Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Repair-Dependent Competition Between Direct Driver Founding and Stochastic Tunneling Before Clonal Establishment

Benjamin Hardisty

hardisty@math.utah.edu

VA Office of Research and Development https://orcid.org/0009-0009-0841-0577

Research Article

Keywords: clonal establishment, mutation–repair balance, stochastic tunneling, multitype branching process, first passage, driver progression

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11643163/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Repair-Dependent Competition Between Direct Driver Founding and Stochastic Tunneling Before Clonal Establishment

Ben Hardisty[1*]

[1]School of Biological Sciences, University of Utah, Salt Lake City, Utah, USA.

Corresponding author(s). E-mail(s): hardisty@math.utah.edu;

Abstract

DNA repair may change not only whether damaged cells survive but also which mutational route produces successful progression. We compare direct two-driver founding in a single division with sequential stochastic tunneling through a transient one-driver lineage. Founders arise from an endogenous homeostatic mutation–repair equilibrium and initiate branching processes stopped by extinction, an irreversible commitment event, or a finite clone-size boundary; only commitment before the other outcomes counts as success. At each fixed size boundary, both routes contribute at order $\boldsymbol{V^2}$ in a rare independent-driver mutation expansion, with $\boldsymbol{V}$ held unchanged across repair comparisons. Exact finite-state coefficients establish a repair-induced reversal of the dominant route after averaging over the homeostatic founder mixture. An opportunity–turnover–establishment identity shows why improved cellular demography need not increase divisions eligible for a second hit. With lesion-dependent post-hit outcomes, a sharp criterion describes overshoot in the unbounded subcritical model, while certified finite-clone examples demonstrate re-entrant dominance. The resulting biological observable is a repair-dependent mixture of successful routes, distinct from clinical cancer incidence.

Keywords: clonal establishment, mutation–repair balance, stochastic tunneling, multitype branching process, first passage, driver progression

MSC Classification: 92D15 , 60J80 , 60J35

1 Introduction

Cancer-associated driver mutations can arise in histologically normal tissues, yet acquiring an alteration, maintaining its lineage, and generating a descendant that reaches a further progression outcome are distinct stochastic events (Martincorena et al. 2015, 2018; Vermeulen et al. 2013). This distinction matters when repair alters both the dynamics of an intermediate clone and the fate of its driver-bearing descendants. Repair can improve a clone's demographic prospects while changing the number of divisions it undergoes before a size cutoff; it may also change the chance that a two-driver descendant commits. Which route, then, contributes more successful progression events?

We compare two routes: direct founding, in which two independent rare driver alterations arise in the same resident division, and sequential stochastic tunneling, in which a transient one-driver lineage produces a two-driver descendant before its extinction or clone-size stopping boundary. For a fixed size threshold M, both route contributions scale as V^2 when the independent per-locus driver-mutation probability V tends to zero; V itself is not varied in comparisons across repair rates. Our central observable is the *ratio of successful-founder coefficients*, which measures how repair redistributes successful progression routes.

Classical multistage and two-hit models established the importance of successive alterations in carcinogenesis (Armitage and Doll 1954; Knudson 1971; Moolgavkar and Knudson 1981). Stochastic-tunneling models showed that intermediate mutants can generate subsequent alterations without first fixing or establishing (Komarova et al. 2003; Iwasa et al. 2004; Weissman et al. 2009). Branching and pathway models have further characterized cancer initiation, driver accumulation, and founder waiting times (Bozic et al. 2010; Zhang and Bozic 2024; Haeno et al. 2007). Repair has also been incorporated into mathematical models of tumor-initiation waiting times (Plotkin and Nowak 2002), while temporally clustered mutagenesis provides a distinct mechanism for changing multistep evolutionary paths (Graser et al. 2026). Here we isolate the effect of *repair* on the relative contributions of successful routes: an exact resident mutation–repair equilibrium determines founder composition, while post-hit commitment competes with extinction and a finite clone-size boundary.

Our answer has three parts. First, an exact three-phenotype calculation at $M = 3$ proves a unique repair-induced reversal, including after averaging over the endogenous homeostatic founder source. Second, an exact demographic identity separates divisions eligible to produce a second driver from premature establishment, while lesion-dependent post-hit commitment can oppose this opportunity effect. Third, certified finite-clone calculations show that the dominant pathway can switch more than once as repair varies. We distinguish genuine crossings of pathway parity from extrema and overshoot of limiting comparisons, neither of which alone implies a reversal. Define the leading-order sequential-route share among successful founders by $f_{\mathrm{seq}}(\nu) = D_M(\nu)/[1 + D_M(\nu)]$. If resident and clonal births have the same rare per-locus driver-mutation probability, this share is independent of that probability at leading order. In the exact population-averaged $M = 3$ example, f_{seq} rises from 0.417 at $\nu = 0$ to 0.637 as $\nu \to \infty$, crossing $1/2$ at the unique repair threshold. The unbounded example underlying Section 5 yields a nonmonotone route mixture

across graded repair levels, with separate finite-M certificates of route reversal. Route classification requires ancestry linking the first two-driver cell to the genotype of its immediate parent at that division; an extant one-driver relative is not an equivalent readout.

2 Model and leading-order pathway competition

2.1 Homeostatic founder source

Rescale time by the pristine division rate. Let $F_i^*(U,\nu)$, $0 \le i \le n_L$, denote the stationary frequencies of a lesion-only resident population of size N. Each lesion is repairable, at total rate $i\nu$ in class i. Reproduction and mortality have demographic factors r_L^i and g_L^i, respectively, where $0 < r_L < 1 < g_L$. The homeostatic multiplier is

$$R_*(U,\nu) = \varepsilon \frac{\sum_{i=0}^{n_L} g_L^i F_i^*(U,\nu)}{\sum_{i=0}^{n_L} r_L^i F_i^*(U,\nu)}, \qquad 0 < \varepsilon < 1. \tag{1}$$

For the small-lesion-mutation equilibrium branch developed in Hardisty (2026), $F_0^* = 1 + O(U)$, $F_i^* = O(U^i)$, and $R_* = \varepsilon + O(U)$. The daughter-lesion mutation kernel $L_{ij}(U)$ is stochastic. Under independent per-locus mutations it is

$$L_{ij}(U) = \binom{n_L - i}{j - i} U^{j-i}(1-U)^{n_L - j} \quad (j \ge i), \tag{2}$$

with $L_{ij} = 0$ otherwise. In particular, offspring reproduction events creating lesion type j, summed over resident parent states, occur at intensity

$$\mathcal{R}_j(U,\nu) = N R_*(U,\nu) \sum_{i=0}^{n_L} r_L^i F_i^*(U,\nu) L_{ij}(U). \tag{3}$$

The normalization $\sum_j \mathcal{R}_j = N R_* \sum_i r_L^i F_i^* = N\varepsilon \sum_i g_L^i F_i^*$ follows immediately. We take $n_D \ge 2$ initially unaltered, irreversible driver loci, mutating independently with probability V per locus at each resident reproductive event. The leading founder immigration rates are

$$\iota_{1,j}(V,\nu) = n_D \mathcal{R}_j(\nu) V + O(V^2), \quad \iota_{2,j}(V,\nu) = \binom{n_D}{2} \mathcal{R}_j(\nu) V^2 + O(V^3). \tag{4}$$

The resident source remains frozen, and founder arrivals are approximated by independent Poisson processes. Here and below U is held fixed unless otherwise stated.

2.2 Finite clone dynamics and stopping outcomes

A clone cell has type $a = (i,d)$, where $i \in \{0,\dots,n_L\}$ counts repairable lesions and $d \in \{1,2\}$ counts fixed drivers. With budding reproduction a type-a cell produces *one additional* daughter at rate $b_a(\nu) > 0$; the parent remains. Death occurs at rate

$\mu_a(\nu) > 0$. On reproduction, the daughter inherits a specified lesion kernel and, from a one-driver mother, acquires the second driver with probability

$$\sigma(V) = (n_D - 1)V + O(V^2). \tag{5}$$

The general finite-state model permits lesion acquisition during clonal reproduction, but the explicit switching and re-entrant examples below set de novo lesion input during clone divisions to zero: ordinary daughters inherit their mother's lesion state. This assumption is distinct from the resident lesion-input probability U that determines founder supply and is not a consequence of taking U small.

Repairable lesion state and irreversible driver count are distinct variables. The two-driver class is terminal with respect to driver count; any further driver alterations are outside this two-hit model. A cell of lesion class $i \geq 1$ repairs one lesion at rate $i\nu$, becoming $(i-1, d)$. A two-driver cell undergoes irreversible commitment at rate $\theta_{i,2}(\nu) \geq 0$, while $\theta_{i,1} = 0$. We assume at least one commitment channel is strictly positive when comparing pathway ratios.

Write $Z = (Z_a)$ for cell counts, $|Z| = \sum_a Z_a$, and fix $M \geq 2$. The transient state space is finite:

$$\mathcal{S}_M = \{z \in \mathbb{N}_0^{2(n_L+1)} : 1 \leq |z| \leq M-1\}. \tag{6}$$

A clone is stopped on first reaching 0 (extinction, $\mathcal{X}$), first reaching M (establishment, $\mathcal{E}$), or undergoing a commitment event ($\mathcal{T}$). A daughter birth from size $M-1$ is counted as establishment, even if the daughter carries the second driver; commitment must occur *before* that boundary. The size-M boundary delimits the independent-clone calculation.

Let $Q_{V,\nu}$ be the transient generator on $\mathcal{S}_M$, including all absorption hazards in its diagonal. Write

$$t_\nu(z) = \sum_{i=0}^{n_L} \theta_{i,2}(\nu) z_{i,2}, \tag{7}$$

and let $e(z)$ and $x(z)$ be the establishment and extinction exit rates. The row sum is $Q_{V,\nu}\mathbf{1} = -(t_\nu + e + x)$. The three-way first-passage resolvent and an occupation-time Feynman–Kac representation are established in Appendix A. Figure 1 summarizes the two routes and their competing stopping outcomes.

2.3 Rare second hits and competing successful-founder pathways

2.3.1 Green-operator tunneling coefficient

Set $V = 0$ and restrict to clones containing only one-driver cells. Their transient generator on $\mathcal{S}_M^{(1)} = \{x \in \mathbb{N}_0^{n_L+1} : 1 \leq |x| \leq M-1\}$ is $Q_\nu^{(1)}$, and their Green operator is

$$G_\nu = (-Q_\nu^{(1)})^{-1} = \int_0^\infty e^{sQ_\nu^{(1)}} \mathrm{d}s. \tag{8}$$

For a one-driver state x, denote by $q_\nu(x + \mathbf{e}_{j,2})$ the probability of commitment before extinction or establishment after appending a two-driver daughter of lesion type j,

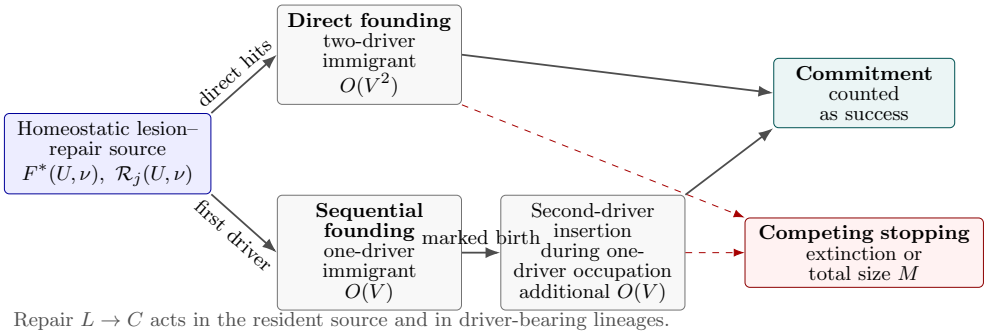


Fig. 1 The two successful-founder routes and their competing outcomes. Rare direct double-driver founders arise in a single resident division; sequential tunneling requires an additional driver-bearing birth within a one-driver lineage. Both routes contribute at order V^2 under independent per-locus mutations. A lineage is counted as successful only if irreversible two-driver commitment occurs before extinction or the total clone-size cutoff M.

with subsequent second-driver mutations suppressed. A birth that reaches size M contributes zero. Let $H_{ij}(\nu)$ be the conditional lesion type of the new double-driver daughter upon a second-driver mutation; its rows sum to one. Define

$$h_\nu(x) = \mathbf{1}_{\{|x|\le M-2\}} \sum_{i=0}^{n_L} b_{i,1}(\nu) x_i \sum_{j=0}^{n_L} H_{ij}(\nu) q_\nu(x + \mathbf{e}_{j,2}). \tag{9}$$

Here h_ν omits the multiplicity factor $n_D - 1$, which is already in $\sigma(V)$.

Theorem 2.1 (Exact coefficient of rare sequential tunneling) *For each finite M and fixed ν, the first commitment probability from a one-driver configuration x has the expansion*

$$p_{V,\nu}(x) = (n_D - 1)V\, A_M(x;\nu) + O(V^2), \qquad A_M(\cdot;\nu) = G_\nu h_\nu. \tag{10}$$

Equivalently,

$$A_M(x;\nu) = \mathbb{E}_x^{0,\nu}\left[\int_0^{\tau_M^0} h_\nu(X_s)\,\mathrm{d}s\right]. \tag{11}$$

The expansion is locally uniform for ν in compact intervals on which the finite-state rates and mutation kernels are sufficiently smooth. The one-driver and post-hit lineages may be supercritical.

Proof For fixed finite M, the full generator depends analytically on the conditional second-driver mutation probability σ near zero: $Q_{\sigma,\nu} = Q_{0,\nu} + \sigma B_\nu + O(\sigma^2)$. The commitment vector does not depend on σ. Differentiate $-Q_{\sigma,\nu} p_{\sigma,\nu} = t_\nu$ to obtain $-Q_{0,\nu}\partial_\sigma p|_0 = B_\nu p_{0,\nu}$. From a one-driver-only configuration, $p_{0,\nu} = 0$, and the first-order births to two-driver states yield exactly h_ν; all first-order changes among one-driver births multiply zero. Restricting the differentiated equation to one-driver configurations gives $-Q_\nu^{(1)}\partial_\sigma p|_0 = h_\nu$. Inverting, inserting $\sigma(V) = (n_D - 1)V + O(V^2)$, and using the standard occupation identity for G_ν proves the theorem. Smooth finite-dimensional inversion proves the remainder assertion. □

2.3.2 Two quadratic-order routes

For a founder of lesion state j, write $A_j(\nu) = A_M(\mathbf{e}_{j,1}; \nu)$ and $q_j(\nu) = p_{0,\nu}(\mathbf{e}_{j,2})$. When $\mathcal{R}_j$ is the source reproductive intensity in (3), define

$$\mathcal{A}_{\text{seq}}(\nu) = n_D(n_D - 1) \sum_{j=0}^{n_L} \mathcal{R}_j(\nu) A_j(\nu), \tag{12}$$

$$\mathcal{A}_{\text{dir}}(\nu) = \binom{n_D}{2} \sum_{j=0}^{n_L} \mathcal{R}_j(\nu) q_j(\nu). \tag{13}$$

These are coefficients of V^2, not already-mutated clone counts.

Theorem 2.2 (Quadratic-order pathway competition) *Under independent rare Poisson immigration from the frozen resident source,*

$$\mathcal{H}(V, \nu) = V^2[\mathcal{A}_{\text{seq}}(\nu) + \mathcal{A}_{\text{dir}}(\nu)] + O(V^3) \tag{14}$$

is the arrival intensity of founders marked by eventual commitment before their own extinction or establishment. *Whenever the direct coefficient is positive, sequential versus direct dominance at leading order is determined by*

$$D_M(\nu) = \frac{\mathcal{A}_{\text{seq}}(\nu)}{\mathcal{A}_{\text{dir}}(\nu)} = \frac{2\sum_j \mathcal{R}_j(\nu) A_j(\nu)}{\sum_j \mathcal{R}_j(\nu) q_j(\nu)}. \tag{15}$$

The expansion holds at fixed M, with ν and U fixed; no uniform-in-M assertion is made.

Proof An exactly one-driver newborn has intensity $n_D\mathcal{R}_j V + O(V^2)$; by Theorem 2.1 it is successful with probability $(n_D-1)VA_j+O(V^2)$. An exactly two-driver newborn has intensity $\binom{n_D}{2}\mathcal{R}_j V^2+O(V^3)$, followed by success probability q_j when further driver hits are suppressed. Poisson marking and addition give (14). The combinatorial identity $\binom{n_D}{2} = n_D(n_D - 1)/2$ gives (15). □

Remark 2.3 The ratio in (15) cancels the number of susceptible loci and the overall source intensity only when founder types and their fate probabilities enter with the corresponding common weights. If lesion-dependent effects distinguish first-hit and direct two-hit founder supply, the general ratio must instead use the actual source-specific vectors, not a universal factor of two. The model currently assumes both driver outcomes sample the same daughter-lesion kernel to leading order.

The intensity in (14) marks founders by eventual success, rather than by an assumed instantaneous commitment event; timing distributions are outside the present pathway-ratio analysis.

3 Repair reverses the dominant pathway

We now ask whether repair can change which successful-founder route is dominant. The principal result of this section is a unique *population-wide* reversal after exact

homeostatic averaging. A small conditional calculation supplies the essential input to its proof.

Consider three clone phenotypes: C (one driver, lesion-free), L (one driver, one repairable lesion), and D (two drivers, lesion-blind). The per-cell rates are

Type	Budding rate	Mortality rate	Commitment rate
C	4	1	0
L	1/10	1	0
D	1/2	1	50

Repair converts $L \to C$ at rate ν. Ordinary daughters inherit the maternal phenotype; a marked second-driver daughter is D regardless of the mother's lesion state. With establishment cutoff $M = 3$, the nine transient clone states have total size one or two. Only births from size-one states can produce a second driver that subsequently commits *before* the size cutoff; births at size two are already absorbed at size three.

For two independently mutable driver loci, let q_D be the commitment probability of a directly founded D cell and $A_{3,L}(\nu)$ the first-order marked-second-hit commitment coefficient from one L cell. The conditional sequential/direct ratio is $D_L(\nu) = 2A_{3,L}(\nu)/q_D$.

Lemma 3.1 (Conditional finite-clone calculation) *In this three-phenotype model,* $q_D = 10400/10607$ *is independent of repair. The exact ratio* $D_L(\nu)$ *is strictly increasing on* $[0, \infty)$*, with*

$$D_L(0) = \frac{198275}{1012024} < 1, \qquad \lim_{\nu\to\infty} D_L(\nu) = \frac{2575}{1469} > 1. \tag{16}$$

Its unique crossing of parity is $\nu_c \approx 1.10964042855$*. All comparisons are at leading order* V^2*, with* $V \downarrow 0$ *at fixed* $M = 3$*.*

Proof The nine-state killed generator and the marked-birth perturbation give $q_D = [(-Q_0)^{-1}t]_{(0,0,1)}$ and $A_{3,L} = [(-Q_0)^{-1}B(-Q_0)^{-1}t]_{(0,1,0)}$. The generator recipe and exact rational formula are recorded in Appendices B and B.1. The differentiated numerator has strictly positive coefficients, which proves monotonicity; the opposite signs relative to unity at the endpoints give the unique crossing. □

3.1 The reversal persists after exact homeostatic averaging

The conditional calculation does not yet determine which route is dominant in a population: repair also changes the types of founders produced by the resident tissue. We now determine those founder weights from the same two-class resident mutation–repair equilibrium.

Take one repairable lesion locus and the exact source parameters

$$U = \frac{1}{2}, \qquad r_L = \frac{1}{2}, \qquad g_L = 2, \qquad \varepsilon = \frac{1}{4}. \tag{17}$$

Write $x(\nu) = F_0^*(U, \nu)$ for the pristine resident frequency. The daughter-lesion kernel is $L_{00} = L_{01} = 1/2$, $L_{11} = 1$, with no further lesion state. Resident balance gives

$$R_* = \frac{2 - x}{2(1 + x)}, \qquad 0 = \frac{1}{2} R_* x - \frac{1}{4} x + \nu(1 - x), \tag{18}$$

so the physical equilibrium branch is

$$(4\nu + 2)x^2 - x - 4\nu = 0, \qquad \boxed{x(\nu) = \frac{1 + \sqrt{64\nu^2 + 32\nu + 1}}{8\nu + 4}}. \tag{19}$$

Thus $x(0) = 1/2$, $x(\nu)$ increases with repair, and $x(\nu) \uparrow 1$ as $\nu \to \infty$. The birth-weighted founder intensities are

$$\mathcal{R}_0 = \frac{N R_* x}{2}, \qquad \mathcal{R}_1 = \frac{N R_*}{2}, \qquad \mathcal{R}_0 : \mathcal{R}_1 = x : 1. \tag{20}$$

Even at arbitrarily high repair, lesion-bearing daughters continue to be supplied by new lesions arising during resident division.

A pristine founder does not change phenotype under repair. Hence its conditional pathway ratio is the constant

$$D_C = \frac{2A_{3,C}}{q_D} = \frac{2575}{1469} = \lim_{\nu \to \infty} D_L(\nu). \tag{21}$$

The last equality also follows from rapid repair of a damaged one-driver founder before its next demographic event, at fixed M.

Theorem 3.2 (Population-wide repair-induced pathway reversal) *Under the source and clone parameters above, with two independent driver loci and $M = 3$, the exact leading-order population-wide ratio is*

$$\boxed{D_3^{\mathrm{pop}}(\nu) = \frac{x(\nu) D_C + D_L(\nu)}{1 + x(\nu)}.} \tag{22}$$

This ratio is strictly increasing for all $\nu \geq 0$ and satisfies

$$D_3^{\mathrm{pop}}(0) = \frac{40878125}{57179356} < 1, \qquad \lim_{\nu \to \infty} D_3^{\mathrm{pop}}(\nu) = \frac{2575}{1469} > 1. \tag{23}$$

Consequently a unique population-wide crossing occurs at

$$\boxed{\nu_c^{\mathrm{pop}} \approx 0.192345334861229.} \tag{24}$$

At leading order V^2 as $V \downarrow 0$ with $M = 3$ fixed, direct founding dominates below this repair rate and sequential tunneling dominates above it, after endogenous homeostatic averaging.

Proof The source weights in (20) and the general pathway formula (15) yield (22). Differentiating the equilibrium quadratic gives

$$x'(\nu) = \frac{4(1 - x(\nu)^2)}{2(4\nu + 2)x(\nu) - 1} > 0$$

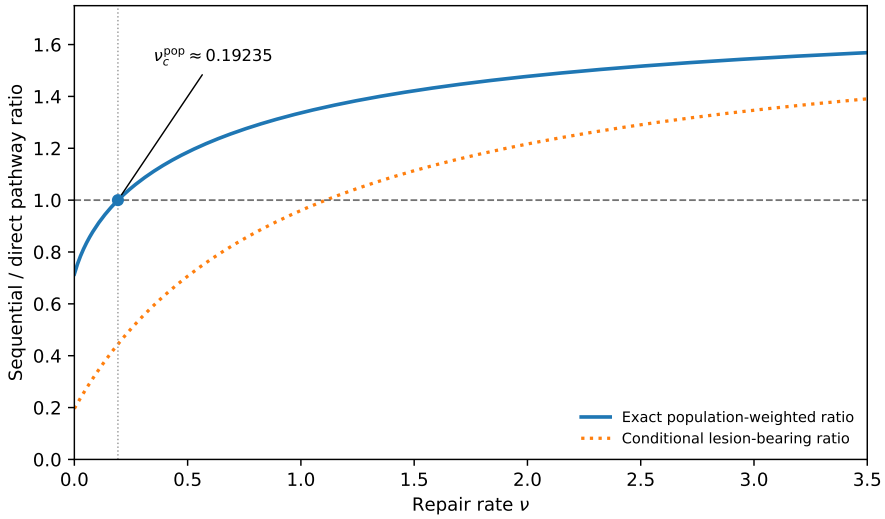


Fig. 2 Exact population-weighted (solid) and conditional lesion-bearing (dotted) sequential/direct route ratios. The population-wide crossover at $\nu_c^{\mathrm{pop}} \approx 0.1923453$ precedes the conditional crossover at $\nu_c \approx 1.1096404$. Both curves represent leading-order V^2 successful-founder comparisons at fixed $M = 3$.

for finite $\nu \geq 0$. By Lemma 3.1, $D_L'(\nu) > 0$ and $D_L(\nu) < D_C$ for finite ν, so

$$\frac{\mathrm{d}}{\mathrm{d}\nu} D_3^{\mathrm{pop}}(\nu) = \frac{x'(\nu)[D_C - D_L(\nu)] + [1 + x(\nu)]D_L'(\nu)}{[1 + x(\nu)]^2} > 0. \tag{25}$$

The endpoint values follow from $x(0) = 1/2$, $x(\infty) = 1$ and (16). By continuity and strict monotonicity there is exactly one crossing; solving (22)$= 1$ yields (24). □

At the population-wide crossing, $x \approx 0.7377$, while the conditional damaged-founder ratio is only $D_L \approx 0.4446 < 1$; the pristine-founder ratio remains $D_C \approx 1.7529 > 1$. Thus population-wide parity does *not* require parity for each founder phenotype. It arises from their distinct route contributions and the repair-dependent weights of the endogenous founder supply. Figure 2 contrasts the two repair thresholds.

Remark 3.3 (Parameter dependence and local robustness) The exact population-wide reversal is established for the stated rate family. Here $U = 1/2$ lies outside the rare-lesion regime. Retaining these vital rates but varying the resident input over $0 < U < 3/4$, define the damaged-founder fraction $\phi(U, \nu) = \mathcal{R}_1(U, \nu)/[\mathcal{R}_0(U, \nu) + \mathcal{R}_1(U, \nu)]$. At zero repair, $\phi(U, 0) = 4U/3$. Because $D_C > D_L(0)$, direct dominance at zero repair requires $U > \frac{3}{4}(D_C - 1)/(D_C - D_L(0)) \approx 0.36267$. Thus these particular rate choices do not give a zero-repair direct advantage at arbitrarily small lesion input. Strict inequalities on opposite sides of the crossover persist under sufficiently small perturbations of the model rates, ensuring local robustness of the *existence* of a crossing, but not necessarily uniqueness or global monotonicity away from this exact parameter family.

4 Mechanisms: mutation opportunity and post-hit rescue

We now explain the exact population-wide reversal of Section 3 and identify why the same repair intervention need not favor the sequential route in every regime. The first mechanism is the number of divisions available to produce a second driver; the second is the subsequent commitment probability of the daughter carrying that driver.

Let τ_M be the time when the *one-driver-only* clone first becomes extinct or reaches size M. Let N_B count all births, N_D all deaths, and N_B^{elig} births occurring when the population size was at most $M-2$. Define $\mathcal{B}_M = \mathbb{E}N_B^{\mathrm{elig}}$, $\mathcal{N}_M = \mathbb{E}N_D$, and $P_M = \mathbb{P}(|Z_{\tau_M}| = M)$.

Theorem 4.1 (Opportunity–turnover–establishment balance) *For a multitype branching system* started from one cell*, with $M \geq 2$, budding births, individual deaths, repair transitions preserving clone size, and stopping at $\{0, M\}$,*

$$\boxed{\mathcal{B}_M = \mathcal{N}_M + (M-1)P_M - 1.} \tag{26}$$

Consequently, on every smooth repair interval,

$$\boxed{\mathcal{B}'_M = \mathcal{N}'_M + (M-1)P'_M.} \tag{27}$$

Proof Pathwise, $|Z_{\tau_M}| = 1 + N_B - N_D$, and its expected value is MP_M. The unique last birth of any established path crosses from $M-1$ to M and is ineligible, so $N_B^{\mathrm{elig}} = N_B - \mathbf{1}_{\{\text{establishment}\}}$. Take expectations and rearrange. Differentiation gives (27). Finite-state transience ensures finite expected birth and death counts and local smoothness. □

4.1 From opportunities to successful progression

For a lesion class j, the sequential coefficient is the expected sum of post-hit success probabilities over eligible potential second-hit births:

$$A_{M,j} = \mathbb{E}_j \sum_{e \in \text{eligible births}} q_\nu(Y_e). \tag{28}$$

Here the marked daughter configuration Y_e is drawn from the stipulated mutation-daughter kernel, and the sum is taken under the unmutated one-driver genealogy. Define $\bar{q}^{\mathrm{hit}}_{M,j} = A_{M,j}/\mathcal{B}_{M,j}$ whenever the denominator is nonzero. This separates mutation opportunity from successful post-hit commitment without pretending that all mutation events have the same probability of success.

Normalize the exact resident founder source as $\pi_j = \mathcal{R}_j / \sum_k \mathcal{R}_k$, and set

$$\overline{\mathcal{B}}_M = \sum_j \pi_j \mathcal{B}_{M,j}, \quad \overline{\mathcal{N}}_M = \sum_j \pi_j \mathcal{N}_{M,j}, \quad \overline{P}_M = \sum_j \pi_j P_{M,j},$$

$$\overline{q}_{\mathrm{seq}} = \frac{\sum_j \pi_j A_{M,j}}{\overline{\mathcal{B}}_M}, \qquad \overline{q}_{\mathrm{dir}} = \sum_j \pi_j q_{M,j}.$$

For positive pathway coefficients, define the logarithmic repair response $\Psi_M(\nu) := \partial_\nu \log D_M(\nu)$. Under the common independent-driver-locus mutation convention,

$$D_M = \frac{2\overline{\mathcal{B}}_M\, \overline{q}_{\mathrm{seq}}}{\overline{q}_{\mathrm{dir}}}, \qquad \boxed{\Psi_M = \frac{\overline{\mathcal{N}}'_M + (M-1)\overline{P}'_M}{\overline{\mathcal{B}}_M} + \partial_\nu \log \overline{q}_{\mathrm{seq}} - \partial_\nu \log \overline{q}_{\mathrm{dir}}.} \tag{29}$$

All derivatives include changes in the *endogenous founder distribution*: when differentiating the weighted balance identity, the source-weight terms cancel in the constant -1 because $\sum_j \pi'_j = 0$. Equation (29) is an exact decomposition of the repair response into mutation-producing opportunity and post-hit success, with all founder weights evaluated from the homeostatic source.

Which mechanism drives the exact reversal? The finite-generator calculation of Section 3 lets us examine the three factors of (29) without introducing any additional model. Table 1 compares the endogenous source-weighted opportunities, average success per eligible insertion, and pathway ratio at zero repair and at the exact population crossover. The direct two-driver success probability remains $q_D = 10400/10607$ at both rates, because the two-driver phenotype in this particular example is lesion-blind.

Table 1 Opportunity and post-hit success in the exact population-weighted $M = 3$ model. Finite-generator values are rounded. Here π_C is the pristine founder fraction (distinct from resident frequency x), $\nu_c^{\mathrm{pop}} = 0.192345334861\ldots$, and $\overline{q}_{\mathrm{seq}} = \sum_j \pi_j A_{3,j}/\overline{\mathcal{B}}_3$.

Repair rate	π_C	$\overline{\mathcal{B}}_3$	$\overline{q}_{\mathrm{seq}}$	D_3^{pop}
$\nu = 0$	0.333333	0.383526	0.913834	0.714911
$\nu = \nu_c^{\mathrm{pop}}$	0.424525	0.539814	0.908170	1.000000

In this specific switching example, eligible second-hit-producing births increase enough to reverse the dominant route, even though the mean success probability per eligible insertion decreases slightly. The change in $\overline{\mathcal{B}}_3$ includes both the changing founder mixture and the changing one-driver clone histories; it is not a claim that mutation opportunities always increase with repair. The balance identity does not, by itself, fix the direction of the opportunity response when the establishment probability rises; the observed positive response in this particular example must be calculated rather than assumed from cellular fitness.

4.2 Lesion-dependent post-hit commitment

The exact $M = 3$ switch isolates repair acting on the one-driver lineage by making the two-driver phenotype lesion-independent. We now change the phenotype assumptions,

rather than the underlying first-passage framework, to ask whether repair of a two-driver descendant can oppose the pre-hit opportunity effect. Explicit comparisons are possible in a different tractable regime: both one-driver phenotypes are strictly sub-critical, while two-driver birth, death, and commitment rates depend on lesion state. The no-new-lesion restriction of the explicit clone models is retained here, isolating post-hit rescue from continued lesion production.

Let one-driver types C, L have birth–death pairs (b_C, d_C) and (b_L, d_L), with $0 < b_i < d_i$, and set

$$a_C = d_C - b_C, \qquad a_L = d_L - b_L, \qquad f_C = \frac{b_C}{a_C}, \qquad f_L^0 = \frac{b_L}{a_L}.$$

The two-driver phenotypes C_2 and L_2 have birth, death, and commitment rates $(B_C, \delta_C, \theta_C)$ and $(B_L, \delta_L, \theta_L)$, with strictly positive rates. Each damaged cell, at either driver stage, repairs to its pristine type at rate ν. Birth is by budding, producing one same-type daughter. Commitment occurs independently at rate θ_i per two-driver cell. The rare-driver mutation coefficient is first extracted at each fixed finite size cutoff; the unbounded expressions below are obtained by taking the subsequent $M \to \infty$ limit.

Write q_C and $q_L(\nu)$ for the unbounded probabilities of eventual commitment from a single C_2 and L_2 cell. A first-event decomposition gives

$$B_C q_C^2 + (\delta_C + \theta_C - B_C) q_C - \theta_C = 0, \qquad 0 < q_C < 1, \tag{30}$$

so that

$$q_C = \frac{2\theta_C}{\delta_C + \theta_C - B_C + \sqrt{(\delta_C + \theta_C - B_C)^2 + 4B_C\theta_C}}. \tag{31}$$

Set

$$F_L(q) = B_L q^2 + (\delta_L + \theta_L - B_L) q - \theta_L. \tag{32}$$

Then $q_L(\nu)$ is the unique root in $(0, 1)$ of

$$F_L(q_L) = \nu(q_C - q_L). \tag{33}$$

If $q_C > q_L(0)$, it follows that $q_L(0) < q_L(\nu) < q_C$, $\nu \geq 0$, and

$$q_L'(\nu) = \frac{q_C - q_L(\nu)}{F_L'(q_L(\nu)) + \nu} > 0. \tag{34}$$

The positivity of the denominator follows by differentiation of the strictly convex first-event quadratic at its positive root.

The unbounded first-order successful second-driver mutation-insertion potentials obey

$$A_C = f_C q_C, \qquad A_L(\nu) = \frac{b_L q_L(\nu) + \nu f_C q_C}{a_L + \nu}. \tag{35}$$

The one-driver family of the newly created two-driver cell evolves independently of all remaining one-driver cells once the size boundary is removed; this is why a successful insertion into a type-i cell is weighted by q_i.

Let $\phi(\nu) \in (0,1)$ denote the lesion-bearing fraction of the homeostatically produced founder population (after common reproduction and population-scale factors cancel). The leading-order ratio of successful sequential and direct two-driver founding is

$$D_\infty(\nu) = 2\frac{(1-\phi)A_C + \phi A_L}{(1-\phi)q_C + \phi q_L}. \tag{36}$$

These identities hold for any continuous founder fraction $\phi : [0,\infty) \to (0,1)$. The high-repair limits below require only that ϕ remain bounded, as is automatic for a population fraction; in particular, they apply to the exact one-locus homeostatic source.

Theorem 4.2 (Sharp high-repair overshoot criterion) *Assume $f_C > f_L^0$ and $q_C > q_L(0)$. Define*

$$T = a_L f_C - b_L > 0, \qquad W(q) = f_C F_L(q) - Tq. \tag{37}$$

Then, for every finite $\nu \geq 0$,

$$D_\infty(\nu) - 2f_C = \frac{2\phi(\nu)W(q_L(\nu))}{(a_L+\nu)[(1-\phi(\nu))q_C + \phi(\nu)q_L(\nu)]}. \tag{38}$$

The high-repair limit is $\lim_{\nu\to\infty} D_\infty = 2f_C$. Furthermore:

1. *If $f_C F_L(q_C) \leq Tq_C$, then $D_\infty(\nu) < 2f_C$ for every finite $\nu \geq 0$.*
2. *If $f_C F_L(q_C) > Tq_C$, there is exactly one $\nu_* > 0$ at which $D_\infty(\nu_*) = 2f_C$. The ratio lies below $2f_C$ for $\nu < \nu_*$ and above it for $\nu > \nu_*$, while converging back to $2f_C$ at infinite repair. In particular, D_∞ is nonmonotone and has at least one finite interior maximum.*

The overshoot criterion is a rate-only inequality and does not depend on the homeostatic mixing function ϕ beyond its positivity.

Proof Equation (35) implies

$$A_L - f_C q_L = \frac{\nu f_C(q_C - q_L) - Tq_L}{a_L + \nu}.$$

Insert this into (36) and invoke (33), yielding (38). At zero repair, $W(q_L(0)) = -Tq_L(0) < 0$. As $\nu \to \infty$, equation (33) forces $q_L \to q_C$, while equation (35) gives $A_L \to A_C$. The polynomial W has positive quadratic coefficient $f_C B_L$ and negative constant coefficient $-f_C\theta_L$. Thus it has exactly one positive root, and $W(q)$ is negative below that root and positive above. The strictly increasing function $q_L(\nu)$ traverses the interval $[q_L(0), q_C)$, proving both alternatives and uniqueness of crossing of the horizontal level $2f_C$. In the second case, continuity and the strict overshoot above the eventual limit imply attainment of a finite interior maximum. The theorem makes no claim of uniqueness of this maximum. □

For biological interpretation, write

$$\frac{F_L(q_C)}{q_C} > a_L - \frac{b_L}{f_C}.$$

The left side is the positive residual of the damaged two-driver first-event equation evaluated at the pristine two-driver commitment probability; the right is the corresponding one-driver opportunity penalty, in units of inverse time. The inequality compares these two mechanisms; dominance reversal additionally requires a crossing of $D_\infty = 1$.

Finite-size implications. The overshoot theorem concerns the unbounded subcritical model. In the fixed-rate, no-new-lesion subclass with positive commitment hazards and strictly subcritical one-driver phenotypes, an estimate uniform over all repair rates gives $\sup_{\nu \geq 0} |D_M(\nu) - D_\infty(\nu)| \leq C(M+1)e^{-cM}$ for sufficiently large M; the complete assumptions and proof are in Appendix C (Theorem C.1). Thus strict dominance signs at finitely many repair rates persist at all sufficiently large finite thresholds. This does *not* transfer derivative signs, certify a particular moderate M, or interchange the limits $V \downarrow 0$ and $M \to \infty$. Section 5 uses independent finite-state certificates at specific thresholds.

5 Re-entrant dominance at finite clone size

The overshoot criterion of Section 4 concerns the high-repair limiting level, not parity of the two pathways. We now establish that repair can reverse the dominant successful-founder route *twice* within the finite-M branching model. This is a statement about crossings of $D_M = 1$, not simply extrema of a response curve.

5.1 Certified finite-threshold reversals

Consider the exact homeostatic source $(U, r_L, g_L, \varepsilon) = (1/2, 1/2, 2, 1/4)$ and $(2 + 4\nu)x^2 - x - 4\nu = 0$, with lesion-bearing founder fraction $\phi = 1/(1+x)$. The one-driver birth and mortality rates are $(b_C, d_C) = (306/755, 6/5)$ and $(b_L, d_L) = (1/5, 3/2)$; the two-driver birth, mortality, and commitment rates are $(B_C, \delta_C, \theta_C) = (12/5, 8/5, 1)$ and $(B_L, \delta_L, \theta_L) = (2, 2, 10^{-6})$. Repair therefore leads to a phenotype with higher reproduction and lower mortality at both driver stages and a higher two-driver commitment hazard. These rates satisfy case 2 of Theorem 4.2: $f_C = 0.51 > f_L^0 = 2/13$, $q_C \approx 0.605174 > q_L(0)$, and $f_C F_L(q_C) \approx 0.373560 > T q_C \approx 0.280196$. Thus the unbounded ratio is nonmonotone; the finite-M parity crossings below require the separate certificates.

At the finite establishment threshold $M = 10$, the independently constructed killed generator gives

$$\begin{aligned} D_{10}(0) &\approx 1.007919 > 1, \\ D_{10}(1/10) &\approx 0.962283 < 1, \\ D_{10}(10) &\approx 1.016642 > 1. \end{aligned} \tag{39}$$

These three strict signs are enclosed by exact rational-residual bounds (Appendix D). Continuity implies at least two distinct repair intensities at which $D_{10} = 1$, with a sequential–direct–sequential dominance sequence as repair increases. No claim of exactly two crossings, or of a complete stationary-point count, is made. Under the *same* biological rates, the certified $M = 9$ values at $\nu = 0, 6, 1000$ are approximately 0.998472, 1.009205, and 0.998640, respectively. They instead witness a

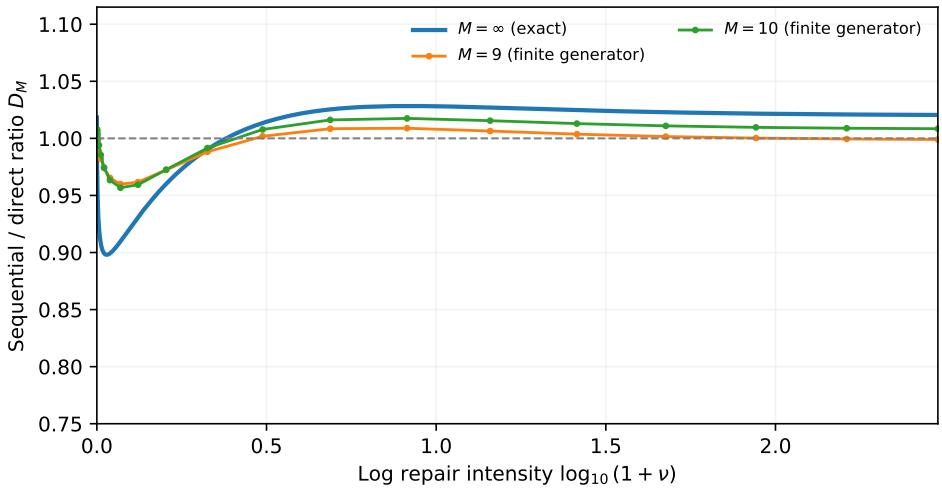


Fig. 3 Repair-dependent sequential-to-direct pathway ratio for the rational-rate lesion-dependent two-driver example. The solid curve is the exact unbounded expression; the other curves are independent finite-state evaluations at $M = 9$ and $M = 10$. The horizontal dashed line marks parity $D_M = 1$. The curves are numerical evaluations; strict parity signs at the stated witness rates have separate rational-residual certificates. The $M = 9$ return to direct dominance is certified at $\nu = 1000$, beyond the plotted repair range. The log horizontal axis is $\log_{10}(1 + \nu)$ so that $\nu = 0$ is visible.

direct–sequential–direct pattern, with rigorous intervals supplied in Appendix D. Thus the choice of finite clone-size stopping boundary can change the qualitative ordering.

Figure 3 visualizes the pathway ratios; the plotted curves interpolate between the separately certified witness values.

5.2 Magnitude and sensitivity

The main $M = 10$ example has a sequential advantage of about 0.79% at $\nu = 0$, a ratio about 3.77% below parity at $\nu = 0.1$, and a sequential advantage of about 1.66% at $\nu = 10$. These percentages quantify departures of the successful-founder pathway ratio from parity.

To assess dependence on the constructed rates, Table 2 varies the effective resident lesion input U and damaged two-driver commitment rate θ_L while keeping the other rates and $M = 10$ fixed. The entries are independently reproducible from the finite generator, and the strict inequalities are also certified by rational residuals (Appendix D.1).

Each row has a ratio strictly above parity at the first and last witness rates and strictly below parity at the intermediate rate. By continuity, at least two pathway reversals therefore persist for each displayed parameter combination; continuity of the finite-state resolvents also preserves these strict signs in a sufficiently small parameter neighborhood of each combination. The sizes of these parameter neighborhoods have not been quantified. The last row is particularly close to parity at its intermediate witness: $D_{10}(0.2) \approx 0.9994643$, just 5.36×10^{-4} below unity.

Table 2 Finite-$M = 10$ pathway ratios for three illustrative combinations of lesion input U and damaged-cell commitment hazard θ_L. Every listed parity sign has an exact rational-residual certificate in Appendix D.1. Entries are rounded.

U	θ_L	ν_{mid}	$D_{10}(0)$	$D_{10}(\nu_{\mathrm{mid}})$	$D_{10}(10)$
0.50	10^{-6}	0.1	1.00792	0.96228	1.01664
0.20	0.001	0.1	1.00696	0.99559	1.01130
0.10	0.010	0.2	1.00403	0.99946	1.00954

The two parameters vary together, so Table 2 does not isolate their separate effects. Details of the sensitivity certificates, including exact interval enclosures for the two additional parameter combinations, are provided in Appendix D.1.

6 Discussion

6.1 Repair redistributes successful progression pathways

Our central result concerns *which route succeeds*, not simply whether repair improves an individual cell's fitness. At fixed rare driver-mutation probability, direct two-driver founding and sequential tunneling occur at the same quadratic order, but expose different parts of their histories to repair. The exact population-weighted three-cell theorem shows a unique change in the dominant route. With lesion-dependent post-hit outcomes, we also obtain a sharp overshoot criterion and finite-clone certificates of re-entrant route dominance. Neither conclusion follows from favorable cellular birth and death rates alone.

6.2 Why the route changes

There are three separable influences: the homeostatic source changes the lesion composition of founders; repair changes the number of intermediate-clone divisions eligible to produce a second driver; and repair changes commitment probabilities after the second hit. In the exact $M = 3$ switching example, repair increases the source-averaged number of eligible births while the average post-hit success per birth falls slightly (Section 4). In other rate regimes, the opportunity–turnover–establishment identity shows that an increase in the chance of hitting the clone-size cutoff does not, by itself, determine the change in eligible divisions. The ratio of successful route contributions consequently cannot be inferred from fitness ordering, pre-hit occupation alone, or post-hit success alone.

6.3 Biological observable and scope

The central prediction concerns the *route mixture among successful founder histories* as repair function varies. In a lineage-resolved system where direct and sequential

origins can be distinguished, the empirical comparison is the relative representation of those routes among lineages reaching a defined progression outcome. The exact population-weighted calculation and certified finite-clone examples establish, respectively, a unique reversal and re-entrant changes under lesion-dependent post-hit outcomes. Together, they motivate lineage-resolved comparisons across graded repair conditions.

The mutation-rate comparison can be made more explicit. Let V_r be the per-locus probability of a driver mutation during a resident division and V_c the corresponding probability during a one-driver clone division. Direct founding contributes at order V_r^2, whereas sequential tunneling contributes at order $V_r V_c$. Consequently, for the same founder source and demographic rates,

$$D_M(\nu; V_r, V_c) = \frac{V_c}{V_r} D_M(\nu) + o(1) \tag{40}$$

as $V_r, V_c \downarrow 0$ with a fixed positive ratio; $D_M(\nu)$ on the right is the equal-rate coefficient ratio in (15). A common multiplicative change in the two mutation probabilities does not by itself change f_{seq} at leading order, but a change in V_c/V_r resets the route comparison and must be measured or controlled.

These proportions also give an illustrative measure of sampling difficulty. For two-sided, equal-size, unadjusted two-proportion tests ($\alpha = 0.05$, power 0.8), assuming independently sampled commitments, perfect route attribution, and no missing observations, comparing $\nu = 0$ with the high-repair limit in Section 3 requires approximately 80 committed lineages *per condition*. In the Section 5 example, separate pairwise comparisons of $\nu = 0$ versus 0.1 and 0.1 versus 10 require roughly 3.5–4.1 $\times 10^3$ per condition in the unbounded limit and 2.1–2.9 $\times 10^4$ at $M = 10$. These are idealized binomial calculations for each contrast, not a jointly powered experimental design; multiplicity adjustments, imperfect classification, and shared ancestry increase the required sampling effort.

The model takes a resident mutation–repair equilibrium as the source of rare founders and follows independent branching lineages to extinction, commitment, or the size-M boundary. Its reversible lesion phenotype is coarse grained, and its parameter examples are mechanistic rather than tissue-fitted. The exact switching and re-entrant examples suppress new lesion acquisition during clonal divisions, although lesion-bearing founders continue to arise from the resident source. This restriction is not implied by small U, and the examples respond differently to its relaxation. In a separate finite-generator sensitivity calculation, let u be the probability that a daughter of a pristine clone cell acquires a new lesion. The $M = 3$ population reversal persists for $0 \le u \le U = 1/2$; the numerically located crossing moves from $\nu \approx 0.19235$ to $\nu \approx 0.18986$. For the $M = 10$ re-entrant parameters, the low-repair sequential advantage disappears at $u \approx 0.00524$ if new lesions may also accompany a marked second-driver birth, or at $u \approx 0.01575$ if the marked daughter retains the maternal lesion state. These exploratory sensitivity thresholds depend on the marked-daughter convention; the certified finite-M re-entrant witnesses remain statements about the $u = 0$ model. *Because irreversible commitment is a specified first-passage outcome rather than a calibrated cancer diagnosis, the results do not by themselves predict*

clinical incidence, age-incidence slopes, or patient outcomes. Connecting route composition to those observables requires time-resolved founding and a biological mapping from commitment to progression and detection. Independent driver-locus mutation is assumed; correlated mutagenesis could alter the relative scaling of the routes. The rare-driver coefficient is extracted at fixed M, without asserting a uniform joint limit as $V \downarrow 0$ and $M \to \infty$.

Appendix A Exact first-passage laws and the Feynman–Kac bridge

Theorem A.1 (Three-way finite-clone resolvent) *For every finite M, all states of $Q_{V,\nu}$ are transient and $-Q_{V,\nu}$ is invertible. The probabilities of commitment, establishment and extinction as the first boundary event are, respectively,*

$$p_{V,\nu} = (-Q_{V,\nu})^{-1} t_\nu, \qquad v_{V,\nu} = (-Q_{V,\nu})^{-1} e, \qquad w_{V,\nu} = (-Q_{V,\nu})^{-1} x, \tag{A1}$$

with $p_{V,\nu} + v_{V,\nu} + w_{V,\nu} = \mathbf{1}$. No subcriticality assumption on a driver type is needed.

Proof Every state has a strictly positive-probability finite sequence of successive cell deaths leading to extinction. Since $\mathcal{S}_M$ is finite, it has no closed communicating class, hence all transient states are absorbed almost surely and $-Q_{V,\nu}$ is a nonsingular M-matrix. First-step analysis gives $-Q_{V,\nu} p = t_\nu$, and analogously for the other exit vectors. Finally, $Q_{V,\nu}\mathbf{1} = -(t_\nu + e + x)$ implies normalization after inversion. □

Theorem A.2 (Occupation-time Feynman–Kac representation) *Let Z^0 be the same finite clone with all commitment clocks suppressed, and let τ_M^0 be its first hitting time of 0 or M. Then, for every $z \in \mathcal{S}_M$,*

$$p_{V,\nu}(z) = 1 - \mathbb{E}_z^{V,\nu} \exp\left(- \int_0^{\tau_M^0} t_\nu(Z_s^0)\, \mathrm{d}s \right). \tag{A2}$$

If $t_\nu(z) = \xi\kappa(z)$ with $\kappa \geq 0$ and the underlying unmarked dynamics held fixed, the no-commitment probability is completely monotone in ξ; the commitment probability is nondecreasing and concave in ξ.

Proof Conditionally on the unmarked path, the integrated commitment intensity is

$$A = \int_0^{\tau_M^0} t_\nu(Z_s^0)\, \mathrm{d}s.$$

The probability of no clock ring on that path is e^{-A}. Thus its complement, averaged over paths, is exactly the probability of commitment before either size boundary. At fixed finite M, the unmarked chain is transient and has an exponentially decaying absorption-time tail, while $|Z_s^0| \leq M - 1$ before stopping; therefore its occupation-time functional has finite moments of all orders. Differentiation of $\mathbb{E}[e^{-\xi A_\kappa}]$ yields alternating signs, proving complete monotonicity and the stated consequences. □

Remark A.3 (The parameter being differentiated matters) Complete monotonicity in a commitment multiplier ξ does *not* imply monotonicity in repair ν, since changing ν changes the law of the unmarked path and possibly its commitment intensity. This distinction is essential for the pathway-switching question.

Appendix B Exact three-cell switch: generator and monotonicity certificate

For $M = 3$, order the states by size and then lexicographically, using configurations $z = (c, \ell, d)$ with $1 \le c + \ell + d \le 2$. With $b = (4, 1/10, 1/2)$, $\mu = (1, 1, 1)$, and $\theta = 50$, define the off-diagonal generator transitions

$$\begin{aligned}
(c,\ell,d) &\longrightarrow (c+1,\ell,d) && \text{at rate } 4c,\\
(c,\ell,d) &\longrightarrow (c,\ell+1,d) && \text{at rate } \ell/10,\\
(c,\ell,d) &\longrightarrow (c,\ell,d+1) && \text{at rate } d/2,\\
(c,\ell,d) &\longrightarrow (c-1,\ell,d) && \text{at rate } c,\\
(c,\ell,d) &\longrightarrow (c,\ell-1,d) && \text{at rate } \ell,\\
(c,\ell,d) &\longrightarrow (c,\ell,d-1) && \text{at rate } d,\\
(c,\ell,d) &\longrightarrow (c+1,\ell-1,d) && \text{at rate } \nu\ell.
\end{aligned}$$

Transitions to population sizes zero or three leave the transient state space. Add the diagonal commitment loss $-50d$ in every transient state and set $t(z) = 50d$. To obtain the perturbation matrix $B = \partial_\sigma Q_\sigma|_{\sigma=0}$, on a state of total size one, transfer rate $4c$ from the usual C-daughter birth transition to the D-daughter transition, and transfer rate $\ell/10$ from the usual L-daughter birth transition to the D-daughter transition. On states of size two every birth exits to establishment and makes no perturbation to the transient generator. This specifies all nine rows without ambiguity. Matrix inversion yields

$$q_D = [(-Q_0)^{-1}t]_{(0,0,1)}, \qquad A_3(\nu) = [(-Q_0)^{-1}B(-Q_0)^{-1}t]_{(0,1,0)}.$$

The accompanying Python scripts independently reconstruct the nine-state switch and five-state opportunity generators, verifying the stated numerical values using direct inversion and exact arithmetic.

B.1 Strict monotonicity and uniqueness of the conditional crossing

The conditional ratio in Lemma 3.1 has the following exact rational form. Write

$$\begin{aligned}
P_4(\nu) &= 40000\nu^4 + 2397000\nu^3 + 15730820\nu^2 + 16064982\nu + 1592283,\\
Q_3(\nu) &= 1000\nu^3 + 8300\nu^2 + 14430\nu + 6771.
\end{aligned}$$

The exact result is

$$D_L(\nu) = \frac{2575P_4(\nu)}{11752(5\nu + 263)Q_3(\nu)}. \tag{B1}$$

Symbolic differentiation and cancellation gives

$$D_L'(\nu) = \frac{7725\, P_6(\nu)}{904(5\nu + 263)^2 Q_3(\nu)^2}, \tag{B2}$$

where

$$\begin{aligned} P_6(\nu) = {} & 5000000\nu^6 + 592200000\nu^5 + 21379970000\nu^4 \\ & + 226294008000\nu^3 + 906566539700\nu^2 \\ & + 1252427781780\nu + 577213388709. \end{aligned}$$

Every coefficient of P_6 is strictly positive, so $D_L'(\nu) > 0$ for all $\nu \geq 0$. The equation $D_L(\nu) = 1$ is equivalently

$$44240000\nu^4 + 2593791000\nu^3 + 14005513900\nu^2 - 3630432990\nu - 16827515571 = 0.$$

Its unique nonnegative root, established by strict increase and endpoint comparison, is

$$\nu_c \approx 1.10964042855. \tag{B3}$$

These formulas can also be verified directly by exact rational arithmetic from the nine-state generator.

Appendix C Uniform finite-establishment estimates: detailed proof

We work with two one-driver phenotypes C, L and two two-driver phenotypes C_2, L_2. At both stages repair changes a damaged cell into the corresponding pristine phenotype at rate ν and preserves clone size. A cell of one-driver type i gives birth by budding at rate $b_i > 0$ and dies at rate $d_i > b_i$; a two-driver cell of type j gives birth at rate $B_j > 0$, dies at rate $\delta_j > 0$, and independently triggers whole-clone commitment at rate $\theta_j > 0$. These demographic and commitment rates are *fixed as ν varies*; there are no newly arising lesions anywhere within a clone, and no additional driver hits after one marked insertion. The parent survives a budding birth, including a marked birth. Throughout, assume $M \geq 2$ and $\nu \in [0, \infty)$. Write

$$a = \min_{i=C,L} (d_i - b_i) > 0, \quad b^* = \max_{i=C,L} b_i, \quad \theta_* = \min_{j=C,L} \theta_j > 0, \quad \delta^* = \max_{j=C,L} \delta_j.$$

For the unmarked one-driver process let $Z_t^{(1)}$ denote its cell-count vector; let τ_k be its first hitting time of total size k, and τ_0 its extinction time. For a post-insertion mixed process let $\tau_{D=0}$ be the first time the last two-driver cell disappears, and τ_{com} its first commitment time. Subsequent second-driver mutations are suppressed, so no commitment is possible after $\tau_{D=0}$.

Theorem C.1 (Uniform finite-establishment approximation) *Assume fixed positive birth and death rates for both one-driver and two-driver phenotypes, strictly subcritical one-driver types $0 < b_i < d_i$, and fixed two-driver commitment hazards $\theta_C, \theta_L > 0$. Repair acts only through size-preserving damaged-to-pristine phenotype transitions at rate ν. Suppose there is no new lesion input within clones, no further driver insertion after the marked second hit, and both pathways use the same total-size establishment threshold M, with boundary-crossing insertions excluded. Then there exist constants $c, C > 0$ and an integer $M_0 \geq 2$, independent of $\nu \geq 0$ and of M, such that for all $M \geq M_0$,*

$$\sup_{\nu \geq 0} |D_M(\nu) - D_\infty(\nu)| \leq C(M+1)e^{-cM}. \tag{C1}$$

The bound is uniform in repair intensity for fixed positive commitment hazards and fixed phenotype rates. It does not assert uniformity as a commitment hazard approaches zero, nor does it justify interchanging the rare-driver and large-establishment-threshold limits. In particular, if D_∞ lies strictly on one side of 1 at $\nu_1 < \nu_3$ and strictly on the other side at $\nu_2 \in (\nu_1, \nu_3)$, then every sufficiently large finite-M ratio has at least two distinct pathway crossings.

Proof The explicit estimate for the numerator and denominator of the pathway ratio is established in Theorem C.4 below. Its strict-sign consequence follows from continuity and the intermediate value theorem at the chosen finite witness rates. Only function-value convergence is used. □

Lemma C.2 (Uniform unmarked pre-hit occupation and entrance tails) *Starting from an arbitrary configuration with n one-driver cells, for every $t \geq 0$ and $\nu \geq 0$,*

$$\mathbb{E}|Z_t^{(1)}| \leq ne^{-at}, \qquad \mathbb{E}N_{\text{births}} \leq \frac{nb^*}{a}.$$

Choose $\eta > 1$ so close to one that $\eta b_i < d_i$ for $i = C, L$. For a single one-driver founder and $m \geq 1$,

$$\mathbb{P}_1(\tau_m < \tau_0) \leq \frac{\eta - 1}{\eta^m - 1}.$$

For $m = 1$ the entrance bound is the identity $\mathbb{P}_1(\tau_1 < \tau_0) = 1$; for $m \geq 2$ it follows from the supermartingale argument below. All inequalities are uniform in $\nu \geq 0$.

Proof For the unmarked generator $\mathcal{G}_1$, repair contributes nothing to the total population and $\mathcal{G}_1|z| = \sum_i (b_i - d_i)z_i \leq -a|z|$. Localization followed by Gronwall's inequality gives the first estimate; nonexplosion follows from domination of total births by a linear-rate Yule process at each fixed ν. The birth-count compensator is bounded by $b^* \int_0^\infty |Z_s^{(1)}| \, \mathrm{d}s$. Tonelli's theorem gives the second estimate. For $g(z) = \eta^{|z|} - 1$, size-preserving repair contributes zero, and

$$\mathcal{G}_1 g(z) = \eta^{|z|} \sum_i z_i(\eta - 1)\left(b_i - \frac{d_i}{\eta}\right) \leq 0.$$

Thus $g(Z^{(1)}_{t \wedge \tau_0 \wedge \tau_m})$ is a nonnegative supermartingale. Optional stopping at a bounded time followed by Fatou's lemma gives $(\eta^m - 1)\mathbb{P}_1(\tau_m < \tau_0) \leq \eta - 1$. □

Lemma C.3 (Killed post-hit boundary tail) *Choose the same $\eta > 1$ sufficiently close to one so that, additionally,*

$$B_j(\eta - 1) + \delta_j(\eta^{-1} - 1) < \theta_j \quad (j = C, L).$$

For every post-insertion configuration y with $n = |y| < M$ and at least one two-driver cell,

$$\mathbb{P}_y\{\tau_M < \tau_{\mathrm{com}} \wedge \tau_{D=0}\} \le \eta^{n-M}.$$

Consequently, writing $q_M(y)$ for commitment before the common size-M boundary and $q_\infty(y)$ for eventual commitment without that boundary,

$$0 \le q_\infty(y) - q_M(y) \le \eta^{n-M}.$$

The bound holds even if either two-driver phenotype is supercritical.

Proof Kill the mixed process at $\tau_{\mathrm{com}} \wedge \tau_{D=0}$ and set $h(y) = \eta^{|y|}$ while a two-driver cell remains and $h = 0$ on either killed state. On a state containing a two-driver cell, the killed generator satisfies

$$\mathcal{K}h(y) \le \eta^{|y|}\Bigg[\sum_{i=C,L} y_{1i}\{b_i(\eta-1) + d_i(\eta^{-1}-1)\} + \sum_{j=C,L} y_{2j}\{B_j(\eta-1) + \delta_j(\eta^{-1}-1) - \theta_j\}\Bigg] \le 0.$$

If a death removes the last two-driver cell, its actual jump is to zero for h, so the displayed ordinary-death contribution is an upper bound. Repair preserves total size and the existence of a two-driver cell, and hence contributes exactly zero, regardless of ν. A nonnegative supermartingale stopped at $t \wedge \tau_M \wedge \tau_{\mathrm{com}} \wedge \tau_{D=0}$ therefore has expectation at most η^n. On the event that τ_M occurs first its value is η^M; letting $t \to \infty$ by Fatou proves the asserted probability bound. Coupling the bounded and unbounded post-hit processes on the same clocks until τ_M shows that their commitment indicators can differ only on this boundary hitting event. After $\tau_{D=0}$ no new two-driver cell can arise. □

Theorem C.4 (Uniform finite-size insertion and ratio errors) *Let $A_{i,M}$ be the coefficient of a single marked second-driver insertion that eventually commits before extinction or size M, starting from one one-driver founder of type i. Define $A_{i,\infty}$ with no size boundary. Take $m = \lfloor M/2 \rfloor$. Then for both founder types and every $\nu \ge 0$,*

$$0 \le A_{i,\infty} - A_{i,M} \le \Delta_M := \frac{b^*}{a}\left[\eta^{m-M} + m\frac{\eta - 1}{\eta^m - 1}\right].$$

For either directly founded two-driver type i, $0 \le q_{i,\infty} - q_{i,M} \le \eta^{1-M}$. Let $s_i(\nu) \ge 0$ be any common normalized founder weights, $\sum_i s_i(\nu) = 1$, used in both pathway comparisons. Define

$$D_M = \frac{2\sum_i s_i A_{i,M}}{\sum_i s_i q_{i,M}}, \qquad D_\infty = \frac{2\sum_i s_i A_{i,\infty}}{\sum_i s_i q_{i,\infty}}.$$

Set $q_* = \theta_*/(\theta_* + \delta^*) > 0$. Choose an integer $M_0 \ge 2$, depending only on η and q_*, such that $\eta^{1-M_0} \le q_*/2$. Then, for every $M \ge M_0$,

$$\sup_{\nu \ge 0} |D_M - D_\infty| \le \frac{4\Delta_M}{q_*} + \frac{4(b^*/a)\eta^{1-M}}{q_*^2} \le C(M+1)e^{-cM},$$

where one may take

$$c = \tfrac{1}{2}\log\eta, \qquad C = \frac{4b^*}{a}\left(\frac{1 + \eta^{3/2}/2}{q_*} + \frac{1}{q_*^2}\right).$$

Proof Construct the finite and unbounded marked-insertion potentials from the same unmarked one-driver genealogy. Each ordinary budding birth of the unmarked genealogy represents one potential rare insertion; under a marked insertion its newborn daughter is instead a two-driver cell, with the one-driver parent retained. Only births from pre-hit size at most $M-2$ are eligible in the finite process. The unbounded potential counts all ordinary births and weights each by the unbounded post-hit success probability. Hence the finite potential is no larger.

Partition the unmarked birth events into (i) those occurring up to *and including* the birth that first attains size m, if that birth occurs, and (ii) births strictly after the first hitting time τ_m. In group (i) the pre-birth size is at most $m-1$, so the hypothetical post-hit state has at most m cells. By Lemma C.3, replacing its unbounded commitment probability by the finite-boundary probability loses at most η^{m-M}; the mean number of such births is at most b^*/a. For group (ii), bound the entire remaining unbounded marked-insertion contribution by the number of ordinary births, since success probabilities are at most one. Lemma C.2 and the strong Markov property bound that number in expectation by

$$\mathbb{P}_1(\tau_m < \tau_0)\frac{mb^*}{a} \le \frac{mb^*}{a}\frac{\eta-1}{\eta^m-1}.$$

This also accounts for all births omitted by the finite boundary, including the boundary-crossing birth itself. Adding the two losses gives Δ_M. The direct-founder inequality is Lemma C.3 applied at $n=1$.

To establish a denominator bound independent of repair, follow the original two-driver cell. While it remains alive its commitment hazard is at least θ_* and its death hazard at most δ^*; repair and budding births do not remove that original cell. By a comparison of competing clocks, the probability it commits before it dies is at least q_*, at every repair intensity. Thus $\sum_i s_i q_{i,\infty} \ge q_*$, and the finite denominator is at least $q_* - \eta^{1-M} \ge q_*/2$. Also $A_{i,\infty} \le b^*/a$. Fraction subtraction therefore gives

$$\begin{aligned}|D_M - D_\infty| &\le \frac{2\Delta_M}{q_*/2} + \frac{2(b^*/a)\eta^{1-M}}{(q_*/2)q_*} \\ &\le \frac{4\Delta_M}{q_*} + \frac{4(b^*/a)\eta^{1-M}}{q_*^2}.\end{aligned}$$

Finally, for $m = \lfloor M/2 \rfloor$ one has $\eta^{m-M} \le \eta^{-M/2}$ and $(\eta-1)/(\eta^m-1) \le \eta^{1-m}$, giving

$$\Delta_M \le \frac{b^*}{a}(1 + \tfrac{1}{2}M\eta^{3/2})\eta^{-M/2}.$$

For $M \ge 2$, $\eta^{1-M} \le \eta^{-M/2}$, so the explicit C and c above complete the proof. No estimate for D'_M is used or claimed. □

Consequences and order of limits. At every fixed M, the rare driver mutation coefficient is obtained first by sending $V \downarrow 0$. Only then is $M \to \infty$ taken. Uniform convergence of the resulting *function values* transfers strict comparison signs at finitely many repair rates to every sufficiently large finite M, and yields intervening parity crossings by continuity. Four strict value witnesses similarly imply a local minimum followed by a later local maximum. These statements require neither derivative convergence nor uniqueness of finite-M roots.

Appendix D Exact finite-clone parity certificates and sensitivity intervals

For $M = 9, 10$, the rational-residual procedure implemented in the companion reproducibility code bounds inverse-generator numerical solutions, giving strict parity signs at the repair rates stated in Section 5.1 and Table D1. At $M = 9$, the four certified repair rates $0, 1/10, 6, 1000$ show a strict down–up–down sequence, proving at least one interior minimum followed by a maximum; the parity signs at $0, 6, 1000$ prove at least two crossings. At $M = 10$, the certified values at $0, 1/10, 10, 100$ likewise show a strict down–up–down sequence, proving at least two extrema, while the parity signs prove at least two crossings. These are existence certificates, not complete polynomial root counts; in particular no exact number of crossings or stationary points is asserted. The named source scripts are available in the public reproducibility release cited below. Neither certification requires the asymptotic boundary bound.

For clarity, the residual enclosures use the following finite-state comparison. Write K for the killed post-hit generator, Q for the killed pre-hit generator, $B \geq 0$ for the mutation-insertion matrix, and $q = (-K)^{-1}t$, $A = (-Q)^{-1}Bq$. For rational candidate vectors $\widetilde{q}, \widetilde{A}$ define the *exact* rational residuals $r_q = t + K\widetilde{q}$ and $r_A = B\widetilde{q} + Q\widetilde{A}$. Because $(-K)\mathbf{1} \geq \theta_* \mathbf{1}$ on the post-hit states, $(-K)^{-1} \geq 0$ gives

$$\|q - \widetilde{q}\|_\infty \leq E_q := \|r_q\|_\infty / \theta_*.$$

In the pre-hit process $(-Q)(|z|/a) \geq \mathbf{1}$ with $a = \min(d_C - b_C, d_L - b_L) > 0$; hence

$$\|A - \widetilde{A}\|_\infty \leq E_A := \frac{M-1}{a}\big(\|r_A\|_\infty + \|B\|_\infty E_q\big).$$

The source equilibrium root is bracketed by exact rational sign tests of its defining quadratic. Combining these enclosures with positive numerator and denominator bounds produces the printed rigorous ratio intervals. The added $M = 9$, $\nu = 1/10$ certificate uses precisely this procedure, with no assumption that the floating-point candidate vectors are accurate.

D.1 Additional exact sensitivity certificates

For the parameter variation in Table 2, the resident source uses $(r_L, g_L, \varepsilon) = (1/2, 2, 1/4)$. For $0 < U < 3/4$ the physical equilibrium satisfies

$$(3 - 2U + 4\nu)x^2 - (3 - 4U)x - 4\nu = 0, \qquad \phi(U, \nu) = \frac{1 + (2U - 1)x}{1 + x}. \tag{D1}$$

In each table row the two outer witnesses lie above pathway parity and the intermediate witness lies below. The exact rational-residual enclosure method described above applies to all three rows of Table 2, without assumptions about the number of additional extrema or crossings. The first row's interval certificates are recorded in the verification code; the two additional rows have the following outward-rounded bounds.

Table D1 Outward-rounded exact rational-residual bounds for the principal finite-clone witnesses, with $(U, \theta_L) = (1/2, 10^{-6})$ and the common rates of Section 5.1. The strict signs relative to 1 prove at least two crossings for each threshold; the wider interval at $M = 9$, $\nu = 1000$ remains strictly below parity.

M	ν	Certified $D_M(\nu)$ interval
10	0	(1.007908453, 1.007928675)
	0.1	(0.962275566, 0.962290025)
	10	(1.016630416, 1.016652652)
	100	(1.009333701, 1.009526991)
9	0	(0.998465662, 0.998477940)
	0.1	(0.964525804, 0.964535790)
	6	(1.009198937, 1.009210991)
	1000	(0.998025412, 0.999254543)

Table D2 Strict rational-residual witness intervals for the two less extreme $M = 10$ parameter combinations. Outward decimal rounding is shown; the full scripts carry the rational bounds. These witnesses prove existence of at least two parity crossings per row of Table 2, not an exact total crossing count.

(U, θ_L)	ν	Certified $D_{10}(\nu)$ interval
(0.20, 0.001)	0	(1.006962844795, 1.006962853976)
	0.1	(0.995593752914, 0.995593763522)
	10	(1.011302751510, 1.011302777955)
(0.10, 0.010)	0	(1.004034950187, 1.004034950969)
	0.2	(0.999464300645, 0.999464301577)
	10	(1.009540598404, 1.009540600411)

The original rate-matrix constructors and rational-residual certification scripts are available in the public reproducibility release (see Code availability). Numerical curves between certified points are illustrative only. These strict interval enclosures demonstrate local persistence of the displayed parity pattern under sufficiently small parameter changes, not a quantified biologically plausible robustness region.

Statements and Declarations

Funding

No funds, grants, or other support were received for conducting this study.

Competing interests

The author has no relevant financial or non-financial interests to disclose.

Data availability

No external datasets were used in this study. Numerical results were generated from the mathematical models specified in the manuscript.

Code availability

The Python code used for numerical verification, rational-residual certification, and figure generation is publicly available in the GitHub repository https://github.com/Priceeqn1/repair-filtered-mutational-histories, under the `paper-ii-driver-tunneling/` directory. The version associated with this manuscript is archived on Zenodo at https://doi.org/10.5281/zenodo.23904950 (release `paper-ii-v1.0.1`). The archive also contains the independently maintained reproducibility materials for *Repair Filters Mutational Histories by Event Count: A Pascal Ordering Theorem with a Lineage Interpretation*.

Author contributions

The author conceived the study, developed the mathematical model, performed the analytical and numerical work, and wrote the manuscript.

Ethics approval

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Materials availability

Not applicable.

Use of generative AI and AI-assisted technologies

During preparation of this manuscript, the author used ChatGPT (OpenAI) and Claude (Anthropic) for language editing, manuscript organization, mathematical discussion, proof checking, computational and coding assistance, and independent numerical verification. The author reviewed the resulting material and takes full responsibility for the mathematical claims, computational results, interpretations, and conclusions.

References

Armitage, P., Doll, R.: The age distribution of cancer and a multi-stage theory of carcinogenesis. British Journal of Cancer **8**, 1–12 (1954) https://doi.org/10.1038/bjc.1954.1

Bozic, I., Antal, T., Ohtsuki, H., Carter, H., Kim, D., Chen, S., Karchin, R., Kinzler, K.W., Vogelstein, B., Nowak, M.A.: Accumulation of driver and passenger mutations during tumor progression. Proceedings of the National Academy of Sciences USA **107**, 18545–18550 (2010) https://doi.org/10.1073/pnas.1010978107

Graser, C.J., Wu, W., Christini, C., Petljak, M., Michor, F.: Punctuated mutagenesis promotes multi-step evolutionary adaptation in human cancers. eLife **14**, RP108058 (2026) https://doi.org/10.7554/eLife.108058.3

Haeno, H., Iwasa, Y., Michor, F.: The evolution of two mutations during clonal expansion. Genetics **177**(4), 2209–2221 (2007) https://doi.org/10.1534/genetics.107.078915

Hardisty, B.: Repair-mediated path filtering of mutation load in multilocus homeostatic populations. bioRxiv. Preprint (2026). https://doi.org/10.64898/2026.10.11.758278

Iwasa, Y., Michor, F., Nowak, M.A.: Stochastic tunnels in evolutionary dynamics. Genetics **166**, 1571–1579 (2004) https://doi.org/10.1534/genetics.166.3.1571

Knudson, A.G.: Mutation and cancer: Statistical study of retinoblastoma. Proceedings of the National Academy of Sciences USA **68**, 820–823 (1971) https://doi.org/10.1073/pnas.68.4.820

Komarova, N.L., Sengupta, A., Nowak, M.A.: Mutation-selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. Journal of Theoretical Biology **223**, 433–450 (2003) https://doi.org/10.1016/S0022-5193(03)00120-6

Martincorena, I., Roshan, A., Gerstung, M., *et al.*: High burden and pervasive positive selection of somatic mutations in normal human skin. Science **348**, 880–886 (2015) https://doi.org/10.1126/science.aaa6806

Martincorena, I., Fowler, J.C., Wabik, A., *et al.*: Somatic mutant clones colonize the human esophagus with age. Science **362**, 911–917 (2018) https://doi.org/10.1126/science.aau3879

Moolgavkar, S.H., Knudson, A.G.: Mutation and cancer: A model for human carcinogenesis. Journal of the National Cancer Institute **66**, 1037–1052 (1981) https://doi.org/10.1093/jnci/66.6.1037

Plotkin, J.B., Nowak, M.A.: The different effects of apoptosis and DNA repair on

tumorigenesis. Journal of Theoretical Biology **214**, 453–467 (2002) https://doi.org/10.1006/jtbi.2001.2471

Vermeulen, L., Morrissey, E., van der Heijden, M., *et al.*: Defining stem cell dynamics in models of intestinal tumor initiation. Science **342**, 995–998 (2013) https://doi.org/10.1126/science.1243148

Weissman, D.B., Desai, M.M., Fisher, D.S., Feldman, M.W.: The rate at which asexual populations cross fitness valleys. Theoretical Population Biology **75**(4), 286–300 (2009) https://doi.org/10.1016/j.tpb.2009.02.006

Zhang, R., Bozic, I.: Accumulation of oncogenic mutations during progression from healthy tissue to cancer. Bulletin of Mathematical Biology **86**, 142 (2024) https://doi.org/10.1007/s11538-024-01372-3