.

Because the error is one-sided, it suffices that $C_{BA} e^{-.09H} < m_J$ and $C_{AB} e^{-.09H} < m_I$. Bounding $e^{.09H}$ from below by a truncation of its nonnegative power series, at $H = 152$ these read $1.66386 \times 10^{-6} < 1.70134 \times 10^{-6}$ and $1.63801 \times 10^{-6} < 1.77058 \times 10^{-6}$ (E). The same bounds fail at $H = 151$. □

This sign argument requires that later molecular behaviour cannot alter the autonomous descendants of the M cells already present. It must not be extended to a continuation violating that assumption. Under the cruder two-sided estimate $|\widetilde{q}_s - q_s| \le C_s e^{-.09H}$ with the errors summed, the same margins require $H = 162$; the one-sided argument saves ten time units.

The honest reading is twofold, and both halves should be stated together. Finite clearing *is* realisable: the counterexample does not secretly depend on an infinite intervention. But the sufficient duration is expensive. The added common erasure increment exposure 44.08 is exactly 15.2 times the finite-course increment 2.9, and the clearing phase lasts 7.6 times as long as the course itself. The initial sequencing still changes how many resistant families are seeded, and clearing does not erase that difference because M is unaffected by it; but the common burden substantially weakens the practical attractiveness of the construction. No optimality of $H = 152$ is asserted, and H is a sufficient duration derived from a bound, not a physical minimum. Replacing the global weight estimate by an enclosed expected molecular cell count at $T+H$ would tighten it; no such optimisation is needed for the methodological theorem.

8 Dimensionless parameters, data requirements and limits

Put $\tau = bt$ with reference division rate $b = .1$. Table 3 gives every source coefficient in these units. Each pulse then has dimensionless duration one and the course has $\tau = 2$.

The correct reading of these numbers is narrow. The initial two-pulse course is short relative to the reference division timescale: its integrated division hazard is two. Calendar duration, cell-cycle statistics, acquisition mechanism and intervention response all remain uncalibrated. It would be wrong to call the course "two cell generations": division is an exponential clock rather than a synchronised cycle, so $bT = 2$ is the expected number of divisions along an immortal tracked lineage, while death and conditioning on survival change that reading. In a pure-birth reference population the doubling time is $\log 2/b$, so $T = 20$ equals $2/\log 2 \approx 2.885$ doubling times — which is again not a population doubling count in the full model with death, state changes and selection. A physical division rate b_{phys} would map one pulse to $1/b_{\text{phys}}$ physical time units, but choosing a time unit cannot by itself establish that the regimen is acute rather than protracted. Finally, the objective extends well beyond the course: the finite-clearing construction of Proposition 7.2 adds another 152 time units before the alternate source begins, so 20 is not the horizon over which eventual rescue is observed.

Table 3: Complete dimensionless synthetic specification. Rates divide by b; probabilities and mark counts do not. Erasure e and additional death c are model rates, not drug concentrations.

Quantity	Dimensional value or formula	Dimensionless value
Activating write rate	$u(.01 + a/2)$	$u(.1 + 5a)$
Repressive write rate	$u(.01 + r/2)$	$u(.1 + 5r)$
Activating erase rate	$a(e + r/4)$	$a(e/b + 2.5r)$
Repressive erase rate	$r(e + a/4)$	$r(e/b + 2.5a)$
Molecular division	.1	1
Base death, AU, AA	.01	.1
Base death, other types	.3	3
Pulse A: (e, c)	$(.3, 0)$	$(3, 0)$
Pulse B: (e, c)	$(.01, .120553)$	$(.1, 1.20553)$
Pulse duration; deadline	$10; 20$	$1; 2$
Clearing: (e, c, μ)	$(.3, 0, 0)$	$(3, 0, 0)$
Mutant birth; death	$.1; .02$	$1; .2$
Course acquisition AU, AA	$\varepsilon\kappa = .1(.4)$	.04 per division
Course acquisition, other	$\varepsilon\kappa = .1(.1)$	.01 per division
Partition probability	1/2 per old mark	1/2
Feedback erasure range	$[.01, .3]$	$[.1, 3]$
Feedback additional death	$[0, .121]$	$[0, 1.21]$

Likewise, a per-division probability of reaching a coarse-grained durable state is not comparable with a per-base DNA mutation rate. A meaningful comparison would require a specified target size, a route to resistance, a heritability assumption and an acquisition mechanism. The values .01 and .04 are synthetic and would need independent justification in any application; calling them phenotypic switching does not supply that justification, particularly as M here is irreversible and autonomous.

The practical content is about what a model must preserve, not about dosing. Our worked example fixes AA founders and compares exactly two equal-duration orders under a declared continuation. To assess an instance of this phenomenon one would need the joint sister-state distribution conditional on maternal state, type-specific division and death hazards, switching rates under intervention, division-associated acquisition, the preparation law, and the survival and reversion behaviour of the durable state. Marginal daughter frequencies and total population means, however well validated, cannot identify the dependence that decides this comparison.

Remaining limits. The result is confined to the literal source, the stipulated acquisition convention, the AA preparation, the two pulse orders, the certified killing interval and the declared continuations. The effect is small. Feedback feasibility is bounded but not optimised. Finite clearing is expensive. Acquisition rates and calendar times are not calibrated. Competition, spatial structure, immune effects, pharmacokinetics, toxicity, density dependence, inherited mutant phenotypes, simultaneous continuation perturbations and global schedule optimisation are all absent. None of these is needed for the counterexample, and none of them repairs the failure it exhibits: exact agreement of the complete mean dynamics does not determine the preferred order for eventual extinction.

A Mechanism details

A.1 Collecting the expansion

Substituting $x = \mathbf{1} - \varepsilon h + \varepsilon^2 k + O(\varepsilon^3)$, $\mu_i = \varepsilon\kappa_i$ and $y = \rho$ into (7) and using $Q_e\mathbf{1} = 0$:

$$\begin{aligned}
(Q_e x)_i &= -\varepsilon(Q_e h)_i + \varepsilon^2 (Q_e k)_i, \\
d_i(1 - x_i) &= \varepsilon d_i h_i - \varepsilon^2 d_i k_i, \\
(1 - \mu_i) D_i(x, x) &= 1 - \varepsilon\kappa_i - 2\varepsilon(\Lambda h)_i + \varepsilon^2\{2(\Lambda k)_i + D_i(h, h) + 2\kappa_i(\Lambda h)_i\} + O(\varepsilon^3), \\
\mu_i \rho(\Lambda x)_i &= \varepsilon\kappa_i\rho - \varepsilon^2\kappa_i\rho(\Lambda h)_i + O(\varepsilon^3), \\
-x_i &= -1 + \varepsilon h_i - \varepsilon^2 k_i.
\end{aligned}$$

The left side is $x' = -\varepsilon h' + \varepsilon^2 k'$. Matching $O(\varepsilon)$ gives $-h' = -(Q_e h)_i + d_i h_i + b\{-2(\Lambda h)_i + \kappa_i\rho - \kappa_i + h_i\}$, that is (12). Matching $O(\varepsilon^2)$ gives $k' = (Q_e k)_i - d_i k_i + b\{2(\Lambda k)_i + D_i(h, h) + (2 - \rho)\kappa_i(\Lambda h)_i - k_i\}$, that is (13). The $O(\varepsilon^2)$ acquisition contribution $b(2 - \rho)\kappa_i(\Lambda h)_i$ involves the joint law only through Λ and therefore cancels from $\Delta = k_I - k_J$.

A.2 The quasi-monotonicity inequalities

The order (16) has least element RR, greatest element AA, six covering pairs and seven nonempty up-sets, namely

$$\begin{gathered}
\{AA\}, \quad \{AU, AA\}, \quad \{UU, AU, AA\}, \quad \{AR, AU, AA\}, \\
\{UU, AR, AU, AA\}, \quad \{UU, UR, AR, AU, AA\}, \quad \text{and all six types.}
\end{gathered}$$

For each of the six covering pairs $i \lessdot j$ and each up-set indicator χ with $\chi_i = \chi_j$, Lemma 3.2 requires $\sum_k (A_{jk} - A_{ik})\chi_k \geq 0$. This is 32 rational inequalities per phase. Exact evaluation confirms all 64; the minimum value attained is 0, at the pair $UR \lessdot UU$ against the ray $\chi = \mathbf{1}_{\{AA\}}$, in both phases. The script `verify_claims.py` in the source package re-runs the enumeration and the arithmetic.

A.3 Why the asymptotic remainder is not usable

Let P, V_2, V_3 bound the first three ε-derivatives of x in sup norm. With logarithmic Jacobian norm at most .09 and the elementary bounds $|f_\varepsilon| \leq .04$, $|f_{x\varepsilon}| \leq .12$, $|f_{xx}| \leq .2$ and $|f_{xx\varepsilon}| \leq .08$, the variational equations and positive scalar comparison give

$$P' = .09P + .04, \qquad V_2' = .09V_2 + .2P^2 + .24P, \qquad V_3' = .09V_3 + .6PV_2 + .36V_2 + .24P^2,$$

all starting from zero, whence $\|x_I - x_J - \varepsilon^2\Delta\| \leq V_3(T)\varepsilon^3/3$. At $T = 20$ the coarse algebraic bound $e^{1.8} < 9$ yields $P \leq 7.2$, $V_2 \leq 2177.28$ and $V_3 < 1.84 \times 10^6$, so the permitted error at $\varepsilon = .1$ exceeds 600. A valid asymptotic statement can therefore be entirely useless for a 10^{-6} decision. This motivated both the exact identity of Proposition 3.4 and the enclosed flows. No larger search for a sharper remainder was undertaken.

B Exact enclosure method and evidence boundary

B.1 Why the flow intervals are rigorous

Write the literal field as $d + Bz + C(z, z)$ with nonnegative quadratic coefficients. If $c_0 = z$, the normalised time derivatives obey

$$(k+1)c_{k+1} = \mathbf{1}_{k=0}d + Bc_k + \sum_{j=0}^{k} C(c_j, c_{k-j}), \tag{34}$$

by differentiating the polynomial ODE. Applying (34) with c_0 equal to the *whole* invariant cube and outward interval arithmetic bounds the normalised twelfth derivative everywhere in each phase. If the largest absolute endpoint is M, Taylor's theorem bounds the degree-11 local remainder by $M(\Delta t)^{12}$. A small final Taylor term is never used as an unproved error estimate.

The Jacobian of the full field is Metzler on the cube, with molecular row sums at most $b - d_i \le .09$ and mutant row sum $\beta - \delta = .08$. The maximal-coordinate differential inequality for the difference of two solutions therefore gives sup-norm Lipschitz factor $e^{.09t}$; the implementation uses the rational upper bound $e^{.09\Delta t} \le 1/(1 - .09\Delta t)$. This estimate applies to the segment between any two probability solutions, so it does not require a componentwise interval box to stay inside the cube.

All interval endpoints are integers divided by 2^{110}. Multiplication and positive integer division round outward, and rational source coefficients are enclosed before use. Each step evaluates a degree-11 Horner polynomial on intervals, projects the midpoint into the cube, and adds the distance from that centre to both endpoints, including the projection, to the local remainder. If the old error is E, the new error satisfies

$$E_{\text{next}} \le \frac{E}{1 - .09\Delta t} + E_{\text{rounding}} + M(\Delta t)^{12}.$$

Errors carry across phase boundaries without reset. With $\Delta t = 1/20$ each schedule takes 400 steps, and each central order-gap enclosure has width below 4×10^{-16}, more than nine orders of magnitude smaller than the margin. This is exact validated arithmetic under the foregoing conventional soundness proof; it is not a machine-verified integrator.

B.2 Scalar and linear consequences

For (33) the interval $[0, .8]$ is invariant for a in (32), whose field derivative $-.3228 - .198a$ is negative there, so the flow is 1-Lipschitz and the same normalised-derivative construction — one coordinate, step .1, degree 11, 200 steps — accumulates local and rounding error without growth. It encloses $a(20)$ in $(.002472578295567906, .002472578295567911)$. The closed-form Riccati solution was evaluated as a diagnostic and matches to floating precision; it is not the certificate.

For the linear rewards of Section 7, the exact source matrices satisfy $\|A\|_\infty \le 3$ with row sums at most .09. The global bounds $9\|v\|_\infty < 1000$ and $9\|w\|_\infty < 1000$ control the reference rewards, and at each step the centre is checked to have norm at most 950, so the local reference trajectory over $\Delta t = .05$ stays below 1000. A degree-15 Taylor evaluation then has local remainder at most $1000 \cdot 3^{16}(.05)^{16}/16!$. Outward arithmetic, explicit centre rounding and the same rational logarithmic norm growth enclose the course means, so the continuation constants U_s and C_s do not rely on an unenclosed floating-point matrix exponential.

Regret, founder and perturbation comparisons are rational operations on the preserved central certificate. The interval theorem of Lemma 4.2 adds only four endpoint flows; it is not a validated high-dimensional sensitivity search. Floating-point roots and curves are illustrations and are inputs to no certified sign conclusion.

C Evidence map and reproduction

Evidence	Claims supported
C: conventional	Branching construction and probability bridge; extinction endpoint; expansion and exact response identities; covariance sign via Nagumo and Harris; curvature bound and the secant argument; feedback envelope; occupation, coupling and clearing bounds.
E: exact arithmetic	Mean-operator identity; clearing weight; the 64 quasi-monotonicity inequalities; central and endpoint flow intervals; scalar floor; linear reward bounds U_s, C_s; regret enclosure, founder increment signs and the two continuation radii.
K: Lean kernel	Finite source normalisation and augmented linearisation; the offspring expansion and covariance-forcing algebra; continuation weight; scalar-envelope and barrier algebra; regret inequalities; decision-preservation and presentation margins. Not the complete stochastic theorem.
N: diagnostics	Table 1 probabilities; the smooth gap curves and their zero crossings; the leading-order transport curves of Figure 2(c). No certified conclusion rests on these.

The probability construction, ODE correspondence, continuation identification, mechanism, curvature argument and feedback theorem are conventional proofs. The numerical claims use exact outward arithmetic under Appendix B. Lean modules under `proofs/EvolutionaryRescue/` check the literal source linearisation, the continuation weight, the scalar-envelope algebra, the regret inequalities and the presentation margins; strict verification treats compiler warnings as errors. The thin `Main.lean` imports those modules. It does *not* formalise the stochastic process, the interval integrator or the complete probability theorem, and no end-to-end Lean formalisation is claimed. Exact rational arithmetic does not by itself make the integrator kernel-verified; that boundary is stated here rather than buried in a repository note.

The source package accompanying this manuscript contains `main.tex`, the figure sources and two scripts: `figures.py`, which rebuilds both figures from the preserved certificates, and `verify_claims.py`, which replays every printed numerical claim, asserting the exact ones and tagging the diagnostics. The preserved original certificate is `results/decision_certificate.json`; the additional endpoint flows, continuation constants, scalar floor and decision consequences are under `results/postproof/`. The campaign workspace `REPRODUCE.md` gives a bounded command that audits the original certificate, rebuilds the selected consequences and figures, strictly verifies the Lean root and compiles this manuscript; the original four-flow producer remains available for independent replay. Reported reproduction receipts record provenance, not mathematical proof.

References

[1] E. B. Gunnarsson, S. De, K. Leder and J. Foo. Understanding the role of phenotypic switching in cancer drug resistance. *Journal of Theoretical Biology* **490** (2020), 110162. https://arxiv.org/abs/2001.07690.

[2] E. B. Gunnarsson, B. V. Magnússon and J. Foo. Optimal dosing of anti-cancer treatment under drug-induced plasticity. arXiv:2412.16391v2 (2025). https://arxiv.org/abs/2412.16391.

[3] T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Mathematical Proceedings of the Cambridge Philosophical Society* **56** (1960), 13–20.

[4] Y. Iwasa, M. A. Nowak and F. Michor. Evolution of resistance during clonal expansion. *Genetics* **172** (2006), 2557–2566. https://doi.org/10.1534/genetics.105.049791.

[5] H. A. Orr and R. L. Unckless. The population genetics of evolutionary rescue. *PLOS Genetics* **10**(8) (2014), e1004551. https://doi.org/10.1371/journal.pgen.1004551.

[6] S. Patil, A. Ahmed, Y. Viossat and R. Noble. Preventing evolutionary rescue in cancer using two-strike therapy. *Genetics* **232**(2) (2026), iyaf255. https://doi.org/10.1093/genetics/iyaf255.