Modeling cancer evolution and inferring its parameters

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Einar Bjarki Gunnarsson

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Kevin Leder
Jasmine Foo

June, 2022

Acknowledgements

I would like to express my sincere gratitude to my advisors, Kevin Leder and Jasmine Foo, for their guidance and support throughout my Ph.D. studies. From my first meetings with them, they have been patient and encouraging, granting me independence and autonomy while always being available for detailed project discussions and to offer valuable feedback and insights. I could not have asked for better advisors.

I would like to thank my Ph.D. committee members, Sherwin Doroudi and William Cooper, for their support and mentorship throughout the years, and Saif Benjaafar for his patience, guidance and support during my first year. I am also grateful to Shuzhong Zhang for his course on optimization and for his advice and encouragement during my first steps in the Ph.D. program. Finally, I would like to thank Arnab Sen for his courses on probability theory and for serving on my master's degree committee.

I gratefully acknowledge support from the College of Science and Engineering Graduate Fellowship, the Leifur Eiríksson Foundation Fellowship, the University of Minnesota Doctoral Dissertation Fellowship, NSF grants CMMI-1552764 (PI: Kevin Leder) and DMS-1349724 (PI: Jasmine Foo), and the Norwegian Centennial Chair grant.

I would like to thank Zicheng and Anna for their help and advice throughout the years, and Chenyu, Aaron and my other fellow research group members for the weekly catch-ups. I am also grateful to my fellow students in the ISyE department, including my roomate Alex, my classmates Bingnan, Fatemeh, Jazeem, Jing, Kang, P.A. and Xiaotang, and my office mates Akshit and Chuan, for sharing the pains and pleasures of Ph.D. life.

Finally, I would like to thank my friends and family in Iceland, especially my parents Sigrún and Gunnar for their endless help and support and for the royal welcome I receive every time I return home. Without them, none of this would have been possible.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Cancer is a group of diseases characterized by uncontrolled proliferation of abnormal cells. It is a leading cause of death worldwide, claiming almost 10 million lives in 2020 [146]. The initiation of cancer usually involves a series of mutations in genes responsible for regulating the cell cycle [4, 97]. As the cancer-initiating cell expands into a tumor, the tumor cells continue to accumulate mutations, which induces substantial genetic heterogeneity within the tumor [25, 40, 111, 154].

In recent years, it has been increasingly recognized that epigenetic mechanisms, which are chemical changes to DNA or the chromatin structure which houses DNA, play an equally important role in cancer evolution [12, 21, 22, 62, 84, 136]. Epigenetic mechanisms can either inhibit or encourage the expression of cancer-related genes, which can have the same effect as genetic mutations that inactivate or amplify the genes. The difference is that epigenetic modifications are reversible and they can occur much faster than genetic mutations [22, 119]. Due to their reversible and rapid nature, epigenetic mechanisms can enable cancer cells to switch dynamically between two or more phenotypic states, which commonly show differential responses to drug treatments [77, 86, 95, 126, 133, 138, 139, 144]. Together, genetic and non-genetic mechanisms induce significant heterogeneity between tumor cells, which makes the tumor highly adaptable to changing conditions and primed to escape drug treatments.

This dissertation consists of four projects which each involves using a mathematical model to study the evolutionary dynamics of cancer. The projects range from addressing the process of cancer initiation to the evolution of drug resistance, and they take both

genetic and non-genetic perspectives. First, in Chapter 2, we study the dynamics of cancer initiation in multilayered tissue under a two-step mutational model of cancer. In Chapter 3, we study the accumulation of neutral mutations during tumor progression, which are mutations that do not affect the division rate of tumor cells. In Chapter 4, we study the role of phenotypic switching in the evolution of stable drug resistance. Finally, in Chapter 5, we develop a statistical framework for inferring the rates of cell division, cell death and phenotypic switching in cancer. For the remainder of this section, we summarize the contents and contributions of each chapter.

Over 80% of cancers originate from the epithelium, which covers the outer and inner surfaces of organs and blood vessels. In stratified epithelium, the bottom layers are occupied by stem and stem-like cells that continually divide and replenish the upper layers. In Chapter 2, we study the spread of premalignant mutant clones and cancer initiation in stratified epithelium, using the biased voter model on stacked two-dimensional lattices [65]. Our main result is an estimate of the propagation speed of a premalignant mutant clone, which is asymptotically precise in the cancer-relevant weak-selection limit. We use our main result to study cancer initiation under a two-step mutational model of cancer, which includes computing the distributions of the time of cancer initiation and the size of the premalignant clone giving rise to cancer. This work quantifies the effect of epithelial tissue thickness on the process of carcinogenesis, thereby contributing to an emerging understanding of the spatial evolutionary dynamics of cancer.

In Chapter 3, we study mutation accumulation in neutrally evolving tumors [76]. More precisely, we study the site frequency spectrum (SFS) of neutral mutations during tumor progression, which is a popular summary statistic of genomic data. While the SFS of a constant-sized population undergoing neutral mutations has been extensively studied in population genetics, the rapidly growing amount of cancer genomic data has attracted interest in the spectrum of an exponentially growing population. Recent theoretical results have generally dealt with special or limiting cases, such as considering only cells with an infinite line of descent, assuming deterministic tumor growth, or taking large-time or large-population limits. In this work, we derive exact expressions for the expected SFS of a cell population that evolves according to a stochastic branching process, first for cells with an infinite line of descent and then for the total population, evaluated either at a fixed time (fixed-time spectrum) or at the stochastic time at which

the population reaches a certain size (fixed-size spectrum). We find that while the rate of mutation scales the SFS of the total population linearly, the rates of cell birth and cell death change the shape of the spectrum at the small-frequency end, inducing a transition between a $1/j^2$ power-law spectrum and a $1/j$ spectrum as cell viability decreases. We show that this insight can in principle be used to estimate the ratio between the rate of cell death and cell birth, as well as the mutation rate, using the site frequency spectrum alone. Although the discussion is framed in terms of tumor dynamics, our results apply to any exponentially growing population of individuals undergoing neutral mutations.

In Chapter 4, we develop a mathematical model to study how phenotypic switching affects the evolution of stable drug resistance in cancer [75]. We highlight unique features of non-genetic resistance, probe the evolutionary consequences of epigenetic drugs, and explore potential therapeutic strategies. We find that even short-term epigenetic modifications and stochastic fluctuations in gene expression can drive long-term drug resistance in the absence of any bona fide resistance mechanisms. We also find that an epigenetic drug that slightly perturbs the average retention of the resistant phenotype can turn guaranteed treatment failure into guaranteed success. Finally, we find that combining an epigenetic drug with an anti-cancer agent can significantly outperform monotherapy, and that treatment outcome is heavily affected by drug sequencing.

Finally, in Chapter 5, we propose a rigorous statistical framework for estimating the rates of cell proliferation and phenotypic switching in cancer, using data from commonly performed cell line experiments, where phenotypes are sorted and expanded in culture. The framework explicitly models the stochastic dynamics of cell division, cell death and phenotypic switching, and it provides likelihood-based confidence intervals for the model parameters. The input data can be either the fraction of cells or the number of cells in each state at one or more time points. Through a combination of theoretical analysis and numerical simulations, we show that when cell fraction data is used, the rates of switching may be the only parameters that can be estimated accurately. On the other hand, using cell number data enables accurate estimation of the net division rate for each phenotype, and it can even enable estimation of the state-dependent rates of cell division and cell death. We conclude by applying our framework to a publicly available dataset and discussing the differences between our approach and existing methods.

Chapter 2

Spread of premalignant mutant clones and cancer initiation in multilayered tissue

2.1 Introduction

According to the widely held multi-stage model of carcinogenesis, cancer arises due to the accumulation of genetic mutations that culminate in malignant cells able to proliferate uncontrollably [4, 5, 96, 97]. Each mutation in this process can afford a small selective advantage, which can allow premalignant cells to expand into clones or "fields" that are further along the evolutionary pathway to cancer than normal cells and thus predisposed to becoming cancerous [39]. The notion that cancer arises on the background of premalignant field expansion is referred to as "field cancerization" or "the cancer field effect". It has important clinical implications, since tumors surrounded by premalignant patches are at increased risk of recurrence following cancer treatment [18, 30]. Premalignant fields often appear histologically normal, making them difficult to distinguish from healthy tissue. This suggests that a mathematical understanding of the spatial evolutionary dynamics of cancer initiation can yield valuable insights into treatment decision-making, including optimal surgical excision margins and post-treatment surveillance protocols.

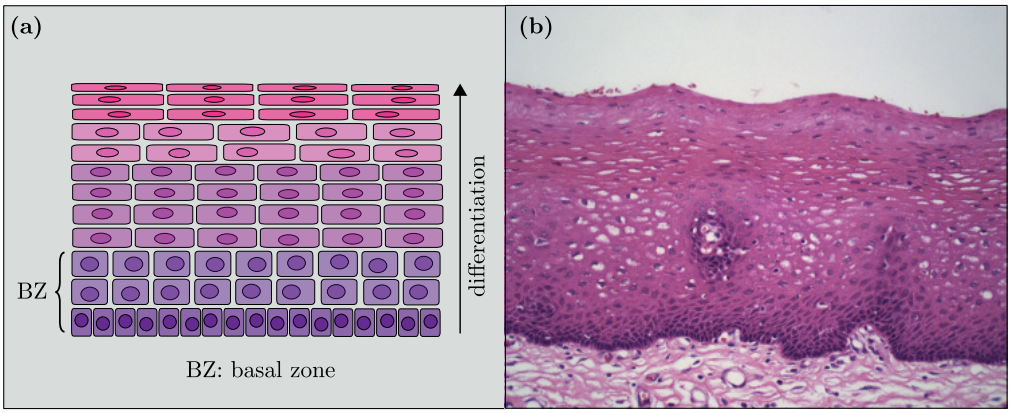

Figure 2.1: **(a)** In stratified squamous epithelium of the esophagus, a basal layer of stem cells and 2-3 layers of proliferative basaloid cells form the basal zone, which continually replenishes the upper layers with differentiated cells. **(b)** Histology of normal stratified squamous epithelium of the esophagus. Figure 4.6 on page 78 of [32].

Over 80% of human cancers originate from the epithelium, which lines the outer and inner surfaces of organs and blood vessels [80]. Simple epithelium consists of a single layer of proliferating cells, whereas in stratified epithelium, stem and stem-like cells proliferate along the bottom layers and continually replenish the upper layers with differentiated cells that lose their ability to proliferate. For example, in *stratified squamous epithelium* of the esophagus, a basal layer of stem cells and 2-3 layers of proliferative basaloid cells form the *basal zone*, which accounts for less than 30% of total epithelial thickness. As cells move upward, they become terminally differentiated keratinocytes with small nuclei that flatten out and eventually get shed at the top layer (Fig. 2.1) [32, 69].[1] Since the accumulation and spread of mutations is driven by the proliferating basal and basaloid cells, the basal zone is the appropriate setting to study the process of carcinogenesis in stratified epithelium.

In this work, we study the spread of premalignant mutant fields in epithelial basal zones, and we examine the effect of basal zone geometry on the process of cancer initiation. Our main result determines the propagation speed of a premalignant mutant clone as a function of a small mutant selective advantage and the number of layers in the basal

[1] Figure 1b is reprinted from GERD, Parakrama T. Chandrasoma, Chapter 4 – Histologic Definition and Diagnosis of Epithelia in the Esophagus and Proximal Stomach, pp. 73-107, Copyright 2018, with permission from Elsevier.

zone, which enables comparison of the evolutionary dynamics between different types of epithelial cancers. We employ a spatially explicit model of cell division and replacement, where cells live on a set of stacked two-dimensional integer lattices, representing a multilayered basal zone. The model dynamics are as follows: Cells of two types, normal and mutant, are arranged on the stacked lattices, with mutant cells dividing more frequently than normal cells. Upon cell division, one daughter cell stays put, and the other replaces a neighboring cell chosen uniformly at random. This model was originally proposed by Williams & Bjerknes [162] in the context of a single two-dimensional epithelial basal layer, and it arose independently within the field of interacting particle systems as the *biased voter model*. Bramson and Griffeath [19, 20] showed in 1980-1981 that under the biased voter model, an advantageous mutant clone eventually assumes a convex, symmetric shape whose diameter grows linearly in time. The Bramson-Griffeath shape theorem extends naturally to our stacked-lattice setting, and it plays a central role in the derivation of our main result.

Once we have determined the propagation speed of premalignant mutant clones in epithelial basal zones, we consider the implications of our result for the dynamics of cancer initiation. The process of carcinogenesis under the multi-step model of cancer has already been well-studied in the non-spatial, homogeneously mixed setting, see e.g. the books by Nowak [121] and Wodarz and Komarova [164]. In the spatial setting, Komarova [98] has analyzed the time of cancer initiation on a one-dimensional lattice under a two-step model of cancer, assuming a neutral or deleterious first-step mutation. Durrett and Moseley [51] extended Komarova's work to two and three dimensions assuming a neutral first step. Durrett, Foo and Leder [46] considered the case of a small selective advantage (weak selection), and they derived the distribution of the time of cancer initiation under a two-step model of cancer in certain parameter regimes. Foo, Leder and Schweinsberg obtained more complete results in [66], where they also studied cancer initiation under a general k-step model. In [63], Foo, Leder and Ryser studied field cancerization under a two-step model, which included computing size-distributions of premalignant fields at the time of cancer initiation. Upon establishing our main result of premalignant mutant propagation, we will adapt analysis from [46] and [63] to gain insights into how cancer initiation and field cancerization is affected by the specific geometric setting of a multilayered basal zone.

The rest of the chapter is organized as follows. In Section 2.2, we propose a model of the spread of premalignant mutant fields along epithelial basal zones and state our main result of their long-run propagation speed. In Section 2.3, we present an outline of the proof of the main result, in which we exploit a duality between the biased voter model and a system of branching coalescing random walks. We use coupling to set up an approximation scheme, based on a pruning procedure of Durrett and Zähle [54], which culminates in simple, coalescence-free, branching random walks that are more readily analyzed. We state ten technical lemmas, the most important of which are Lemmas 2.7 and 2.10 that provide lower and upper bounds for the propagation speed of a premalignant mutant clone. In Section 2.4, we show how our main result follows from these two lemmas, and in Section 2.5, we discuss the implications of our main result for cancer initiation and field cancerization in multilayered epithelial basal zones. In Section A.1, we present proofs of the ten lemmas, and we discuss how the Bramson-Griffeath shape theorem in two dimensions can be extended to our stacked-lattice setting.

Notation. In our exposition, we make use of the following asymptotic notation, where a is taken to be 0 or ∞ depending on the context:

$f(x) = o(g(x))$ and $g(x) = \omega(f(x))$ as $x \to a$ if $f(x)/g(x) \to 0$ as $x \to a$.

$f(x) = O(g(x))$ and $g(x) = \Omega(f(x))$ as $x \to a$ if $\limsup_{x\to a} |f(x)/g(x)| < \infty$.

$f(x) = \Theta(g(x))$ as $x \to a$ if $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$ as $x \to a$.

$f(x) \sim g(x)$ as $x \to a$ if $f(x)/g(x) \to 1$ as $x \to a$.

2.2 Model of spread of premalignant mutant fields and statement of main result

Let $\mathbb{Z}_w := \mathbb{Z} \mod w$ denote the additive group of integers modulo $w \geq 1$. We represent an epithelial basal zone as the set $\mathbb{Z}^2 \times \mathbb{Z}_w$ of w layers of two-dimensional integer lattices, with a periodic boundary condition along the third dimension. For each site $x \in \mathbb{Z}^2 \times \mathbb{Z}_w$, we define its *neighborhood* as $\mathcal{N}(x) := \{x \pm e_i : i = 1, 2\}$ for $w = 1$ and $\mathcal{N}(x) := \{x \pm e_i : i = 1, 2, 3\}$ for $w > 1$, where e_i is the i-th unit vector, and addition along the third coordinate is carried out modulo w. Note that for $w = 1$, each site has

four neighbors, for $w = 2$, each site has five neighbors, and for $w > 2$, each site has six neighbors.

To model the spread of premalignant mutant fields, we define the biased voter model on $\mathbb{Z}^2 \times \mathbb{Z}_w$ as follows: Each site in $\mathbb{Z}^2 \times \mathbb{Z}_w$ is occupied by either a type-0 cell, representing a normal cell, or a type-1 cell, representing a premalignant mutant cell. Type-1 cells have a fitness advantage $\beta > 0$ over type-0 cells, meaning that type-1 cells divide at exponential rate $1+\beta$, while type-0 cells divide at exponential rate 1. Upon cell division at a site $x \in \mathbb{Z}^2 \times \mathbb{Z}_w$, one daughter cell stays put, while the other daughter cell replaces a neighboring cell at a site $y \in \mathcal{N}(x)$ chosen uniformly at random (Fig. 2.2a).

We assume throughout that the fitness advantage β is small. In [14], for example, Bozic et al. show that data on multiple cancer types (glioblastoma, pancreatic cancer and colon cancer) is consistent with an average selective advantage of $\beta = 0.004$ per mutational step. They further argue that this estimate should be more broadly relevant across cancer types, given the considerable overlap of pathways through which the selective mutations act.

A few comments are in order on the biological significance of our modeling assumptions. First, we allow cells on the top layer of $\mathbb{Z}^2 \times \mathbb{Z}_w$ to replace cells on the bottom layer, and vice versa (Fig. 2.2a). This assumption simplifies the analysis, as it means that the top and bottom layers have the same neighborhood structure as the intermediate layers, but it is not biologically realistic. In Appendix A.2, we use simulation to show that this assumption does not significantly affect type-1 propagation when β is small. Secondly, the model dynamics are driven by cell division, with cell division preceding cell death, and the model assumes that type-0 and type-1 cells are equally likely to be replaced by a dividing cell. When β is small, we expect that the exact dynamics of cell division and cell death are not important for long-run type-1 propagation, and we plan to discuss this in future work. Finally, we assume that cells are arranged on the lattice $\mathbb{Z}^2 \times \mathbb{Z}_w$ throughout, whereas in real tissue, the spatial structure may be different, and the structure may change as the premalignancy progresses. As with any model of biological or physical phenomena, our model is simplified and not intended to capture the full complexity of the system. Our focus here is on two important parameters of the process: The fitness advantage $\beta > 0$ of mutant cells and the tissue thickness $w \geq 1$.

Let ξ_t^A denote the set of sites in $\mathbb{Z}^2 \times \mathbb{Z}_w$ occupied by type-1 cells at time t, given

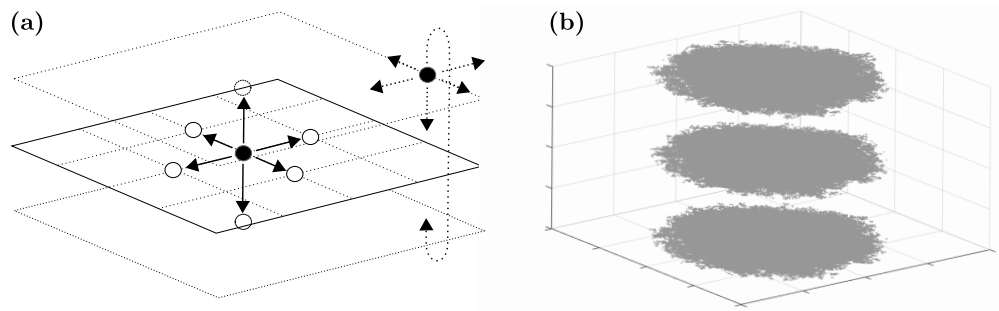

Figure 2.2: **(a)** Model dynamics for the $w = 3$ case (basal zone consists of three layers). When the black cell divides, one daughter cell stays put, and the other replaces one of six neighboring cells chosen uniformly at random. A cell on the top (resp. bottom) layer can replace a cell on the same layer, the layer immediately below (resp. above) or on the bottom (resp. top) layer. **(b)** Simulation of the model for $w = 3$ and $\beta = 0.1$, stopped at 50,000 cells. By the Bramson-Griffeath shape theorem (2.4), the premalignant population eventually takes on a convex, symmetric shape, and representative simulations suggest that this limiting shape is a union of two-dimensional Euclidean disks.

the initial condition $\xi_0^A = A$ with $A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$. Our baseline assumption is that the system starts out with a single type-1 cell at the origin, i.e. $A = \{0\}$. Define

$$\tau_\varnothing^0 = \tau_\varnothing^0(\beta) := \inf\{t \geq 0 : \xi_t^0 = \varnothing\}$$

as the time of extinction of the process starting from the origin, with $\inf \varnothing = \infty$. The discrete-time jump process embedded in $(|\xi_t^0|)_{t\geq 0}$, with $|\cdot|$ denoting cardinality, is a simple, biased random walk on the nonnegative integers with absorption at 0. The walk moves up with probability $(1+\beta)/(2+\beta)$ and down with probability $1/(2+\beta)$. It follows by the gambler's ruin formula that a type-1 mutant expands into a successful type-1 clone with probability

$$\mathbb{P}(\tau_\varnothing^0 = \infty) = \beta/(1+\beta). \tag{2.1}$$

If the mutant is successful, the Bramson-Griffeath shape theorem on $\mathbb{Z}^2$ [19, 20] can be extended to show that the mutant clone on $\mathbb{Z}^2 \times \mathbb{Z}_w$ eventually assumes a convex, symmetric shape whose diameter grows linearly in time. To carry out the extension, we need to introduce a notion of spatial scaling and distance in our stacked-lattice setting.

To that end, we go to the larger space $\mathbb{R}^2 \times \mathbb{Z}_w$. For $t \in \mathbb{R}$ and $A \subseteq \mathbb{R}^2 \times \mathbb{Z}_w$, we define the scalar multiplication operation tA as multiplication along the first two coordinates:

$$tA := \{(tx_1, tx_2, x_3) : (x_1, x_2, x_3) \in A\}. \tag{2.2}$$

The distance of a point $x = (x_1, x_2, x_3) \in \mathbb{R}^2 \times \mathbb{Z}_w$ from the origin is defined in terms of its two-dimensional projection as

$$||x|| = ||(x_1, x_2, x_3)|| := \sqrt{x_1^2 + x_2^2}. \tag{2.3}$$

Note that $\mathbb{R}^2 \times \mathbb{Z}_w$ is not a vector space since $(s+t)x = sx + tx$ does not hold in general for $s, t \in \mathbb{R}$ and $x \in \mathbb{R}^2 \times \mathbb{Z}_w$. However, it does hold up to addition by a vector parallel to $e_3 = (0, 0, 1)$, which is sufficient for establishing a shape theorem (Section A.1.11).

Using the above definitions, we can modify Bramson and Griffeath's arguments to show that there exists a set $D = D(\beta)$ on $\mathbb{R}^2 \times \mathbb{Z}_w$ so that for any $\varepsilon > 0$,

$$\mathbb{P}\big(\exists t_* < \infty : (1-\varepsilon)tD \cap (\mathbb{Z}^2 \times \mathbb{Z}_w) \subseteq \xi_t^0 \subseteq (1+\varepsilon)tD,\ t \geq t_* \ \big|\ \tau_\varnothing^0 = \infty\big) = 1. \tag{2.4}$$

The set D can be written as $D = \bigcup_{i \in \mathbb{Z}_w} (X \times \{i\})$, where $X \subseteq \mathbb{R}^2$ is convex, and X has the same symmetries as those of $\mathbb{Z}^2$ that leave the origin fixed. For example, $(x_1, x_2) \in X$ implies $(-x_1, x_2), (x_1, -x_2) \in X$ (reflection across an axis) and $(-x_2, x_1) \in X$ (rotation by 90 degrees). The shape theorem does not offer a more explicit description of D, but representative simulations suggest that it is a union of two-dimensional Euclidean disks (Fig. 2.2b). In Section A.1.11, we discuss how to adapt Bramson and Griffeath's arguments from $\mathbb{Z}^2$ to $\mathbb{Z}^2 \times \mathbb{Z}_w$, and we provide an implicit definition of the set D in terms of the process $(\xi_t^0)_{t \geq 0}$.

To determine the rate of expansion of the mutant clone ξ_t^0, we denote the radius of $D = D(\beta)$ by $c_w(\beta)$, and define it in terms of the projection of D onto the x-axis as

$$\{x \in \mathbb{R} : (x, y, z) \in D\} =: [-c_w(\beta), c_w(\beta)]. \tag{2.5}$$

We furthermore define

$$p_w := \begin{cases} 1, & w = 1, \\ 4/5, & w = 2, \\ 2/3, & w > 2, \end{cases} \tag{2.6}$$

as the probability that a cell giving birth on $\mathbb{Z}^2 \times \mathbb{Z}_w$ replaces a cell occupying the same layer. Note that each cell has four neighboring cells occupying the same layer, independently of w. The difference between the $w = 1$, $w = 2$ and $w > 2$ cases lies in the fact that there are zero, one and two neighboring cells occupying other layers, respectively.

Our main result determines $c_w(\beta)$ as a function of a small selective advantage $\beta > 0$ and tissue thickness $w \geq 1$. Intuitively, it is clear that $c_w(\beta) \to 0$ as $\beta \to 0$. In order to determine $c_w(\beta)$ for small β, we therefore compute its rate of decrease as $\beta \to 0$.

Theorem 2.1. *For $\beta > 0$ and $w \geq 1$, let $c_w(\beta)$ be the radius of the asymptotic shape $D = D(\beta)$ of the biased voter model $(\xi_t^0)_{t\geq 0}$ on $\mathbb{Z}^2 \times \mathbb{Z}_w$, conditioned on the event that it does not die out, as defined by* (2.4) *and* (2.5)*. Then, as $\beta \to 0$,*

$$c_w(\beta) \sim p_w \sqrt{\pi w \beta} / \sqrt{\log(1/\beta)} = a_w / \sqrt{h(\beta)},$$

where p_w is defined as in (2.6)*, $a_w := p_w \sqrt{\pi w}$ and $h(\beta) := (1/\beta) \cdot \log(1/\beta)$.*

Our choice of analyzing type-1 propagation in the direction of the unit vector $e_1 = (1, 0, 0)$ is arbitrary. Our arguments apply to any unit vector of the form $(\cos\theta, \sin\theta, 0)$, so propagation is the same in any direction along the first two coordinates in the $\beta \to 0$ regime.

In Theorem 1 of [46], Durrett, Foo and Leder compute the propagation speed of the biased voter model on $\mathbb{Z}^2$ as $\sqrt{\pi(\beta/4)}/\sqrt{\log(1/\beta)} = (1/2)\sqrt{\pi\beta}/\sqrt{\log(1/\beta)}$ in the $\beta \to 0$ regime. In contrast, our result for the $w = 1$ case on $\mathbb{Z}^2 \times \mathbb{Z}_w$ is $c_1(\beta) \sim \sqrt{\pi\beta}/\sqrt{\log(1/\beta)}$ as $\beta \to 0$, which is larger by a factor of 2. This discrepancy is due to a scaling error in the calculations in [46]. On page 1396 of [46], the authors change the time scale of the process in a way that reduces the fitness advantage of type-1 cells on $\mathbb{Z}^d$ to $\beta/(2d)$. If the result for $\mathbb{Z}^2$ in [46] is read with this in mind, it is consistent with the $w = 1$ case

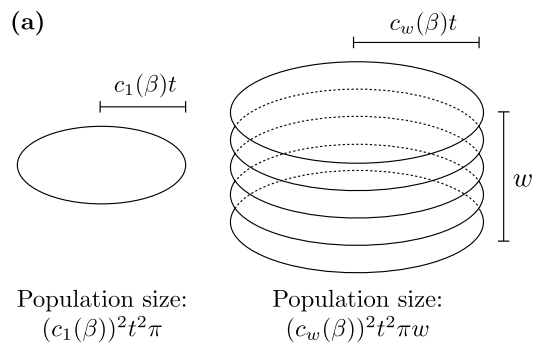

(b)

w	$c_w(\beta)/c_1(\beta)$	$t_1(N)/t_w(N)$
1	1.00	1.00
2	1.13	1.60
3	1.15	2.00
4	1.33	2.67
5	1.49	3.33
6	1.63	4.00
7	1.76	4.67

Figure 2.3: **(a)** For sufficiently large t, the population size of ξ_t^0 at time t is approximately $(c_w(\beta))^2t^2\pi w$, where $c_w(\beta)$ is the radius of the asymptotic shape $D = D(\beta)$. **(b)** The ratio $c_w(\beta)/c_1(\beta)$ shows how the radius of the asymptotic shape of ξ_t^0 changes with increasing w, and the ratio $t_1(N)/t_w(N)$, with $t_w(N)$ defined as in (2.7), shows how quickly the population reaches a given size N as w increases. The latter ratio grows according to $p_w w$.

of our Theorem 2.1.

Note that if ξ_t^0, conditional on $\{\tau_\varnothing^0 = \infty\}$, is a union of two-dimensional disks of radius $c_w(\beta)t$ at time t, the total area across all w layers is $(c_w(\beta))^2t^2\pi w$ (Fig. 2.3a). Thus, given N sufficiently large, the time it takes for the type-1 population to reach size N is approximately

$$t_w(N) = (p_w\pi w)^{-1}h(\beta)^{1/2}N^{1/2}. \tag{2.7}$$

This implies that going from $w = 1$ layer to $w > 1$ layers accelerates population growth by $p_w w$. For example, population growth is twice as fast for $w = 3$ layers as for $w = 1$ layer, and over three times as fast for $w = 5$ layers (Fig. 2.3b).

2.3 Outline of proof of main result

2.3.1 Duality

The biased voter model $(\xi_t^A)_{t\geq 0}$ admits a simple graphical construction, which allows us to define the entire system $\{(\xi_t^A)_{t\geq 0} : A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w\}$ on a common probability space, using a countable family of Poisson processes. For a description of this construction, see e.g. Section 2 of [52], Section 3 of [48] or Appendix A of [46]. By tracing the ancestry

of particles in ξ_t^A backwards in time, the biased voter model gives rise to a system of *branching coalescing random walks* $(\tilde{\zeta}_t^B)_{t\geq 0}$ on $\mathbb{Z}^2\times\mathbb{Z}_w$, which satisfy the duality relation

$$\mathbb{P}(\xi_t^A \cap B \neq \varnothing) = \mathbb{P}(\tilde{\zeta}_t^B \cap A \neq \varnothing), \quad A, B \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w. \tag{2.8}$$

The process $(\tilde{\zeta}_t^B)_{t\geq 0}$ can be described as follows: Each particle performs a simple, symmetric random walk (SSRW) on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate 1, i.e. each particle jumps at rate 1 to a randomly chosen neighboring site. Furthermore, each particle gives birth to a new particle at rate β, with the parent particle staying put and the daughter particle placed at a randomly chosen neighboring site. Any time two particles meet, they coalesce into a single particle. The following elementary properties of ξ_t^A and $\tilde{\zeta}_t^B$ are easily verified:

- *Additivity*: For each $A, B \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$ and $t \geq 0$,

$$\xi_t^{A\cup B} = \xi_t^A \cup \xi_t^B \quad \text{and} \quad \tilde{\zeta}_t^{A\cup B} = \tilde{\zeta}_t^A \cup \tilde{\zeta}_t^B. \tag{2.9}$$

- *Monotonicity:* For $A \subseteq B$ and $t \geq 0$,

$$\xi_t^A \subseteq \xi_t^B \quad \text{and} \quad \tilde{\zeta}_t^A \subseteq \tilde{\zeta}_t^B. \tag{2.10}$$

- *Translation invariance:* For each $A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$ and $x \in \mathbb{Z}^2 \times \mathbb{Z}_w$,

$$(\xi_t^A)_{t\geq 0} \stackrel{d}{=} (x + \xi_t^{A-x})_{t\geq 0} \quad \text{and} \quad (\tilde{\zeta}_t^A)_{t\geq 0} \stackrel{d}{=} (x + \tilde{\zeta}_t^{A-x})_{t\geq 0}. \tag{2.11}$$

- *Symmetry:* For each $A, B \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$:

$$(\xi_t^A)_{t\geq 0} \stackrel{d}{=} (-\xi_t^{-A})_{t\geq 0} \quad \text{and} \quad (\tilde{\zeta}_t^A)_{t\geq 0} \stackrel{d}{=} (-\tilde{\zeta}_t^{-A})_{t\geq 0}. \tag{2.12}$$

Due to the duality relation (2.8), we can use the dual process $\tilde{\zeta}_t$ to study the propagation speed of the biased voter model ξ_t. Direct analysis of $\tilde{\zeta}_t$ is complicated by its coalescing nature. However, it turns out that when β is small, most coalescence events in $\tilde{\zeta}_t$ will be between parent and daughter shortly after the daughter's birth. Before elucidating this property further, we need to establish some fundamental properties of the dual process.

2.3.2 Fundamental properties of dual process

We begin our analysis of the dual process $\tilde{\zeta}_t$ by determining the long-run position of individual particles. In the following lemma, we extend the local central limit theorem (LCLT) for the discrete-time SSRW on $\mathbb{Z}^2$ to the multilayered setting $\mathbb{Z}^2 \times \mathbb{Z}_w$. Since our arguments apply to $\mathbb{Z}^d \times \mathbb{Z}_w$ for any $d \geq 1$, we state and prove the result for the general case. The proof is simple and proceeds as follows. First, we decompose the discrete-time SSRW $(S_n)_{n\geq 0}$ on $\mathbb{Z}^d \times \mathbb{Z}_w$ into walks on $\mathbb{Z}^d$ and $\mathbb{Z}_w$, respectively, and use a large deviations estimate to bound the number of steps in each direction. We then apply the LCLT on $\mathbb{Z}^d$ to the $\mathbb{Z}^d$-walk (Theorem 2.1.3 of [102]), and a convergence theorem for finite Markov chains to the $\mathbb{Z}_w$-walk (Theorem 4.9 of [104]). Note that for $w = 1$, the SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ is equivalent to the SSRW on $\mathbb{Z}^d$, which is periodic. For $w > 1$, on the other hand, the SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ is periodic if and only if w is even.

Lemma 2.1 (LCLT on $\mathbb{Z}^d \times \mathbb{Z}_w$)**.** *Let $(S_n)_{n\geq 0}$ be the discrete-time SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ with $S_0 = 0$. Set $b_1 := 2$, $b_w := 2$ if $w > 1$ is even and $b_w := 1$ if $w > 1$ is odd, and define*

$$p_{w,d} := \begin{cases} 1, & w = 1, \\ 2d/(2d+1), & w = 2, \\ d/(d+1), & w > 2, \end{cases} \tag{2.13}$$

as the probability that $(S_n)_{n\geq 0}$ takes a step in the $\mathbb{Z}^d$-direction, with $p_{w,2} = p_w$ as defined in (2.6). Then, for any $n \geq 1$ and $x \in \mathbb{Z}^d \times \mathbb{Z}_w$ so that $\mathbb{P}(S_n = x) > 0$,

$$n^{d/2}\mathbb{P}(S_n = x) = (b_w/w)\big(d/(2\pi p_{w,d})\big)^{d/2} + o(1). \tag{2.14}$$

Proof. Section A.1.1. □

Let $(Z_t)_{t\geq 0}$ be the continuous-time SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ with jump rate $\alpha > 0$ and $Z_0 = 0$. Since in continuous time, the coordinates move independently, we can decompose the walk into independent components $Z_t = (\hat{Z}_t, Z_t^w)$, where $(\hat{Z}_t)_{t\geq 0}$ is the SSRW on $\mathbb{Z}^d$ with jump rate $\alpha p_{w,d}$, and $(Z_t^w)_{t\geq 0}$ is the SSRW on $\mathbb{Z}_w$ with jump rate $\alpha(1 - p_{w,d})$,

with $p_{w,d}$ defined as in (2.13). For $\hat{x} \in \mathbb{Z}^d$, set

$$p_t(\hat{x}) := \big(d/(2\pi\alpha p_{w,d}t)\big)^{d/2} \exp\big(-d\|\hat{x}\|^2/(2\alpha p_{w,d}t)\big).$$

By Theorem 2.1.3 of [102], there exists $c > 0$ so that for all $\hat{x} \in \mathbb{Z}^d$ and all $t > 0$,

$$|\mathbb{P}(\hat{Z}_t = \hat{x}) - p_t(\hat{x})| \leq c/(\alpha p_{w,d}t)^{(d+2)/2}. \tag{2.15}$$

Expression (2.15) and independence of coordinates imply a continuous-time version of (2.14),

$$\lim_{t\to\infty}(\alpha t)^{d/2}\mathbb{P}(Z_t = x) = (1/w)\big(d/(2\pi p_{w,d})\big)^{d/2}, \quad x \in \mathbb{Z}^d \times \mathbb{Z}_w,$$

since $(Z_t^w)_{t\geq 0}$ converges to the uniform distribution on $\mathbb{Z}_w$.

Whenever a new particle is born into the dual process $\tilde{\zeta}_t$, the parent and daughter perform independent SSRWs Z_t^1 and Z_t^2 on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate 1, started at neighboring sites. We next determine the asymptotic tail of T_0, the time at which Z_t^1 and Z_t^2 first meet, which we can equivalently view as the time of the first visit to the origin of the SSRW $\bar{Z}_t := Z_t^1 - Z_t^2$ with jump rate 2. Due to the recurrence of the SSRW in two dimensions, the two walks Z_t^1 and Z_t^2 are guaranteed to meet in finite time. In the following lemma, we compute the rate of decrease of $\mathbb{P}(T_0 > t)$ as $t \to \infty$. In the proof, we generalize an argument given by Dvoretzky and Erdös for the SSRW on $\mathbb{Z}^2$ in [55]. The only modification necessary is to substitute the LCLT on $\mathbb{Z}^2$ by the LCLT on $\mathbb{Z}^2 \times \mathbb{Z}_w$ (Lemma 2.1).

Lemma 2.2 (Asymptotic tail of T_0)**.** *Let $(Z_t)_{t\geq 0}$ be the SSRW on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate $\alpha > 0$, started at a nearest neighbor of the origin. Set $T_0 := \inf\{t \geq 0 : Z_t = 0\}$ and define*

$$\mu_w := p_w \pi w = \begin{cases} \pi, & w = 1, \\ (4/5)\pi w, & w = 2, \\ (2/3)\pi w, & w > 2, \end{cases}$$

where p_w is the probability given by (2.6)*. Then*

$$\mathbb{P}(T_0 > t) \sim \mu_w / \log t, \quad t \to \infty.$$

Proof. Section A.1.2. □

Recall that each particle in $\tilde{\zeta}_t$ gives birth to a new particle at rate β, so the mean time between births along a particular lineage is $1/\beta$. Set

$$\tau(\beta) := (1/\beta)\big(1/\sqrt{\log(1/\beta)}\big). \tag{2.16}$$

By Lemma 2.2, a new particle avoids coalescence with its parent particle during the first $\tau(\beta)$ time units of its existence with probability

$$\mathbb{P}(T_0 > \tau(\beta)) \sim \mu_w / \log(1/\beta), \quad \beta \to 0.$$

Thus, most new particles coalesce with their parents before time $\tau(\beta)$, and since $\tau(\beta) = o(1/\beta)$, they are unlikely to produce their own offspring before coalescing. Ignoring such particles should simplify the process considerably without affecting its long-run growth. This is the basic idea of a pruning procedure suggested by Durrett and Zähle [54], which we use to set up an approximation scheme to prove our main result (see Sections 2.3.4 and 2.3.5 below). The specific form of $\tau(\beta)$ in (2.16) ensures not only that inconsequential particles are ignored, but also that new particles that do avoid coalescence up until time $\tau(\beta)$ are neither too far away from nor too close to their parent particles at that time.

It turns out that a separate approximation scheme is required for establishing an upper bound and a lower bound on the propagation speed of $\tilde{\zeta}_t$. Before describing the scheme in detail, we discuss it at a high level and provide some intuition for our main result.

2.3.3 Overview of approximation scheme and intuition for main result

Since particles in $\tilde{\zeta}_t$ branch at rate β, branching events become less and less frequent as $\beta \to 0$, and most events produce particles that coalesce with their parents shortly

after birth. If we "reject" branching events where new particles are lost to coalescence quickly, the rate of "accepted" events along a particular lineage is of order

$$\beta \cdot \big(\mu_w / \log(1/\beta)\big) = \mu_w / h(\beta), \quad \beta \to 0,$$

where $h(\beta) := (1/\beta) \cdot \log(1/\beta)$. Assume for the moment that particles in $\tilde{\zeta}_t$ are sufficiently spread out that we can ignore other coalescence events. We then obtain a branching random walk with branching rate $\mu_w/h(\beta)$ and average number of particles $\exp\big((\mu_w/h(\beta))t\big)$ alive at time t. If we project onto the x-axis, each particle performs a SSRW on $\mathbb{Z}$ with jump rate $p_w/2$, where p_w is defined as in (2.6). For large t, its position has approximate distribution

$$(\sqrt{p_w \pi t})^{-1} \cdot \exp\big(-x^2/(p_w t)\big), \quad x \in \mathbb{Z},$$

and the particle intensity (average number of particles) at $x \in \mathbb{Z}$ is approximately

$$(\sqrt{p_w \pi t})^{-1} \cdot \exp\big((\mu_w/h(\beta))t - x^2/(p_w t)\big).$$

If we set $|x| = ct$ for $c > 0$, this quantity is nonzero in the $t \to \infty$ limit as long as $c^2 < p_w \mu_w / h(\beta)$ i.e. $c < a_w/\sqrt{h(\beta)}$, since $a_w = p_w \sqrt{\pi w} = \sqrt{p_w \mu_w}$. This suggests a long-run expansion rate of $a_w/\sqrt{h(\beta)}$ per unit time, which is our main result.

To make this argument rigorous, we need to show that for small β, the dual process $\tilde{\zeta}_t$ sufficiently resembles a branching random walk (BRW) with branching rate $\mu_w/h(\beta)$. Since the time between accepted branching events is of order $h(\beta) \to \infty$ as $\beta \to 0$, and in this time, fluctuations in the movement of individual particles are of order $\sqrt{h(\beta)}$, it makes sense to speed up time by $h(\beta)$ and reduce space by $\sqrt{h(\beta)}$. We therefore introduce the scaled dual process

$$\tilde{\zeta}_t^\beta := h(\beta)^{-1/2} \cdot \tilde{\zeta}_{h(\beta)t}, \tag{2.17}$$

and our goal is to show that for small β, this process sufficiently resembles a BRW with branching rate μ_w. Recall that by the definition (2.2) of scalar multiplication on $\mathbb{R}^2 \times \mathbb{Z}_w$, the spatial scaling by $h(\beta)^{-1/2}$ only affects the first two coordinates.

ξ_t	Biased voter model	Section 2.2	Type-0 particles divide at rate 1, type-1 at rate $1+\beta$. A neighbor selected uniformly at random is replaced.
$\tilde{\zeta}_t$	Dual process	Section 2.3.1	Branching coalescing random walk (BCRW). Particles jump at rate 1, branch at rate β.
$\bar{\phi}_t$	Unaltered BRW	Section 2.3.4	Branching random walk (BRW) obtained by ignoring all coalescence events in dual process $\tilde{\zeta}_t$.
$\mathring{\phi}_t$	Pruned BRW, upper bound	Section 2.3.4	BRW obtained by ignoring new particles in $\bar{\phi}_t$ that coincide quickly with their parent particles.
ϕ_t	Simple BRW, upper bound	Section 2.3.4	BRW obtained by modifying particle paths in $\mathring{\phi}_t$ to uncondition movement at branching events.
$\hat{\zeta}_t$	Pruned dual, lower bound	Section 2.3.5	BCRW obtained by ignoring new particles that coalesce quickly with *any* particle in $\hat{\zeta}_t$.
$\mathring{\psi}_t$	Pruned BRW, lower bound	Section 2.3.5	BRW obtained by ignoring new particles in $\bar{\phi}_t$ that coincide quickly with their parent particles.
ψ_t	Simple BRW, lower bound	Section 2.3.5	BRW obtained by modifying particle paths in $\mathring{\psi}_t$ to uncondition movement at branching events.

Table 2.1: List of the particle processes used in the proofs of Lemmas 2.3 to 2.10.

In the upper bound proof, the main work resides in identifying which branching events to accept, and in analyzing parent-daughter interactions under the accepted events. In the lower bound proof, we can only approximate $\tilde{\zeta}_t^\beta$ with a branching random walk on finite time intervals. We therefore discretize time and space and apply a percolation argument to obtain the long-run propagation speed of $\tilde{\zeta}_t^\beta$.

2.3.4 Upper bound argument

To prove an upper bound, we couple the dual process $\tilde{\zeta}_t$ with a pruned branching random walk $\mathring{\phi}_t$, which we in turn couple with a simpler BRW ϕ_t. We then analyze the propagation speed of ϕ_t to obtain an upper bound on the propagation speed of $\tilde{\zeta}_t$. For reference, we list the processes used in the proof of our main result along with a short description in Table 2.1.

Definition of pruned BRW $\mathring{\phi}_t$

Consider a branching random walk $\bar{\phi}_t$ obtained by ignoring all coalescence events in the dual process $\tilde{\zeta}_t$. In other words, particles in $\bar{\phi}_t$ jump at rate 1 and branch at rate

β, and whenever two particles meet, both are retained. One particle follows the path of the coalesced particle in the dual process, and the other performs a new SSRW on $\mathbb{Z}^2 \times \mathbb{Z}_w$ independently of all other particles. By construction, each particle path in $\tilde{\zeta}_t$ also appears in $\bar{\phi}_t$, but $\bar{\phi}_t$ contains additional paths. Since $\bar{\phi}_t$ allows multiple particles to occupy the same site, it should be viewed as a sequence of sites in $\mathbb{Z}^2 \times \mathbb{Z}_w$, as opposed to a subset of $\mathbb{Z}^2 \times \mathbb{Z}_w$. We will call $\bar{\phi}_t$ the *unaltered* BRW to distinguish it from the *pruned* BRW $\mathring{\phi}_t$, which we define now.

For a given branching event in $\bar{\phi}_t$, let T_0 be the time at which the new particle first coincides with its parent, and let S be the time at which the new particle first produces its own offspring. Recall that S is exponentially distributed with mean $1/\beta$, and it is independent of T_0. We categorize the branching events in $\bar{\phi}_t$ as follows:

- **Type-0:** $T_0 \leq \min\{S, \tau(\beta)\}$: The new particle quickly coincides with its parent.
- **Type-1:** $S \leq \min\{T_0, \tau(\beta)\}$: The new particle quickly produces its own offspring.
- **Type-2:** $\tau(\beta) \leq \min\{S, T_0\}$: The new particle neither coincides with its parent nor produces its own offspring before time $\tau(\beta)$.

We refer to $[0, T_0]$, $[0, S]$ and $[0, \tau(\beta)]$, respectively, as the *decision period* for each type of event. The pruned BRW $\mathring{\phi}_t$ is defined as follows (Fig. 2.4):

- A new particle born through a type-0 branching event in $\bar{\phi}_t$ is ignored in $\mathring{\phi}_t$.
- A new particle born through a type-1 event is introduced to $\mathring{\phi}_t$ at time S after birth in $\bar{\phi}_t$, at the location it then occupies in $\bar{\phi}_t$. Its offspring is viewed as a new branching event in $\bar{\phi}_t$ and is evaluated according to the same rules as outlined here.
- A new particle born through a type-2 event is introduced to $\mathring{\phi}_t$ at time $\tau(\beta)$ after birth in $\bar{\phi}_t$, at the location it then occupies.

Once a new particle is introduced to $\mathring{\phi}_t$, it follows the same path as in $\bar{\phi}_t$. Let $\mathring{\phi}_t^{(k)}$ for $k = 0, 1, 2$ be the subprocess of $\bar{\phi}_t$ containing offsprings of particles in $\mathring{\phi}_t$ that have just been born through a type-k branching event and whose decision period has not yet

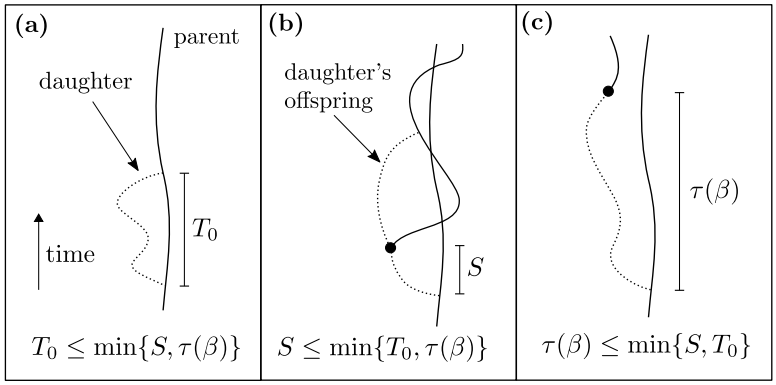

Figure 2.4: Categorization of branching events in the unaltered BRW $\bar{\phi}_t$ and definition of the pruned process $\mathring{\phi}_t$. **(a)** If a new particle in $\bar{\phi}_t$ coincides quickly with its parent, it is not introduced to $\mathring{\phi}_t$. **(b)** If the new particle produces its own offspring quickly, it gets introduced to $\mathring{\phi}_t$ at the time and location at which it gives birth. **(c)** If the new particle neither coincides with its parent nor has its own offspring too quickly, it gets introduced to $\mathring{\phi}_t$ at time $\tau(\beta)$ after birth at the location it then occupies.

passed. Then

$$\tilde{\zeta}_t \subseteq \mathring{\phi}_t \cup \mathring{\phi}_t^{(0)} \cup \mathring{\phi}_t^{(1)} \cup \mathring{\phi}_t^{(2)}, \tag{2.18}$$

i.e. $\mathring{\phi}_t$ upper bounds the dual process $\tilde{\zeta}_t$ if we add newborn particles whose fate has not been decided yet. Expression (2.18) allows us to relate the propagation speed of $\tilde{\zeta}_t$ to that of $\mathring{\phi}_t$. Before doing so, we need more information on the branching dynamics of $\bar{\phi}_t$ and $\mathring{\phi}_t$.

Branching in $\bar{\phi}_t$ and $\mathring{\phi}_t$

In the following lemma, we show that in the $\beta \to 0$ regime, almost all branching events of the unaltered BRW $\bar{\phi}_t$ are type-0. In other words, only a small proportion of branching events is accepted to produce the pruned BRW $\mathring{\phi}_t$. We also show that type-2 branching events are much more frequent than type-1 events, meaning that most particles introduced to $\mathring{\phi}_t$ neither coincide with their parent nor produce their own offspring by time $\tau(\beta)$. We finally produce moment bounds on the distance traveled by a new particle during its decision period, as well the separation between parent and daughter throughout the decision period.

Lemma 2.3. *Let $(Z_t^1)_{t\geq 0}$ and $(Z_t^2)_{t\geq 0}$ be independent SSRWs on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate 1, started at 0 and a nearest neighbor of 0. Set $T_0 := \inf\{t \geq 0 : Z_t^1 = Z_t^2\}$, and let S be an exponential random variable with mean $1/\beta$, independent of $(Z_t^1)_{t\geq 0}$ and $(Z_t^2)_{t\geq 0}$. Then*

(1) $\alpha_0(\beta) := \mathbb{P}(T_0 \leq \min\{S, \tau(\beta)\}) \to 1$ *as* $\beta \to 0$.

(2) $\alpha_1(\beta) := \mathbb{P}(S \leq \min\{T_0, \tau(\beta)\}) = \Theta\big(1/(\log(1/\beta))^{3/2}\big)$ *as* $\beta \to 0$,

(3) $\alpha_2(\beta) := \mathbb{P}(\tau(\beta) \leq \min\{S, T_0\}) \sim \mu_w/\log(1/\beta)$ *as* $\beta \to 0$.

Furthermore, there exists $C > 0$ so that for sufficiently small β,

(4) $\mathbb{E}\big[\sup_{t\leq\tau(\beta)} ||Z_t^1||^j \,\big|\, T_0 \leq \min\{S, \tau(\beta)\}\big] \leq Cj!\tau(\beta)^{j/2}, \quad j \geq 1,$

(5) $\mathbb{E}\big[\sup_{t\leq\tau(\beta)} ||Z_t^1||^j \,\big|\, S \leq \min\{T_0, \tau(\beta)\}\big] \leq Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2}, \quad j \geq 1,$

(6) $\mathbb{E}\big[\sup_{t\leq\tau(\beta)} ||Z_t^1||^j \,\big|\, \tau(\beta) \leq \min\{S, T_0\}\big] \leq Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2}, \quad j \geq 1.$

Each of (4)-(6) continues to hold if Z_t^1 is replaced by Z_t^2 or $\bar{Z}_t = Z_t^1 - Z_t^2$.

Proof. Section A.1.3. □

From dual $\tilde{\zeta}_t$ to pruned BRW $\mathring{\phi}_t$

Using (2.18) and Lemma 2.3, we can obtain the following relationship between the propagation speed of the dual process $\tilde{\zeta}_t$ and the pruned BRW $\mathring{\phi}_t$ (Lemma 2.4). By (2.18), $\mathring{\phi}_t$ upper bounds $\tilde{\zeta}_t$ if we add newborn particles whose fate has not been decided yet. By (4)-(6) in Lemma 2.3, these newborn particles will not be too far from their parent particles in $\mathring{\phi}_t$, so adding them should not materially affect the propagation speed of $\mathring{\phi}_t$. As motivated by the discussion in Section 2.3.3, we perform our analysis using the scaled processes $\tilde{\zeta}_t^{\beta,0} := h(\beta)^{-1/2}\tilde{\zeta}^0_{h(\beta)t}$ and $\mathring{\phi}_t^{\beta,0} := h(\beta)^{-1/2}\mathring{\phi}^0_{h(\beta)t}$, where the 0 means that the processes are started with a single particle at the origin.

Lemma 2.4. *Fix $a > 0$ and $0 < \rho < 1$, and define $A_r := [r, \infty) \times \mathbb{R} \times \mathbb{Z}_w$. For each $\delta > 0$, there exist $M > 0$ and $\beta_0 > 0$ so that*

$$\mathbb{P}\big(\tilde{\zeta}_t^{\beta,0} \cap A_{b+at} \neq \varnothing\big) \leq 4\,\mathbb{P}\big(\mathring{\phi}_t^{\beta,0} \cap A_{b+\rho at} \neq \varnothing\big) + Me^{-\delta t}, \quad \beta \leq \beta_0,\ t > 0,\ b \in \mathbb{R}.$$

Proof. Section A.1.4. □

Definition of simple BRW ϕ_t

At each accepted branching event of the pruned BRW $\mathring{\phi}_t$ (type-1 or type-2), the new particle is introduced with a time delay, and the location at which it is introduced is conditioned on it not coinciding too quickly with its parent. The parent's path during the decision period is likewise influenced by this conditioning. We next couple $\mathring{\phi}_t$ with a simpler BRW ϕ_t where we modify particle paths as follows:

- At each accepted branching event of $\mathring{\phi}_t$, the paths of parent and daughter during the decision period are replaced by two independent SSRWs started at the parent's location. From the end of the decision period onward, the two new walks make the same transitions as parent and daughter make in $\mathring{\phi}_t$. An illustration of this procedure is shown in Figure 2.5, and a more formal mathematical description is given in the proof of Lemma 2.5 (Section A.1.5).

With these modifications, both parent and daughter follow independent, unconditioned paths at each branching event of ϕ_t. We must make further modifications, however, since the path followed by the parent at a type-0 branching event in $\mathring{\phi}_t$, in which case the daughter is not introduced to $\mathring{\phi}_t$, is conditioned on coinciding quickly with the daughter. This conditioning will affect particle paths in ϕ_t if not addressed. We therefore make the following modifications:

- At a type-0 branching event in $\mathring{\phi}_t$, the parent follows a path $(Z_t^1)_{t\geq 0}$ conditioned on $\{T_0 \leq \min\{S, \tau(\beta)\}\}$ during $[0, \tau(\beta)]$, in the notation of Lemma 2.3.
  - With probability $\alpha_0(\beta)$, with $\alpha_0(\beta)$ defined as in Lemma 2.3, we make no modification to the parent's path.
  - With probability $\alpha_1(\beta)$, we replace the parent's path on $[0, \tau(\beta)]$ with a path $(Z_t^1)_{t\geq 0}$ conditioned on $\{S \leq \min\{T_0, \tau(\beta)\}\}$. From time $\tau(\beta)$ onward, the new path makes the same transitions as the parent in $\mathring{\phi}_t$.
  - With probability $\alpha_2(\beta)$, we replace the parent's path on $[0, \tau(\beta)]$ with a path $(Z_t^1)_{t\geq 0}$ conditioned on $\{\tau(\beta) \leq \min\{S, T_0\}\}$. From time $\tau(\beta)$ onward, the new path makes the same transitions as the parent in $\mathring{\phi}_t$.

With these modifications, we remove any effect of daughter particles not introduced to

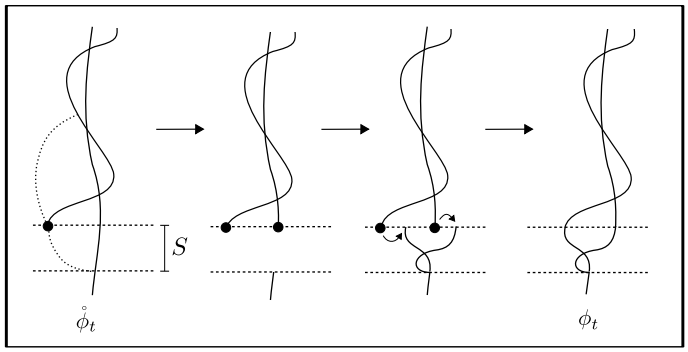

Figure 2.5: The simple BRW ϕ_t is obtained from the pruned BRW $\mathring{\phi}_t$ by making modifications to the paths followed by parent and daughter at each branching event of $\mathring{\phi}_t$. In the figure, we show the procedure for a type-1 branching event. Two new independent SSRWs are started at the parent's location at the beginning of the decision period, which replace the paths of parent and daughter during the decision period. From the end of the decision period onward, the paths followed by parent and daughter in $\mathring{\phi}_t$ are shifted to meet the two new walks in ϕ_t.

$\mathring{\phi}_t$ on the paths followed by their parent particles in ϕ_t. By the above construction, ϕ_t has the following three simplifying properties:

- Particles follow independent, unconditioned SSRWs at all times.
- Time between branching events is exponentially distributed.
- New particles are born to their parents' locations.

From pruned BRW $\mathring{\phi}_t$ to simple BRW ϕ_t

Working with the scaled versions $\mathring{\phi}_t^\beta := h(\beta)^{-1/2}\mathring{\phi}_{h(\beta)t}$ and $\phi_t^\beta := h(\beta)^{-1/2}\phi_{h(\beta)t}$, we show in the following lemma that the path followed by an arbitrary particle in ϕ_t^β is never too far away from the corresponding path in $\mathring{\phi}_t^\beta$ when β is small. Since ϕ_t^β is defined by perturbing particle paths in $\mathring{\phi}_t^\beta$ at branching events, we need to show that the accumulated perturbation up until time t is not too large. To do so, we first establish an upper bound on the number of perturbations by time t, and we then use the moment bounds established in (4)-(6) of Lemma 2.3 to bound the accumulated perturbation.

Lemma 2.5. *For a particle chosen uniformly at random from $\phi_t^{\beta,0}$, let $(Y_s^\beta)_{s\le t}$ be the path followed by this particle and its ancestors, and let $(\mathring{Y}_s^\beta)_{s\le t}$ be the corresponding path in $\mathring{\phi}_t^{\beta,0}$. Then, for any $r>0$ and $\delta>0$, there exist $M>0$ and $\beta_0>0$ so that*

$$\mathbb{P}\big(\sup_{s\le t}||Y_s^\beta-\mathring{Y}_s^\beta||>rt\big)\le Me^{-\delta t},\quad \beta\le\beta_0,\ t>0.$$

Proof. Section A.1.5. □

Using Lemma 2.5, we can obtain the following relationship between the propagation speed of $\mathring{\phi}_t^\beta$ and ϕ_t^β (Lemma 2.6). In the proof, we use the fact that the mean number of particles alive at time t in ϕ_t^β is $\exp((\mu_w+o(1))t)$, and that the error in approximating $\mathring{\phi}_t^\beta$ with ϕ_t^β on a particle-by-particle basis is sufficiently small by Lemma 2.5 to ensure a small total error.

Lemma 2.6. *Fix $a>0$ and $0<\rho<1$, and define $A_r:=[r,\infty)\times\mathbb{R}\times\mathbb{Z}_w$. For each $\delta>0$, there exist $M>0$ and $\beta_0>0$ so that*

$$\mathbb{P}\big(\mathring{\phi}_t^{\beta,0}\cap A_{b+at}\neq\varnothing\big)\le\mathbb{P}\big(\phi_t^{\beta,0}\cap A_{b+\rho at}\neq\varnothing\big)+Me^{-\delta t},\quad \beta\le\beta_0,\ t>0,\ b\in\mathbb{R}.$$

Proof. Section A.1.6. □

Upper bound result for ξ_t

With the above ingredients, we can establish the following upper bound result on the propagation speed of the biased voter model $(\xi_t^0)_{t\ge0}$ on $\mathbb{Z}^2\times\mathbb{Z}_w$ conditioned on nonextinction (Lemma 2.7). The proof is split into three key steps. First, we remove the conditioning on nonextinction by waiting until ξ_t^0 has covered a sufficiently large box. We then introduce duality using (2.8) and use Lemmas 2.4 and 2.6 to pass from the dual process $\tilde{\zeta}_t$ to the simple BRW ϕ_t. We finally obtain the desired result by analyzing the tail of ϕ_t.

Lemma 2.7. *Define $\tau_\varnothing^A:=\min\{t\ge0:\xi_t^A=\varnothing\}$ for $A\subseteq\mathbb{Z}^2\times\mathbb{Z}_w$. For each $\kappa>1$, there exists a family of random variables $(S_\beta)_{\beta>0}$, with $\mathbb{P}(S_\beta<\infty|\tau_\varnothing^0=\infty)=1$ for*

each $\beta > 0$, so that

$$\lim_{\beta\to 0} \liminf_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1 \notin \xi^0_{S_\beta + h(\beta)t} \,\big|\, \tau^0_\varnothing = \infty\big) = 1,$$

where $a_w := p_w\sqrt{\pi w}$.

Proof. Section A.1.7. □

2.3.5 Lower bound argument

To prove a lower bound, we couple the dual process $\tilde{\zeta}_t$ with a pruned dual process $\hat{\zeta}_t$ and a pruned BRW $\mathring{\psi}_t$, which we in turn couple with a simpler BRW ψ_t. Unfortunately, as mentioned previously, the scaled versions of these processes only behave similarly on finite time intervals. We therefore discretize time and space and apply a percolation argument to determine the propagation speed of ψ_t in the discretized spacetime.

Definition of pruned dual $\hat{\zeta}_t$

To define the pruned dual process $\hat{\zeta}_t$, we start with the dual process $\tilde{\zeta}_t$. For a given branching event in $\hat{\zeta}_t$, let T_0' be the time at which the new particle coalesces with *any other particle in* $\hat{\zeta}_t$, and let S be the time at which the new particle first produces its own offspring. We categorize the branching events in $\hat{\zeta}_t$ as follows:

- **Type-0:** $T_0' \leq \min\{S, \tau(\beta)\}$: The new particle quickly coalesces with another particle.
- **Type-1:** $S \leq \min\{T_0', \tau(\beta)\}$: The new particle quickly produces its own offspring.
- **Type-2:** $\tau(\beta) \leq \min\{T_0', S\}$: The new particle neither coalesces with another particle nor produces its own offspring before time $\tau(\beta)$.

The pruned process $\hat{\zeta}_t$ is obtained by only accepting type-2 branching events. In case of acceptance, the new particle is introduced to $\hat{\zeta}_t$ at time $\tau(\beta)$ after birth in $\tilde{\zeta}_t$, at the location it then occupies. From the time of introduction to $\hat{\zeta}_t$, the new particle follows the same path as in the dual process $\tilde{\zeta}_t$, and it coalesces with other particles in $\hat{\zeta}_t$. By

construction, the pruned process $\hat{\zeta}_t$ lower bounds $\tilde{\zeta}_t$ in the following sense:

$$\tilde{\zeta}_t \supseteq \hat{\zeta}_t, \quad t \geq 0. \tag{2.19}$$

From pruned dual $\hat{\zeta}_t$ to pruned BRW $\mathring{\psi}_t$

As in the proof of the upper bound, we consider an unaltered BRW $\bar{\phi}_t$ obtained from ignoring all coalescence events in the dual process $\tilde{\zeta}_t$. We classify its branching events into type-0, type-1 and type-2 as in Section 2.3.4, using T_0, the time at which a new particle first coincides with its *parent particle*. We then define the pruned BRW $\mathring{\psi}_t$ by only accepting type-2 branching events. Note that the difference between $\hat{\zeta}_t$ and $\mathring{\psi}_t$ is that the latter process accepts a new particle that coincides with a particle other than its parent before time $\tau(\beta)$, and it ignores any coalescence (whether with the parent or another particle) that occurs after the new particle has been introduced.

The pruned BRW $\mathring{\psi}_t$ may not appear useful for determining a lower bound on the propagation of $\hat{\zeta}_t$, as its growth is not checked to the same degree by coalescence. However, working with the scaled versions $\hat{\zeta}_t^\beta := h(\beta)^{-1/2}\hat{\zeta}_{h(\beta)t}$ and $\mathring{\psi}_t^\beta := h(\beta)^{-1/2}\mathring{\psi}_{h(\beta)t}$, we can show that if $\hat{\zeta}_t^\beta$ is started with sufficient spacing between the initial particles, then as $\beta \to 0$, the only coalescence that occurs during a finite time interval will be between parent and daughter during a decision period. In other words, $\hat{\zeta}_t^\beta$ behaves like $\mathring{\psi}_t^\beta$ on finite time intervals (in the scaled spacetime). We obtain the following lemma, whose proof follows from an induction argument similar to the one given on pages 1758-1759 of Durrett and Zähle [54].

Lemma 2.8. *Set $d(\beta) := \beta^{-1/2}(\log(1/\beta))^{-1}$. Let $\mathcal{A} = \mathcal{A}(\beta)$ denote the collection of finite subsets of $\mathbb{Z}^2 \times \mathbb{Z}_w$ in which points are pairwise separated by at least $d(\beta)$. Set $A^\beta := h(\beta)^{-1/2}A$ for $A \in \mathcal{A}$ and $\mathcal{A}^\beta := h(\beta)^{-1/2}\mathcal{A}$. Then, for any $K > 0$ and $T > 0$,*

$$\sup_{A\in\mathcal{A},\, |A|\leq K} \mathbb{P}\big(\{(\hat{\zeta}_t^{\beta,A^\beta})_{t\leq T} \neq (\mathring{\psi}_t^{\beta,A^\beta})_{t\leq T}\} \cup \{\mathring{\psi}_T^{\beta,A^\beta} \notin \mathcal{A}^\beta\}\big) \to 0, \quad \beta \to 0.$$

Proof. Section A.1.8. □

Propagation of pruned dual $\hat{\zeta}_t$

In the following lemma (Lemma 2.9), we show that for any $2/3 < \theta < 1$, we can find L and K so that if $\hat{\zeta}_t^\beta$ is started with K particles in a box of diameter $\theta a_w L$, then L time units later, there will be at least K particles in an adjacent box of diameter $\theta a_w L$, with high probability. This suggests that during $[0, L]$, the propagation speed of $\hat{\zeta}_t^\beta$ is at least θa_w, which translates into the desired lower bound of $\theta a_w/\sqrt{h(\beta)}$ in the unscaled spacetime. The specific form of the result (2.20) of Lemma 2.9 enables us to define a lower-bounding percolation process (Fig. 2.6a) using a comparison theorem from Section 4 of [52], in which we discretize time into blocks of length L and space into boxes of diameter $\theta a_w L$. A similar construction is carried out in Durrett and Zähle [54], except we must shorten their time blocks from length L^2 to length L to obtain a tight lower bound on the propagation speed of $\hat{\zeta}_t^\beta$. To prove Lemma 2.9, we use Lemma 2.8 to approximate $\hat{\zeta}_t^\beta$ with $\mathring{\psi}_t^\beta$, and we then approximate $\mathring{\psi}_t^\beta$ with a simpler BRW ψ_t^β as in the proof of the upper bound. We finally analyze ψ_t^β to obtain the result.

Lemma 2.9. *For $0 < \theta < 1$ and $L > 0$, define (Fig. 2.6b)*

$$I_0^\theta := [-(1/2)\theta a_w L, (1/2)\theta a_w L]^2 \times \mathbb{Z}_w,$$
$$I_k^\theta := I_0^\theta + k \cdot \theta a_w L e_1,\ k \in \mathbb{Z},$$

and

$$\mathcal{A}^{\beta,\theta,K,k} := \{A^\beta \in \mathcal{A}^\beta : |A^\beta \cap I_k^\theta| \geq K\},$$

with $\mathcal{A}^\beta$ defined as in Lemma 2.8. Let $(\hat{\zeta}_t^{\beta,A^\beta,\theta})_{t\geq 0}$ denote a pruning of $(\hat{\zeta}_t^{\beta,A^\beta})_{t\geq 0}$ with particles killed as soon as they exit the box I_Δ^θ with $I_\Delta^\theta := [-2\theta a_w L, 2\theta a_w L]^2 \times \mathbb{Z}_w$. Then, for any $2/3 < \theta < 1$ and $\varepsilon > 0$, there exist $L = L(\theta) > 0$, $K = K(\theta, \varepsilon) > 0$ and $\beta_0 = \beta_0(\theta, \varepsilon) > 0$ such that for any $A^\beta \in \mathcal{A}^{\beta,\theta,K,0}$ with $|A^\beta| = K$, and any $\beta \leq \beta_0$,

$$\mathbb{P}(\hat{\zeta}_L^{\beta,A^\beta,\theta} \in \mathcal{A}^{\beta,\theta,K,k}) \geq 1 - \varepsilon, \quad k \in \{-1, 1\}. \tag{2.20}$$

Proof. Section A.1.9. □

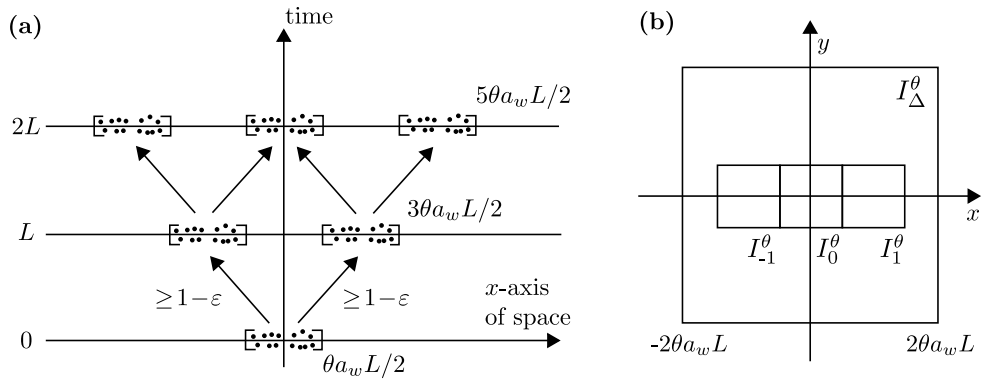

Figure 2.6: **(a)** Graphical depiction of the percolation construction embedded in Lemma 2.9. **(b)** Illustration of the two-dimensional projection of the sets I^θ_k for $k \in \{-1, 0, 1\}$ and the set I^θ_Δ defined in Lemma 2.9.

Lower bound result for ξ_t

With the above ingredients, we can establish the following lower bound on the propagation speed of the biased voter model $(\xi^0_t)_{t\geq 0}$ on $\mathbb{Z}^2 \times \mathbb{Z}_w$ conditioned on nonextinction (Lemma 2.10). The proof is split into three key steps. In the first two steps, we remove the conditioning on nonextinction and introduce duality. In the final step, we use Lemma 2.9 to define a lower-bounding percolation process, which we then analyze using results from [47] and [107].

Lemma 2.10. *Define $\tau^A_\varnothing := \min\{t \geq 0 : \xi^A_t = \varnothing\}$ for $A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$. For each $2/3 < \rho < 1$, there exists a constant $L > 0$ and a family of random variables $(S_\beta)_{\beta>0}$, with $\mathbb{P}(S_\beta < \infty | \tau^0_\varnothing = \infty) = 1$ for each $\beta > 0$, so that*

$$\lim_{\beta\to 0} \liminf_{n\to\infty} \mathbb{P}\big(\xi^0_{S_\beta + 2nLh(\beta)} \cap \big[2n\rho a_w L h(\beta)^{1/2}, \infty\big) e_1 \neq \varnothing \,\big|\, \tau^0_\varnothing = \infty\big) = 1,$$

where $a_w := p_w\sqrt{\pi w}$.

Proof. Section A.1.10. □

Note that due to the discretization involved in the proof, Lemma 2.10 only addresses the propagation of $(\xi_t)_{t\geq 0}$ along a subsequence of timepoints. This turns out to be sufficient when combined with the Bramson-Griffeath shape theorem (2.4), as we show next.

2.4 Proof of main result

In this section, we complete the proof of our main theorem (Theorem 2.1) by showing how it follows from the lower and upper bound results of Lemmas 2.7 and 2.10, together with the shape theorem (2.4).

Proof of Theorem 2.1. Fix $0 < \delta < 1$ and let $(\beta_i)_{i\geq 1}$ be a sequence of real numbers converging to 0. Apply Lemma 2.7 with $\kappa := 1+\delta$ to obtain finite random variables $(S_{\beta_i})_i$ and an integer i_0 so that for $i \geq i_0$,

$$\liminf_{n\to\infty} \mathbb{P}\big(\lceil (1+\delta) a_w h(\beta_i)^{1/2} n\rceil e_1 \notin \xi^0_{S_{\beta_i}+h(\beta_i)n} \,\big|\, \tau^0_\varnothing = \infty\big) \geq 1/2.$$

It follows that for $i \geq i_0$,

$$\mathbb{P}\big(\lceil (1+\delta) a_w h(\beta_i)^{1/2} n\rceil e_1 \notin \xi^0_{S_{\beta_i}+h(\beta_i)n} \text{ for infinitely many } n \,\big|\, \tau^0_\varnothing = \infty\big) \geq 1/2. \quad (2.21)$$

Fix $i \geq i_0$. By the shape theorem (2.4),

$$\mathbb{P}\big(\exists t_* < \infty : [\![-\lfloor (1-\delta) c_w(\beta_i) t\rfloor, \lfloor (1-\delta) c_w(\beta_i) t\rfloor]\!] e_1 \subseteq \xi^0_t,\ t \geq t_* \,\big|\, \tau^0_\varnothing = \infty\big) = 1, \quad (2.22)$$

where $c_w(\beta_i)$ is the radius of the asymptotic shape $D = D(\beta_i)$ as defined in (2.5), and $[\![m,n]\!] = \{m, m+1, \ldots, n\}$ for integers $m < n$. Assume now, by way of contradiction, that

$$c_w(\beta_i) \geq \big((1+2\delta)/(1-2\delta)\big) \cdot a_w h(\beta_i)^{-1/2}.$$

For sufficiently large n (which depends on the outcome ω),

$$\lceil (1+\delta) a_w h(\beta_i)^{1/2} n\rceil / (S_{\beta_i} + h(\beta_i) n) \leq (1+2\delta) a_w h(\beta_i)^{-1/2},$$
$$\lfloor (1-\delta) c_w(\beta_i)(S_{\beta_i} + h(\beta_i) n)\rfloor / (S_{\beta_i} + h(\beta_i) n) \geq (1-2\delta) c_w(\beta_i),$$

which implies

$$\lceil (1+\delta) a_w h(\beta_i)^{1/2} n\rceil \leq \lfloor (1-\delta) c_w(\beta_i)(S_{\beta_i} + h(\beta_i) n)\rfloor.$$

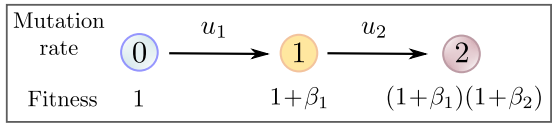

Figure 2.7: In the two-step mutational model of cancer (Section 2.5), the mutation rates are u_1 and u_2 for type-0 and type-1 cells, respectively. Type-1 cells have fitness advantage β_1 over type-0 cells, and type-2 cells have fitness advantage β_2 over type-1 cells.

Conditional on $\{\tau^0_\varnothing = \infty\}$, we must then have $\lceil (1+\delta) a_w h(\beta_i)^{1/2} n \rceil e_1 \in \xi^0_{S_{\beta_i}+h(\beta_i)n}$ for all but finitely many n with probability 1 by (2.22), which contradicts (2.21). We can therefore conclude that

$$c_w(\beta_i) \leq \big((1+2\delta)/(1-2\delta)\big) \cdot a_w h(\beta_i)^{-1/2}, \quad i \geq i_0.$$

Sending $i \to \infty$, and noting that the subsequence $\{\beta_i\}_{i\geq 1}$ is arbitrary, we obtain

$$\limsup_{\beta\to 0} c_w(\beta) \big/ \big(a_w h(\beta)^{-1/2}\big) \leq (1+2\delta)/(1-2\delta).$$

Sending $\delta \to 0$ then yields $\limsup_{\beta\to 0} c_w(\beta)/(a_w h(\beta)^{-1/2}) \leq 1$. Applying a similar argument to the lower bound result of Lemma 2.10 will show that

$$\liminf_{\beta\to 0} c_w(\beta) \big/ \big(a_w h(\beta)^{-1/2}\big) \geq 1,$$

and we can conclude that $c_w(\beta) \sim a_w h(\beta)^{-1/2}$ as desired. □

2.5 Application to cancer initiation and field cancerization

We now use our main result to explore the dynamics of cancer initiation and field cancerization under a two-step mutational model of cancer. In this section, as in [51], [46] and [63], we assume finite tissue of the form $\mathbb{Z}^2_L \times \mathbb{Z}_w$, where L is chosen so that the total number of cells in the tissue is N, with N typically of order at least 10^6. We now impose the same periodic boundary condition along the first two dimensions as along the third dimension.

Suppose each site in $\mathbb{Z}^2_L \times \mathbb{Z}_w$ is initially occupied by a normal cell (type-0). Each

type-0 cell mutates to a premalignant type-1 cell, with fitness advantage $\beta_1 > 0$ over normal cells, at exponential rate u_1. A type-1 cell gives rise to a successful type-1 clone (one that does not go extinct) with probability $\beta_1/(1+\beta_1)$ by (B.12), in which case its long-run expansion rate $c_w(\beta_1)$ is given by our main theorem (Theorem 2.1). Each type-1 cell mutates to a cancer cell (type-2), with fitness advantage $\beta_2 > 0$ over type-1 cells, at rate u_2 (Fig. 2.7). As before, a type-2 cell gives rise to a successful clone with probability $\beta_2/(1+\beta_2)$. We let σ_2 denote the time at which the first successful type-2 cell arises in the population, which we consider the time of cancer initiation. To simplify the following discussion, we assume that $\beta_1 = \beta_2 =: \beta$.

In [63], Foo, Leder and Ryser analyze an approximated version of the above model for the $w = 1$ case. They assume that cells occupy a spatial continuum, and that type-1 clones grow deterministically with radial growth rate $c_1(\beta)$. Under this simplified model, the dynamics of cancer initiation are governed by the value of the metaparameter

$$\Gamma := N^3(u_1\beta)^3 c_w(\beta)^{-2}(u_2\beta)^{-1}. \tag{2.23}$$

When Γ is small, the first successful cancer cell (type-2) typically arises within the first successful type-1 clone. As Γ increases, it becomes possible for cancer to initiate from one of several successful type-1 clones, and when Γ is large, it may even arise from an unsuccessful type-1 clone before it goes extinct (Fig. 2.8). A more detailed discussion of these regimes can be found in [46], [63] and [66].

Fortunately, the analysis in [63] carries over to the more general $w > 1$ case, with the assumption that type-1 clones grow deterministically as a union of two-dimensional

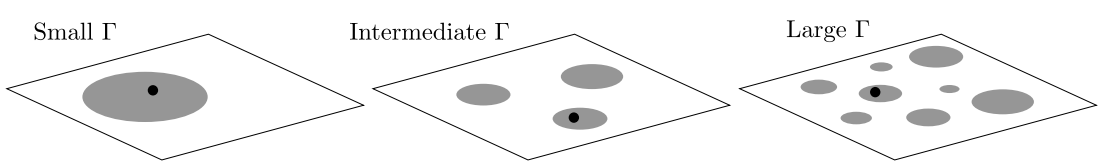

Figure 2.8: The dynamics of cancer initiation on $\mathbb{Z}_L^2 \times \mathbb{Z}_w$ under the two-step mutational model of Section 2.5 are determined by the metaparameter Γ as defined in (2.23). For small Γ, the first successful cancer cell (type-2) arises from the first successful premalignant clone (type-1), but as Γ increases, cancer can originate from one of several successful type-1 clones (intermediate Γ), or even an unsuccessful type-1 clone before it goes extinct (large Γ).

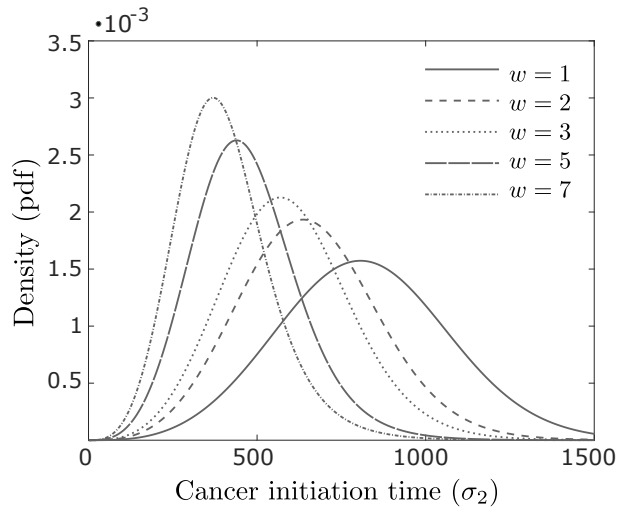

Figure 2.9: The density of the cancer initiation time σ_2 for a few values of the tissue thickness w. Other parameter values are $N = 10^6$, $u_1 = 10^{-6}$, $u_2 = 10^{-5}$ and $\beta = 0.01$.

disks (Fig. 2.2b) with radial growth rate $c_w(\beta) = p_w\sqrt{\pi w\beta}/\sqrt{\log(1/\beta)}$. We begin by considering the metaparameter Γ. Note first the asymmetric role of the mutation rates u_1 and u_2: Increasing u_1 increases the likelihood that multiple successful type-1 clones arise prior to cancer initiation (large-Γ regime), whereas increasing u_2 has the reverse effect (small-Γ regime). Note also the asymmetric role of β and w: As β increases, Γ increases according to $\beta\log(1/\beta)$ for small β, while as w increases, Γ decreases according to $1/(p_w^2 w)$. Both parameters affect how quickly type-1 clones expand, but β also affects the success probability of type-1 clones by (B.12). Thus, whereas a larger w means faster type-1 clonal expansion and a greater chance that cancer initiates within the early clones (small-Γ regime), for larger β, faster type-1 expansion is counterbalanced by the fact that more successful type-1 clones arise, which turns out to push the dynamics toward several successful type-1 clones (large-Γ regime).

We next consider the distribution of σ_2, the time of cancer initiation. Its density is given by (4) in [63], with the substitution $\gamma_2 := \pi w$ (area of stacked unit disks in $\mathbb{Z}^2 \times \mathbb{Z}_w$) and $c_w(\beta) = p_w\sqrt{\pi w\beta}/\sqrt{\log(1/\beta)}$. Predictably, as the tissue thickness w increases, faster type-1 expansion translates into earlier cancer initiation (Fig. 2.9). In Figure 2.3b of Section 2.2, we noted that premalignant population growth is over three times as fast on $w = 5$ layers as on $w = 1$ layer, whereas cancer initiation speeds

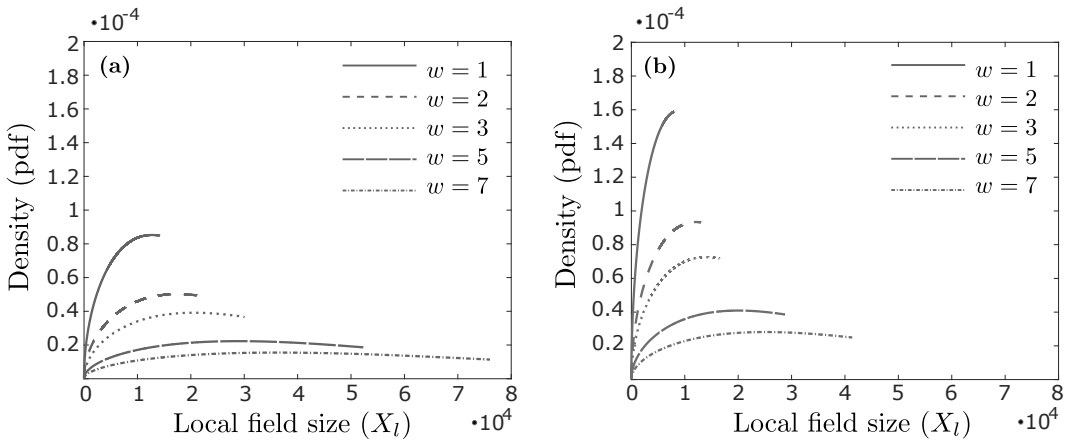

Figure 2.10: Distribution of the local field size X_l, conditioned on $\{\sigma_2 \in dt\}$ with $t = \mathbb{E}[\sigma_2]$, for **(a)** $\beta = 0.01$ and **(b)** $\beta = 0.05$. Other parameters are $N = 10^6$, $u_1 = 10^{-6}$ and $u_2 = 10^{-5}$.

up around twofold over this range according to Figure 2.9. To see why, note that the probability of the event $\{\sigma_2 \in dt\}$ depends on the "total mass" or "spacetime volume" of type-1 particles up until time t, i.e. the time-integral of the size of the type-1 population. Under our deterministic growth assumption, a successful type-1 clone that originates at time 0 grows to size $(c_w(\beta))^2 s^2 \pi w$ by time s, and it reaches spacetime volume V by time

$$t_w(V) = 3^{1/3}(p_w \pi w)^{-2/3} h(\beta)^{1/3} V^{1/3}.$$

Thus, going from $w = 1$ to $w > 1$ layers should accelerate cancer initiation by around $(p_w w)^{2/3}$, which is consistent with a twofold increase from $w = 1$ to $w = 5$. Of course, while these calculations give us some idea of what to expect, the dynamics are more complex in general. For small u_1 or small β, for example, it may take a long time for the first successful type-1 clone to arise, and for larger values, cancer may originate from one of several successful type-1 clones originating at distinct times.

We finally consider the size X_l of the *local field*, i.e. the premalignant clone from which the first successful cancer cell arises, at the time of cancer initiation. The density of X_l, conditioned on the event $\{\sigma_2 \in dt\}$, is given by (8) in [63]. We focus here on the case $t = \mathbb{E}[\sigma_2]$ when cancer initiates at its expected time. In Figure 2.10, we show

how the conditional distribution of X_l changes with tissue thickness w, given fitness advantage $\beta = 0.01$ (Fig. 2.10a) and $\beta = 0.05$ (Fig. 2.10b). Since we condition on $\{\sigma_2 \in dt\}$, and type-1 clones are assumed to grow deterministically, the support of X_l is finite and reflects the maximum possible size of a type-1 clone at time t, which is $(c_w(\beta))^2 t^2 \pi w$.

In Figure 2.10, we see that as w increases, the local field size X_l increases and varies across a wider range. This reflects the fact that increasing w pushes the dynamics toward the small-Γ regime, in which cancer initiates within one or a few large clones. When Figures 2.10a and 2.10b are compared, we see that increasing β results in a smaller local field, and the local field size appears less sensitive to β than to w. As noted above, increasing β both leads to faster expansion of type-1 clones and improved viability of these clones, with counteracting effects on the metaparameter Γ. The fact that X_l decreases with increasing β is consistent with Γ increasing, moving the dynamics toward a greater number of smaller premalignant fields.

The above discussion reveals how our main result enables prediction of how premalignant fields evolve and how they give rise to cancer, given information on the tissue thickness w and the fitness advantage β. We have seen how the number of premalignant patches, the time of cancer initiation σ_2 and the local field size X_l is significantly affected by w, and how w and β affect the dynamics in distinct ways that would be difficult to anticipate without the aid of a mathematical model. These insights are furthermore clinically relevant, since premalignant fields often appear histologically normal, making them difficult to distinguish from normal tissue. Thus, the capability to make quantitative predictions on the spatial evolutionary history of the tumor can yield valuable insights into e.g. optimal excision margins under surgery, and into when and where recurrence can be expected to occur following treatment.

Chapter 3

Exact site frequency spectra of neutrally evolving tumors: A transition between power laws reveals a signature of cell viability

3.1 Introduction

The study of genetic variation driven by neutral mutations has a long history in population genetics [91]. Usually, the population is assumed to be of a large constant size N, and reproduction follows either the Wright-Fisher model (nonoverlapping generations) or the Moran model (overlapping generations) [49]. Neutral mutations occur at rate u per individual per time unit, and each new mutation is assumed to be unique (the *infinite-sites model* of Kimura [92]). This framework gives rise to a sample-based theory of tracing genealogies of extant individuals backwards in time via the *coalescent* [93, 94]. A popular summary statistic of genomic data is the *site frequency spectrum* (SFS), which records the frequencies of mutations in a population or population sample. Under the Moran model with neutral mutations, the expected number of mutations found in j cells of a sample of size n is $\mathbb{E}[\xi_j] = (Nu)(1/j)$ [49], and any linear combination of the

form $\sum_{j=1}^{n-1} jc_j\xi_j$ with $\sum_{j=1}^{n-1} c_j = 1$ is an unbiased estimator of $\theta := Nu$, the population-scaled mutation rate [1, 170].[1] Prominent estimators of this form include Watterson's θ_W [157], Tajima's θ_π [147], Fu and Li's ξ_1 [67] and Fay and Wu's θ_H [58], and these estimators form the basis of several statistical tests of neutral evolution vs. evolution under selection [1, 170]. In this way, the site frequency spectrum has provided a simple means of understanding the evolutionary history of populations from genomic data.

Cancer can be viewed as its own evolutionary process, operating at the somatic level. Cancer initiation is usually understood to be a series of mutational events that culminates in malignant cells able to proliferate uncontrollably [4, 5, 97, 122]. Such "driver" mutations are complemented by more frequent neutral or "passenger" mutations [14, 149], that have no functional role in the evolution to malignancy, but contribute to the genetic diversity characteristic of cancer [25, 111, 154]. The dominant paradigm of tumor progression has been that of sequential clonal expansion of driver mutations. However, several recent works suggest that a neutral evolution model, under which all driver mutations are already present in the tumor-initiating cell, is sufficient to explain the intratumoral heterogeneity in many cancers, see e.g. Sottoriva et al. [142], Ling et al. [108] and Williams et al. [160], and the reviews by Venkatesan and Swanton [151] and [40]. A simple test of neutral tumor evolution based on the site frequency spectrum was proposed in Williams et al. [160], which has since generated debate e.g. surrounding its significance level and statistical power (see e.g. Bozic et al. [17], McDonald et al. [110], Tarabichi et al. [148], with author responses in Heide et al. [79], Werner et al. [158]). The authors of Williams et al. [160] subsequently suggested a Bayesian framework for detecting tumor subclones evolving under selection [161], and more recent approaches to that problem include Dinh et al. [44] and Caravagna et al. [27]. These works and the surrounding debate are indicative both of the fact that increased attention is being paid to the role of neutral evolution in cancer, and that efforts are just underway to develop robust methods of inferring the evolutionary history of tumors [150].

While the constant-sized models of population genetics are appropriate for understanding early cancer development in small tissue compartments, exponential growth

[1] In Theorem 1.33 of Durrett [49], the result is given as $\mathbb{E}[\xi_j] = (2Nu)(1/j)$ for a population of size $2N$.

models are more relevant for understanding long-run tumor progression [50, 124]. In this work, we will employ a stochastic branching process model in which tumor cells divide at rate r_0 and die at rate d_0, with net birth rate $\lambda_0 := r_0 - d_0 > 0$, and w neutral mutations accumulate on average per cell division. Let $p_0 := d_0/r_0$ be the *extinction probability* of the tumor, and let $q_0 := 1 - p_0 = \lambda_0/r_0$ be its *survival probability* (Section 3.2.1). We will show in (3.14) of Section 3.3.2 that if we only consider cells with an infinite line of descent, then at the time the number of cells becomes N, the expected number of mutations found in j cells is $\xi_j = (w/q_0)N \cdot 1/(j(j+1))$ for $2 \leq j \leq N-1$. This SFS differs from the SFS of the constant-sized Moran model of population genetics in two important ways: The spectrum now follows a $1/j^2$ power law as opposed to a $1/j$ law, and it now depends on the growth parameters r_0 and d_0 via the survival probability q_0. Cumulative versions of the $1/j^2$ spectrum have previously been established by Durrett [50, 53], Bozic et al. [16] and Williams et al. [160], as we outline in more detail in Section 5.9. In addition, Williams et al. [160], Bozic et al. [16] and Ling et al. [108] have used the $1/j^2$ spectrum to infer the ratio w/q_0 from tumor data, but extracting information about the mutation parameter w and the growth parameter q_0 separately seemingly requires different tools. In a recent work by Werner et al. [159], the authors measured pairwise mutational differences between the ancestors of spatially separated tumor bulk samples, and they developed a coalescent-based approach for estimating w and a function of p_0 given by (3.5) below.

Tumor evolution is commonly characterized by low cell viability, i.e. a large extinction probability p_0. For example, in a modeling study of cancer recurrence, Avanzini and Antal [7] collected clinical estimates of the tumor volume doubling time and the time between cell divisions in metastatic breast cancer, colorectal cancer, head & neck cancer, lung cancer and prostate cancer. Based on the collected data, they computed a typical net growth rate (λ_0) and division rate (r_0) for metastases of each cancer type, which lead them to estimate $p_0 = 1 - \lambda_0/r_0$ as $0.90, 0.97, 0.95, 0.97$ and 0.76, respectively. Similarly, in an investigation of targeted combination therapy, Bozic et al. [15] estimated an average net growth rate of $\lambda_0 = 0.01$ per day for 21 melanoma lesions, which they combined with a typical division rate of $r_0 = 0.14$ per day [130] to compute a typical death rate of $d_0 = 0.13$ per day. These estimates suggest a typical extinction probability of $p_0 = 0.93$ for the melanoma lesions. Finally, Bozic et al. [16] used their

cumulative version of the $1/j^2$ spectrum to estimate w/q_0 from the SFS of mutations at cell frequency $24\% - 50\%$ in colorectal cancer.[2] For microsatellite stable (MSS) tumor samples, they combined their median estimate of w/q_0 with an independent estimate of the mutation rate [85] to obtain p_0 as 0.997. It should be emphasized that even for a given cancer type, there is substantial heterogeneity between individual tumors (as the clinical data collected in Table 1 of Avanzini and Antal [7] indicate), so these values should only be taken as rough estimates. However, these simple estimates do suggest that low cell viability is broadly relevant to tumor evolution. Low cell viability, and the corresponding high cell turnover, induces a large mutational burden and high genetic diversity, which enhances the adaptability of the tumor under treatment. It is therefore important to understand how low-viability tumors behave, and to explore how they can potentially be identified from genomic data.

Prior theoretical works on the expected SFS of an exponentially growing tumor population offer only a limited understanding of how p_0 affects the spectrum, as these works generally consider only cells with an infinite line of descent, or they consider special cases such as deterministic growth of the tumor bulk or no cell death, which is reasonable when p_0 is small. Moreover, many prior results are given in the large-time or large-population limit, and for practical reasons, the focus is often on mutations of frequency 10% and higher. These results are discussed in more detail in Section 5.9 below. Our goal in this work is to gain a more complete understanding of the SFS of an exponentially growing tumor with neutral mutations. We seek to understand how the spectrum behaves both at small and large frequencies, for all values of p_0, and for any population size N. We obtain separate results for cells with an infinite line of descent and for the total population, evaluated either at a fixed time or at the stochastic time at which the population reaches a certain size, each of which is relevant to tumor data analysis depending on the context. We observe that while the SFS of cells with an infinite line of descent depends on the mutation rate w and the extinction probability p_0 only via the ratio w/q_0, the two parameters decouple in the SFS of the total population.

[2] Bozic et al. [16] estimated w/q_0 from the *variant allele frequency* (VAF) spectrum, which records the proportion of chromosomes carrying the mutations, as opposed to the proportion of cells carrying the mutations. They considered allele frequencies $12\% - 25\%$, which translates to cell frequencies $24\% - 50\%$ in the simplified setting where all cells are diploid and no cell is mutated at the same site on both chromosomes.

	Description
r_0	Division rate of tumor cells (per unit time)
d_0	Death rate of tumor cells (per unit time)
p_0	Extinction probability of the tumor (and of a single-cell derived clone)
q_0	Survival probability of the tumor (and of a single-cell derived clone)
w	Mutation rate (expected number of mutations per cell division)
$\tilde{t}_N$	Fixed time at which the skeleton subpopulation has expected size N
$\tilde{\tau}_N$	Stochastic time at which the skeleton subpopulation first reaches size N
t_N	Fixed time at which a surviving tumor cell population has expected size N
τ_N	Stochastic time at which the tumor cell population first reaches size N

Table 3.1: Notation used in the chapter.

In fact, as p_0 increases from 0 to 1, the small-frequency end of the spectrum transitions from the $1/j^2$ power law characteristic of pure-birth exponential growth to the $1/j$ power law characteristic of a constant-sized population. We investigate simple metrics that quantify this transition, and use one of them to propose a simple estimator for p_0, which we subsequently evaluate using idealized synthetic single-cell sequencing data.

The rest of the chapter is organized as follows. In Section 3.2, we formulate our branching process model with neutral mutations, define the *skeleton subpopulation* of cells with an infinite line of descent, and establish relevant notation. In Section 3.3, we analyze the SFS of skeleton cells, and in Section 4.4, we analyze the SFS of the total cell population. In Section 3.5, we use our theoretical results to propose and evaluate a simple estimator for p_0. In Section 5.9, we summarize our results, and discuss in detail how they relate to the existing literature. The proofs of all of our theoretical results are found in appendices at the end of the dissertation.

3.2 Model description

3.2.1 Branching process dynamics

We assume that the tumor evolution follows a branching process model in continuous time. Cells divide into two cells at rate $r_0 > 0$ per unit time and die at rate $d_0 \geq 0$ per unit time, which means that in a small time interval of length Δt, a cell divides

with probability $r_0\Delta t$ and dies with probability $d_0\Delta t$. Assume $r_0 > d_0$ and define $\lambda_0 := r_0 - d_0 > 0$ as the net growth rate. Let $Z_0(t)$ denote the size of the tumor population at time t and assume $Z_0(0) = 1$, i.e. the tumor expands from a single tumor-initiating cell. Define

$$\Omega_\infty := \{Z_0(t) > 0 \text{ for all } t > 0\}$$

as the event that the tumor does not go extinct, and

$$p_0 := \mathbb{P}(\Omega_\infty^c) = \mathbb{P}(Z_0(t) = 0 \text{ for some } t > 0) \tag{3.1}$$

as the *extinction probability* of the tumor. This probability can be computed as $p_0 = d_0/r_0$ with $0 \leq p_0 < 1$, see e.g. Section 3 of Durrett [50]. Note that any clone derived from a single tumor cell gives rise to its own branching process with the same growth parameters r_0 and d_0 and the same extinction probability. We also define

$$q_0 := 1 - p_0 = \lambda_0/r_0 \tag{3.2}$$

as the *survival probability* of the tumor or of a single-cell derived clone.

3.2.2 Decomposition into skeleton cells and finite-family cells

On the nonextinction event Ω_∞, the cells alive at time $t > 0$ can be split into two categories, one consisting of cells with an infinite line of descent, i.e. cells that start clones that do not go extinct, and the other consisting of cells whose descendants eventually go extinct. We refer to the former cells as *skeleton* cells and the latter as *finite-family* cells. An arbitrary tumor cell is a skeleton cell with probability q_0, so in the long run, the proportion of skeleton cells in the population is q_0. We can think of skeleton cells as forming the trunk and scaffold branches of the genealogical branching tree, with finite-family clones growing out from the skeleton as lateral branches, see Figure 3.1a.

3.2.3 Mutation accumulation

We next add neutral mutations under the infinite-sites model. Prior to a cell division, each parental DNA molecule is unwound and separated into two complementary strands.

Each parental strand serves as a template for the construction of a new complementary daughter strand. The end result is two copies of the DNA molecule, each consisting of one parental and one daughter strand. Errors in nucleotide pairing during this process can result in one or more point mutations per daughter strand. We assume that these errors amount to $w/2$ mutations on average per daughter strand, for a total of w mutations on average per cell division. Note that the only assumptions we make on the distribution of the number of mutations is that it is nonnegative and integer-valued with a finite mean. The point mutation rate has been estimated as $5 \cdot 10^{-10}$ per base pair per cell division [85], and it is commonly higher in cancer due to genomic instability [25]. Since the number of base pairs is of order 10^7 in the exome and 10^9 in the genome, it makes sense to allow w to be any positive number, i.e. $w \in (0, \infty)$. In many works, the convention is to allow at most one mutation per cell division, introducing a probability $u \in (0, 1)$ of a new mutation. Since our analysis only depends on the mean number of mutations w per cell division, it includes this case with $w := u$. We assume that the mutation rate is constant throughout tumor evolution, which ignores e.g. the possibility of an elevated mutation rate over time due to genomic instability.

While our focus in this work is on discrete mutation accumulation, we will also present all of our results in terms of continuous mutation accumulation, another common and biologically relevant assumption. In the continuous model, neutral mutations occur at rate $\nu > 0$ per cell per unit time, at any time throughout the lifetime of the cell. In other words, each cell undergoes a neutral mutation in a small time interval of length Δt with probability $\nu \Delta t$. The continuous model differs from the discrete model in that at most one mutation occurs at a time, and this mutation occurs in between cell divisions with probability 1. However, as we will show, the mean behavior of the two models is similar when $\nu = wr_0$, since in the discrete model, each cell accumulates $wr_0 \Delta t$ mutations on average in a small time interval of length Δt.

3.2.4 Clonal and subclonal mutations

Before proceeding, we need to make a distinction between *clonal* and *subclonal* mutations. A mutation is clonal if it is shared by all tumor cells, while it is subclonal if there is at least one tumor cell without it. As an example, say the tumor-initiating cell divides into two cells, A and B, and that cell A acquires a new mutation. This mutation

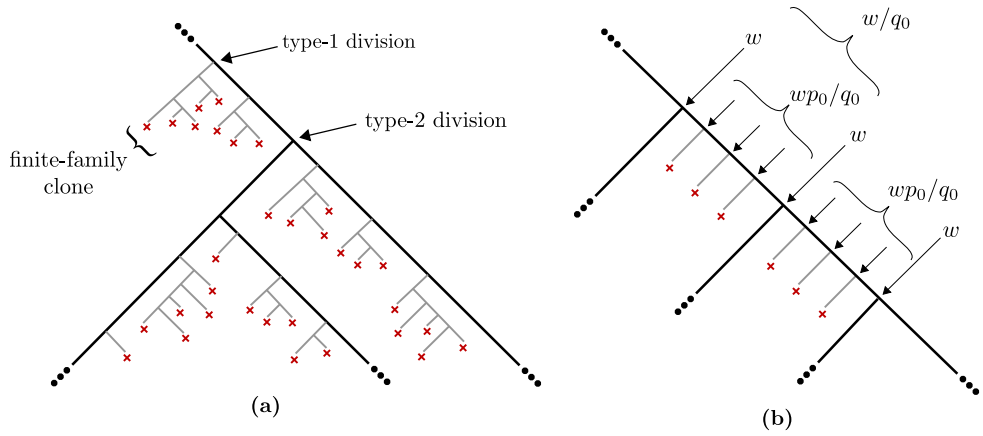

Figure 3.1: Categorization of cells into skeleton cells and finite-family cells, and mutation accumulation on the skeleton. **(a)** The skeleton subpopulation, which consists of cells with an infinite line of descent, can be thought of as forming the trunk and scaffold branches of the genealogical branching tree (bold branches). Each skeleton cell divides into one skeleton cell and one finite-family cell at rate $2r_0p_0 = 2d_0$ (type-1 skeleton division, see (3.3) in Section 3.3.1), in which case a finite-family clone grows out from the skeleton as a lateral branch (light gray branches). A skeleton cell divides into two skeleton cells at rate $r_0q_0 = \lambda_0$ (type-2 skeleton division, see (3.4) in Section 3.3.1), in which case another scaffold branch is added. **(b)** In between two type-2 skeleton divisions, the expected number of mutations that accumulate due to type-1 divisions is wp_0/q_0 by (3.8). Each type-2 division adds w mutations on average and starts a new skeleton population size level. Thus, the expected number of mutations that accumulate on the skeleton per size level is $w + wp_0/q_0 = w/q_0$.

is initially subclonal. However, if the clone started by cell B dies out, the mutation in cell A becomes clonal from that point onward. While this example demonstrates how clonal mutations can arise post-tumor-initiation, all mutations that accumulate prior to initiation, as the cancer precursor cell evolves to malignancy, also become clonal. For this reason, the clonal mutations usually tell us more about the events preceding cancer than the dynamics post-initiation, and they can in fact outnumber the subclonal mutations [149]. Nevertheless, clonal mutations do appear in the SFS of mutations post-initiation, and they play distinct roles in the fixed-time and fixed-size spectrum, which is why we pay them special attention below.

3.3 Site frequency spectrum of skeleton cells

In this section, we establish the expected fixed-time and fixed-size spectrum of skeleton cells. The reason we are interested in analyzing skeleton cells separately is twofold:

- When $p_0 = 0$ (no cell death), all cells are skeleton cells, and the SFS of the total population is the SFS of the skeleton. More generally, when p_0 is small, the total population spectrum is well-approximated by the simpler skeleton spectrum.
- When the extinction probability p_0 is large, finite-family cells affect the SFS of the total population at the small-frequency end. However, the large-frequency end is still characterized by the skeleton spectrum, as we demonstrate in Section 4.4 below.

3.3.1 Effective rates of cell division and mutation

When the tumor is conditioned on nonextinction, the probability that the tumor-initiating cell divides during the first Δt units of time, and that exactly one of the two daughter cells survives, i.e. starts a clone that does not go extinct, is

$$\frac{\mathbb{P}(\text{division in } [0, \Delta t], \text{ one offspring survives})}{\mathbb{P}(\Omega_\infty)} = \frac{r_0 \Delta t \cdot 2p_0(1-p_0)}{1-p_0} = 2r_0 p_0 \Delta t, \quad (3.3)$$

and the probability of a division in $[0, \Delta t]$ where both daughter cells survive is

$$\frac{\mathbb{P}(\text{division in } [0, \Delta t], \text{ both offspring survive})}{\mathbb{P}(\Omega_\infty)} = \frac{r_0 \Delta t \cdot (1-p_0)^2}{1-p_0} = r_0 q_0 \Delta t. \quad (3.4)$$

Since each skeleton cell starts a clone that does not go extinct, we can conclude that a skeleton cell divides into one skeleton cell and one finite-family cell at rate $2r_0p_0 = 2d_0$ per unit time (type-1 skeleton division), and it divides into two skeleton cells at rate $r_0q_0 = \lambda_0$ per unit time (type-2 skeleton division). The probability that a skeleton division is type-2 is

$$\kappa_0 := \frac{r_0 q_0}{r_0 q_0 + 2r_0 p_0} = \frac{1-p_0}{1+p_0}. \quad (3.5)$$

This probability can also be computed directly as follows: $(1-p_0)^2$ is the probability that both daughter cells survive, and $1-p_0^2$ is the probability that at least one of them does, so $(1-p_0)^2/(1-p_0^2) = (1-p_0)/(1+p_0)$ is the probability that a skeleton division is type-2.

Let $\tilde{Z}_0(t)$ denote the number of skeleton cells at time t, conditional on the nonextinction event Ω_∞. Since type-1 divisions do not affect the size of the skeleton, we can think of the type-2 divisions as the "effective" divisions. More precisely, $(\tilde{Z}_0(t))_{t\geq 0}$ is a pure-birth exponential growth process, known as a *Yule process*, with birth rate λ_0 and mean size $\mathbb{E}[\tilde{Z}_0(t)] = e^{\lambda_0 t}$ at time t [50, 123]. Type-1 divisions do contribute to neutral mutation accumulation however. Indeed, each type-1 division adds $w/2$ mutations on average to the skeleton, and each type-2 division adds w mutations on average. The rate at which mutations accumulate on the skeleton is then

$$(w/2) \cdot 2r_0 p_0 + w \cdot r_0(1-p_0) = wr_0 \tag{3.6}$$

per skeleton cell per unit time, which equals the mutation rate for the original, unconditioned process $(Z_0(t))_{t\geq 0}$. The mutation rate per type-2 division, or the *effective* mutation rate, is on the other hand

$$wr_0/\lambda_0 = w/q_0 \tag{3.7}$$

per unit time. We can also think of mutations as accumulating across skeleton population size levels as follows. A type-2 division increases the size of the skeleton population by one, and it adds w mutations on average. Upon the type-2 division, the number of type-1 divisions before the next type-2 division has the geometric distribution with support $\{0, 1, 2 \ldots\}$ and success probability κ_0 given by (3.5). It follows that in between the two type-2 divisions, the expected number of mutations that accumulate on the skeleton is

$$(w/2) \cdot (1/\kappa_0 - 1) = (w/2) \cdot 2p_0/q_0 = wp_0/q_0. \tag{3.8}$$

At each population size level, the skeleton therefore accumulates w mutations on average due to the type-2 division that starts the level, and wp_0/q_0 mutations on average due

to type-1 divisions that occur before the next type-2 division that changes levels. We thus obtain

$$w + wp_0/q_0 = w/q_0 \tag{3.9}$$

mutations per level (Figure 3.1b). The effective mutation rate w/q_0 plays a key role in the SFS of the skeleton, with the continuous-time viewpoint in (3.7) applying to the fixed-time spectrum, and the population-size-level viewpoint in (3.9) applying to the fixed-size spectrum.

3.3.2 Expected fixed-time and fixed-size skeleton spectrum

Let $\tilde{S}_j(t)$ denote the number of mutations that are found in $j \geq 1$ skeleton cells at time t, conditional on the nonextinction event Ω_∞. This is the site frequency spectrum of skeleton cells. For any integer $N \geq 1$, define

$$\tilde{t}_N := \log(N)/\lambda_0 \tag{3.10}$$

as the (fixed) time at which the skeleton has expected size N, i.e. $e^{\lambda_0 \tilde{t}_N} = N$, and define

$$\tilde{\tau}_N := \inf\{t \geq 0 : \tilde{Z}_0(t) = N\} \tag{3.11}$$

as the (stochastic) time at which the skeleton reaches size N. In Proposition 3.1 below, we provide the expected SFS of the skeleton evaluated both at time $\tilde{t}_N$ (fixed-time spectrum) and at time $\tilde{\tau}_N$ (fixed-size spectrum). Both the fixed-time and fixed-size spectrum can be relevant to tumor data analysis depending on the context. For example, *in vitro* cell culture experiments and *in vivo* mouse experiments are often conducted over a fixed time period, in which case the fixed-time spectrum would apply. In the clinic, however, the size of a tumor sample is more readily estimated than its age, in which case the fixed-size spectrum is more relevant [99]. It is therefore useful to understand both spectra and to what extent they differ.

Proposition 3.1. *(1) Define $\tilde{t}_N$ as in* (3.10)*. Then, for any $N \geq 1$ and any $j \geq 1$,*

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] &= (w/q_0)N \cdot \textstyle\int_0^{1-1/N}(1-y)y^{j-1}dy \\ &= (w/q_0)N \cdot \big(1-\tfrac{1}{N}\big)^j\big(\tfrac{1}{j(j+1)} + \tfrac{1}{N}\tfrac{1}{j+1}\big).\end{aligned} \tag{3.12}$$

For fixed $j \geq 1$, then as $N \to \infty$,

$$\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] \sim (w/q_0)N \cdot 1/(j(j+1)), \tag{3.13}$$

where $f(y) \sim g(y)$ as $y \to \infty$ means $\lim_{y\to\infty} f(y)/g(y) = 1$.

(2) Define $\tilde{\tau}_N$ as in (3.11)*. Then, for any $N \geq 2$,*

$$\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] = \begin{cases}(w/q_0)N \cdot 1/(j(j+1)) - (wp_0/q_0)\delta_{1,j}, & 1 \leq j \leq N-1, \\ wp_0/q_0 = w/q_0 - w, & j = N,\end{cases} \tag{3.14}$$

where $\delta_{\ell,m} = 1$ if $\ell = m$ and $\delta_{\ell,m} = 0$ otherwise.

Proof. Appendix B.1. □

Analogous results for continuous mutation accumulation are presented in Appendix B.3. In Figure 3.2, we compare our fixed-time (3.12) and fixed-size (3.14) results with simulation results for $w = 1$, $p_0 = 0$ and $N = 100$. In this example, there are no clonal mutations, since $p_0 = 0$. The fundamental difference between the fixed-time and fixed-size spectrum is that the skeleton size at time $\tilde{t}_N$ is variable, while it is always N at time $\tilde{\tau}_N$. The fixed-time spectrum therefore has nonzero mass at $j > N$, due to instances in which the skeleton is larger than N at time $\tilde{t}_N$. It is however natural to ask how the fixed-time spectrum restricted to $j = 1, \ldots, N$ relates to the fixed-size spectrum. By (3.14), the fixed-size spectrum follows the power law $(w/q_0)N \cdot 1/(j(j+1))$ exactly on $j = 2, \ldots, N-1$, and asymptotically as $N \to \infty$ for $j = 1$. By (3.13), the fixed-time spectrum converges to the same power law for fixed $j \geq 1$ as $N \to \infty$, which means that it follows this power law when N is large and $j \ll N$. In Figure 3.3, we compare the fixed-time (3.12) and fixed-size (3.14) spectrum for $N = 10^3$ and $p_0 = 0.9$. As expected, the two spectra agree on $j \ll N$, while the fixed-time spectrum deviates from the fixed-size spectrum at the very largest frequencies (for j of order N). In Figure 3.3b,

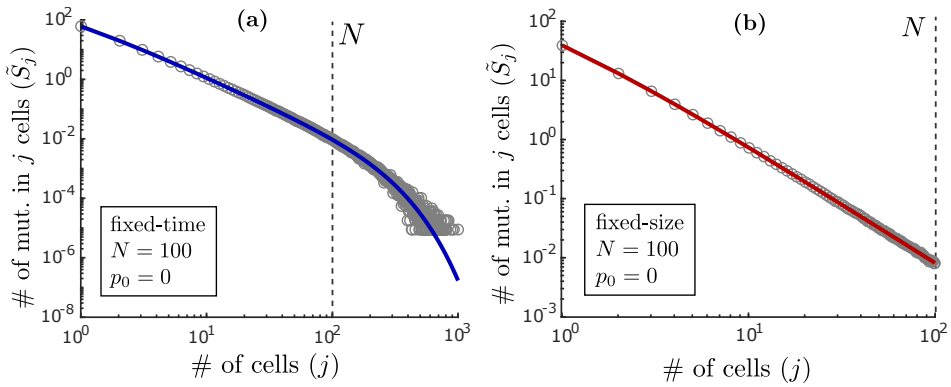

Figure 3.2: Comparison between the expected SFS of the skeleton, as derived in Proposition 3.1, and simulation results. **(a)** The expected fixed-time spectrum (3.12) of Proposition 3.1 (solid blue line) shows good agreement with the average spectrum of simulated tumors (grey dots). In this example, the extinction probability is $p_0 = 0$, the mutation rate is $w = 1$, and the expected size of the skeleton is $N = 100$, which is also the expected size of the tumor since $p_0 = 0$. We generated 10^5 tumors with $p_0 = 0$ and $w = 1$ and stopped each simulation at the fixed time $\tilde{t}_N$ with $N = 100$ as defined by (3.10). At each cell division, the number of mutations acquired by each daughter cell was generated as a Poisson random variable with mean $w/2$. **(b)** The expected fixed-size spectrum (3.14) of Proposition 3.1 (solid red line) shows good agreement with the average spectrum of simulated tumors (grey dots). We again generated 10^5 tumors with $p_0 = 0$ and $w = 1$, but this time, we stopped each simulation when the tumor reached size $N = 100$, i.e. at the stochastic time $\tilde{\tau}_N$ defined by (3.11). The fundamental difference between the fixed-time and fixed-size spectrum is that the skeleton size at time $\tilde{t}_N$ is variable, while it is always N at time $\tilde{\tau}_N$. As a result, the fixed-size spectrum is restricted to $j = 1, \ldots, N$, while the fixed-time spectrum has nonzero mass at values $j > N$.

we show that almost all mutations are found on $j \ll N$. In this example, the difference between the two spectra is within 1% on the range $j = 1, \ldots, 150$, on which 99.3% of mutations are found.

We next note the sharp discontinuity at $j = N$ in the fixed-size spectrum of Figure 3.3a, which does not appear in the fixed-time spectrum. This is due to the distinct ways in which clonal mutations manifest in the two spectra. In the fixed-time spectrum, clonal mutations can appear at any value of j, depending on the skeleton size at time $\tilde{t}_N$. In the fixed-size spectrum, all these mutations are concentrated at $j = N$, which creates a significant point mass at $j = N$. Note that by (3.14), the expected number

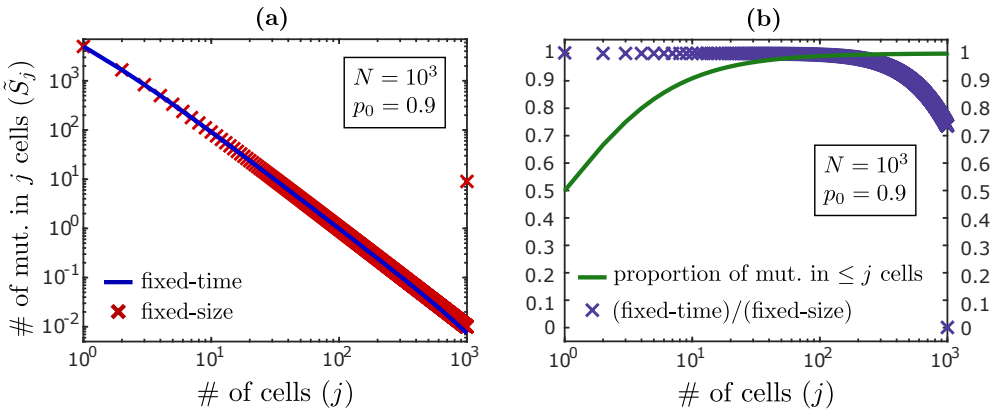

Figure 3.3: Comparison between the expected fixed-time spectrum (3.12) and fixed-size spectrum (3.14) of the skeleton as derived in Proposition 3.1. **(a)** The fixed-time spectrum (3.12) of the skeleton (blue curve) is a good approximation of the fixed-size spectrum (3.14) (red crosses) for mutations at frequencies $j \ll N$, given parameters $N = 10^3$, $p_0 = 0.9$ and $w = 1$. The two spectra diverge at the very largest frequencies, and the difference is substantial at $j = N$, since clonal mutations are concentrated at $j = N$ in the fixed-size spectrum, while they are scattered in the fixed-time spectrum. **(b)** Here, we show the ratio of the fixed-time spectrum to the fixed-size spectrum (purple crosses), and the proportion of mutations found in $\leq j$ skeleton cells in the fixed-size spectrum (green curve). The fixed-time and fixed-size spectrum are virtually the same on $j \ll N$, where almost all mutations are found.

of mutations found in $j = N - 1$ skeleton cells is given by $(w/q_0)(1/(N-1))$ in the fixed-size spectrum, which is of order $1/N$ as $N \to \infty$, while clonal mutations are given by the constant wp_0/q_0, independently of N.

The observed difference between the fixed-time and fixed-size spectrum at the very largest frequencies reflects the unbounded range of the fixed-time spectrum. When N is large, this difference can be alleviated by computing the SFS of mutations found in a given *proportion* of cells (as opposed to a given *number* of cells), as this normalizes both spectra to the frequency range $[0, 1]$. To justify this claim, we need to combine our results with results previously obtained by Durrett [50, 53] and Bozic et al. [16], as we discuss in detail in Section 5.9 below.

3.3.3 Proportion of mutations found in one cell

We conclude this section by computing two simple and important metrics derived from the site frequency spectrum of the skeleton. Note first that by Proposition 3.1, the expected number of mutations found in a single skeleton cell is

$$\begin{aligned} \mathbb{E}[\tilde{S}_1(\tilde{t}_N)] &= (1/2)(w/q_0)(N^2-1)/N \sim (1/2)(w/q_0)N, \\ \mathbb{E}[\tilde{S}_1(\tilde{\tau}_N)] &= (1/2)(w/q_0)N - wp_0/q_0 \sim (1/2)(w/q_0)N, \end{aligned} \tag{3.15}$$

as $N \to \infty$. Next, let $\tilde{M}_j(t) := \sum_{k \geq j} \tilde{S}_k(t)$ denote the number of mutations found in $\geq j$ skeleton cells at time t. Again, by Proposition 3.1, the expected total number of mutations on the skeleton is

$$\begin{aligned} \mathbb{E}[\tilde{M}_1(\tilde{t}_N)] &= (w/q_0)(N-1) \sim (w/q_0)N, \\ \mathbb{E}[\tilde{M}_1(\tilde{\tau}_N)] &= (w/q_0)(N-1) \sim (w/q_0)N, \end{aligned} \tag{3.16}$$

as $N \to \infty$. Expressions (3.15) and (3.16) suggest that for large N, half the mutations discovered at time $\tilde{t}_N$ or $\tilde{\tau}_N$ are found in only one cell. This is a consequence of the pure-birth exponential growth of the skeleton. Indeed, note that if we only consider the effective type-2 skeleton divisions, the total number of divisions required to reach generation k is $\sum_{j=0}^{k-1} 2^j = 2^k - 1$. The expected total number of mutations in generation k is then $(w/q_0)(2^k - 1)$, which is (3.16) with $N = 2^k$ the number of cells in generation k. An additional 2^k divisions are required to reach generation $k+1$, so each generation roughly doubles the total number of mutations. Of course, our model is stochastic, it operates in continuous time, and generations may overlap, but this simple discrete argument gives intuition as to why half the mutations are found in one cell, and more generally why most mutations are found at the smallest frequencies.

3.4 Site frequency spectrum of total population and transition between power laws

When the extinction probability p_0 is small, the SFS of the total population $(Z_0(t))_{t\geq 0}$ is well-approximated by the SFS of the skeleton $(\tilde{Z}_0(t))_{t\geq 0}$. However, tumor evolution is commonly characterized by a large extinction probability, as was discussed in the introduction. In this section, we investigate the expected fixed-time and fixed-size spectrum of the total population $(Z_0(t))_{t\geq 0}$ for all values of p_0. We show that as p_0 increases, the small-frequency end of the spectrum starts to deviate from the skeleton spectrum of Section 3.3, and that as p_0 approaches 1, it transitions to the spectrum of a constant-sized population.

3.4.1 Expected fixed-time and fixed-size total population spectrum

Let $S_j(t)$ denote the number of mutations found in $j \geq 1$ cells at time t, the site frequency spectrum of the total cell population. We wish to compute the mean of $S_j(t)$ conditioned on the tumor surviving to time t. We can compute the probability of this survival event as

$$\mathbb{P}(Z_0(t) > 0) = q_0 e^{\lambda_0 t}/(e^{\lambda_0 t} - p_0), \quad t \geq 0,$$

see (B.6) of Appendix B.2, and the expected size of a tumor that survives to time t as

$$\mathbb{E}[Z_0(t)|Z_0(t) > 0] = (e^{\lambda_0 t} - p_0)/q_0, \quad t \geq 0. \tag{3.17}$$

Note that $\mathbb{E}[Z_0(t)|Z_0(t) > 0] \sim e^{\lambda_0 t}/q_0$ as $t \to \infty$, which means that the long-run expected growth of a tumor conditioned on survival is exponential, and the initial value of the exponential growth function is given by $1/q_0$. An interesting consequence of the conditioning on survival is that if r_0 and d_0 are increased by the same amount (so that $\lambda_0 = r_0 - d_0$ stays fixed), the survival probability $q_0 = \lambda_0/r_0$ will decrease, while the size of a tumor conditioned on survival will increase. Now, for any integer $N \geq 1$, define

$$t_N := \log(q_0 N + p_0)/\lambda_0 \tag{3.18}$$

as the (fixed) time at which a surviving tumor has expected size N, i.e. $(e^{\lambda_0 t_N} - p_0)/q_0 = N$, and define

$$\tau_N := \inf\{t \geq 0 : Z_0(t) = N\} \tag{3.19}$$

as the (stochastic) time at which the tumor reaches size N, with $\inf \varnothing = \infty$. In Proposition 3.2, we provide the expected SFS of $(Z_0(t))_{t\geq 0}$ at time t_N and τ_N, conditioned on the survival events $\{Z_0(t_N) > 0\}$ and $\{\tau_N < \infty\}$, respectively. In this case, we cannot obtain an explicit expression for the fixed-size spectrum, and provide instead a computational expression. Note however that for the case $p_0 = 0$ of no cell death, the fixed-size spectrum follows the power law $wN \cdot 1/(j(j+1))$ exactly on $j = 2, \ldots, N-1$, by (3.14) of Proposition 3.1.

Proposition 3.2. *(1) Define t_N as in* (3.18)*. Then, for any $N \geq 1$ and any $j \geq 1$,*

$$\mathbb{E}[S_j(t_N)|Z_0(t_N) > 0] = wN \int_0^{1-1/N} (1 - p_0 y)^{-1}(1-y)y^{j-1}dy. \tag{3.20}$$

For fixed $j \geq 1$, then as $N \to \infty$,

$$\begin{aligned} \mathbb{E}[S_j(t_N)|Z_0(t_N) > 0] &\sim wN \cdot \int_0^1 (1 - p_0 y)^{-1}(1-y)y^{j-1}dy \\ &= wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{(j+k)(j+k+1)}, \end{aligned} \tag{3.21}$$

where $f(y) \sim g(y)$ as $y \to \infty$ means $\lim_{y\to\infty} f(y)/g(y) = 1$.

(2) Define τ_N as in (3.19)*, let $\mathcal{S} := \{(\ell, m) : \ell, m \geq 0$ and $\ell + m \leq N\}$ and $A := \{(0,0)\} \cup \{(r,s) : r, s \geq 0$ and $r + s = N\}$. Then, for any $N \geq 2$ and any $1 \leq j \leq N$,*

$$\mathbb{E}[S_j(\tau_N)|\tau_N < \infty] = (w/q_0) \cdot \sum_{k=1}^{N-1} (1 - p_0^{N-k}) \cdot h_{(1,k)}^{(j,N-j)}, \tag{3.22}$$

where for each $(r,s) \in A$, the vector $\left(h_{(\ell,m)}^{(r,s)}\right)_{(\ell,m)\in\mathcal{S}}$ solves the system

$$(\ell + m)(1 + p_0)h_{(\ell,m)}^{(r,s)} = \ell h_{(\ell+1,m)}^{(r,s)} + \ell p_0 h_{(\ell-1,m)}^{(r,s)} + m h_{(\ell,m+1)}^{(r,s)} + m p_0 h_{(\ell,m-1)}^{(r,s)} \tag{3.23}$$

for $\ell, m \geq 1$ and $\ell + m < N$, with boundary conditions given by (B.14) *in Appendix B.2.*

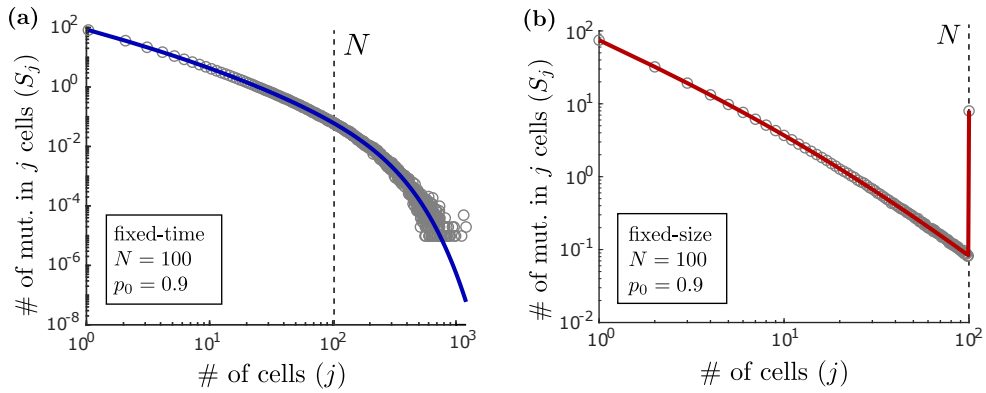

Figure 3.4: Comparison between the expected SFS of the total population, as derived in Proposition 3.2, and simulation results. **(a)** The expected fixed-time spectrum (3.20) of Proposition 3.2 (solid blue line) shows good agreement with the average spectrum of simulated tumors (grey dots). In this example, the extinction probability is $p_0 = 0.9$, the mutation rate is $w = 1$, and the expected tumor size is $N = 100$. We generated 10^5 tumors with $p_0 = 0.9$ and $w = 1$ and stopped each simulation at the fixed time t_N with $N = 100$ as defined by (3.18). At each cell division, the number of mutations acquired by each daughter cell was generated as a Poisson random variable with mean $w/2$. **(b)** The expected fixed-size spectrum (3.22) of Proposition 3.2 (solid red line) shows good agreement with the average spectrum of simulated tumors (grey dots). We again generated 10^5 tumors with $p_0 = 0.9$ and $w = 1$, but this time, we stopped each simulation when the tumor reached size $N = 100$, i.e. at the stochastic time τ_N defined by (3.19). Note the discontinuity in the fixed-size spectrum at $j = N$, which is due to clonal mutations.

Proof. Appendix B.2. □

Analogous results for continuous mutation accumulation are presented in Appendix B.3. In Figure 3.4, we compare our fixed-time (3.20) and fixed-size (3.22) results with simulation results for $w = 1$, $p_0 = 0.9$ and $N = 100$. The fundamental difference between the fixed-time and fixed-size spectrum is the same as we observed in Section 3.3.2. In Figure 3.5a, we compare the two spectra in more detail for $N = 100$ and $p_0 \in \{0, 0.5, 0.7, 0.9\}$. The fixed-time spectrum is a good approximation of the fixed-size spectrum on $j \ll N$ for all but $p_0 = 0.9$, in which case there is a significant difference even on $j \ll N$. In Figure 3.5b, we show that this difference reduces as N increases, with the expected number of mutations found in one cell being 1.46% off for $p = 0.9$ and $N = 1000$. The number of cells in 1 cm^3 of tumor tissue is around $10^7 - 10^9$ [41], in which

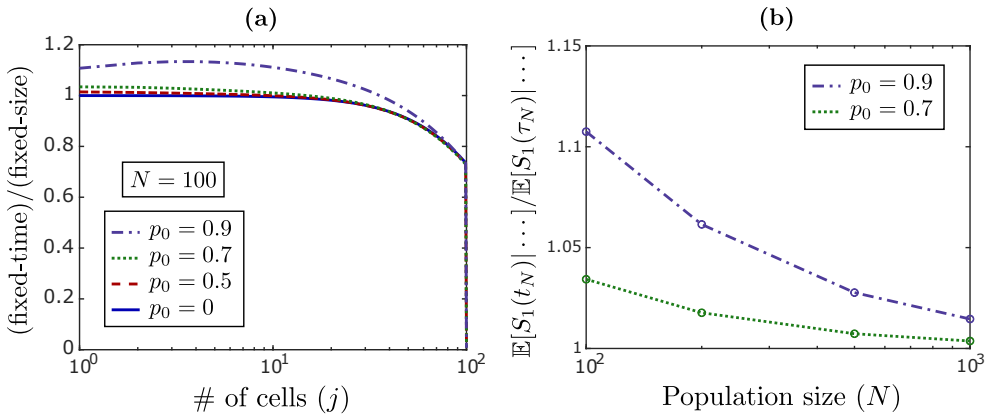

Figure 3.5: Comparison between the expected fixed-time spectrum (3.20) and fixed-size spectrum (3.22) of the total population as derived in Proposition 3.2. **(a)** Here, we show the ratio between the fixed-time spectrum (3.20) and the fixed-size spectrum (3.22) for $N = 100$, $p_0 \in \{0, 0.5, 0.7, 0.9\}$ and $w = 1$. The fixed-time spectrum is a good approximation of the fixed-size spectrum on $j \ll N$ for smaller values of p_0, but the two spectra start to diverge as p_0 increases. **(b)** Here, we show the ratio between the expected number of mutations found in one cell for the fixed-time spectrum and the fixed-size spectrum, as a function of the population size N, for $p_0 \in \{0.7, 0.9\}$ and $w = 1$. As N increases, the difference between the two spectra reduces both for $p_0 = 0.7$ and $p_0 = 0.9$.

case the fixed-time spectrum can generally be expected to be a good approximation of the fixed-size spectrum on $j \ll N$. The two spectra may diverge, however, when applied to smaller tumor samples (e.g. 1 mm^3 or smaller) or when p_0 is very close to 1.

In Appendix B.5, we present heuristic calculations that indicate how the fixed-size spectrum can be approximated by the fixed-time spectrum on $j \ll N$ as $N \to \infty$. The key insight is that once the tumor has reached a large size, its growth becomes essentially deterministic with exponential rate λ_0, which allows us to approximate the probabilities $h_{(1,k)}^{(j,N-j)}$ in (3.22) by continuous-time probabilities following Iwasa et al. [81]. For added intuition, we refer to our discussion in Section 3.4.5 below and accompanying calculations in Appendix B.8, where we conjecture a law of large numbers for the fixed-time and fixed-size spectrum, whose limits agree in the mean with the asymptotic expected fixed-time spectrum (3.21).

For each $j = 1, \ldots, N$, computing $\mathbb{E}[S_j(\tau_N)|\tau_N < \infty]$ according to (3.22) requires solving a linear system of the form (3.23), which has order $N^2/2$ equations. A more

general version of this system arises in the study of the number of wild-type and mutant cells under a two-type (Luria-Delbrück) population model stopped at a certain size, see e.g. Komarova et al. [99]. The coefficient matrix of the system is sparse and banded, and it has a certain structure which allows one to solve it in $O(N^3)$ arithmetic operations [71], compared to $O\big((N^2)^3\big) = O(N^6)$ for Gaussian elimination. Solving the system directly is still not computationally feasible for the largest values of N, in which case one must develop an approximate solution. See Komarova et al. [99] for a partial differential equations approach. Our observations suggest another simple approach: For N sufficiently large, one can approximate the fixed-size spectrum by the fixed-time spectrum on $j \ll N$. For j of order N, one can apply the $1/j^2$ skeleton law of (3.14), since as we discuss next, the skeleton will be responsible for the large-frequency mutations.

3.4.2 Transition between power laws

By Proposition 3.1 of Section 3.3.2, the SFS of the skeleton depends on the mutation rate w and the extinction probability p_0 only through the effective mutation rate w/q_0. By Proposition 3.2, however, the two parameters decouple in the SFS of the total population. The mutation rate w scales the total population spectrum linearly, and the same is true of a large population size N, but the dependence on p_0 is more complex. To better understand how p_0 affects the spectrum, recall first that by (3.21), the asymptotic expected fixed-time spectrum is given by

$$\mathbb{E}[S_j(t_N)|Z_0(t_N) > 0] \sim wN \cdot \textstyle\sum_{k=0}^{\infty} \frac{p_0^k}{(j+k)(j+k+1)}, \quad N \to \infty.$$

In Appendix B.6, we show that for fixed $0 < p_0 < 1$, sending $j \to \infty$ in this expression yields

$$wN \cdot \textstyle\sum_{k=0}^{\infty} \frac{p_0^k}{(j+k)(j+k+1)} \sim (w/q_0)N \cdot 1/(j(j+1)), \quad j \to \infty, \tag{3.24}$$

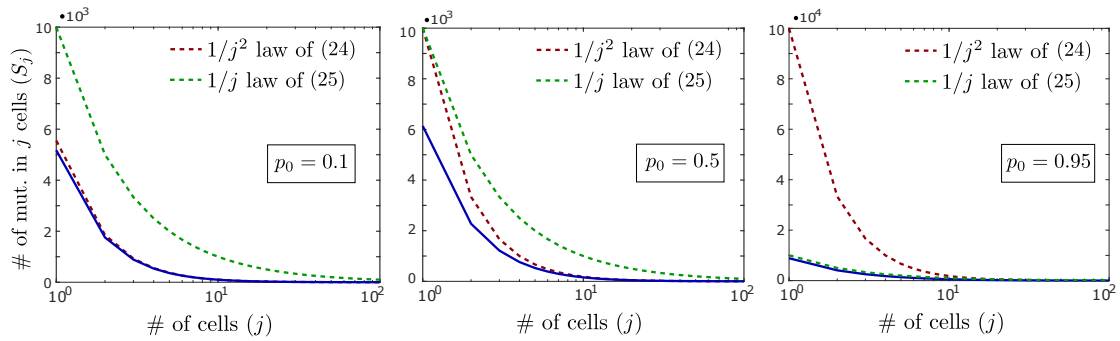

Figure 3.6: The small-frequency end of the total population spectrum transitions between two power laws as p_0 increases from 0 to 1. When p_0 is small, the expected fixed-time spectrum (3.20) of Proposition 3.2 (solid blue line) approximately follows the the $1/j^2$ law of (3.24) (dotted red line). As p_0 increases to 1, the small-frequency end of the spectrum transitions to the $1/j$ law of (3.25) (dotted green line), while the $1/j^2$ law increases according to $(w/q_0)N \cdot 1/(j(j+1))$. Note that the y-axis is on a linear scale, and that the scale of the rightmost panel is ten times larger than the scale of the other two. Also note that the $1/j$ law of (3.25) is fixed as a function of p_0 and is therefore the same curve in all panels. Parameters are $N = 10^4$ and $w = 1$, and the spectrum is shown only at the smallest frequencies $j = 1, \ldots, 100$.

which is the $1/j^2$ power law given in (3.13)-(3.14) of Proposition 3.1. We also show that for fixed $j \geq 1$, sending $p_0 \to 1$ in the same expression yields

$$wN \cdot \sum_{k=0}^{\infty} \frac{p_0^k}{(j+k)(j+k+1)} \sim wN \cdot 1/j, \quad p_0 \to 1, \tag{3.25}$$

which is the $1/j$ power law of the constant-sized Moran model of population genetics, see Theorem 1.33 of Durrett [49]. When referring to results from Durrett [49], note that there, the population is assumed to have size $2N$, which is a common convention in population genetics.

Recall that when $p_0 = 0$, the SFS of the total population $(Z_0(t))_{t\geq 0}$ is the SFS of the skeleton $(\tilde{Z}_0(t))_{t\geq 0}$. Expressions (3.24) and (3.25) suggest that when $p_0 > 0$, the SFS of $(Z_0(t))_{t\geq 0}$ continues to follow the $1/j^2$ skeleton law at the large-frequency end, while a deviation starts to occur at the small-frequency end. In fact, as p_0 approaches 1, the small-frequency end transitions to the $1/j$ law of the constant-sized Moran model. This is illustrated in Figure 3.6 for $p_0 \in \{0.1, 0.5, 0.95\}$ and $N = 10^4$. Note that the y-axis is on a linear scale, and that the scale of the rightmost panel is ten times larger

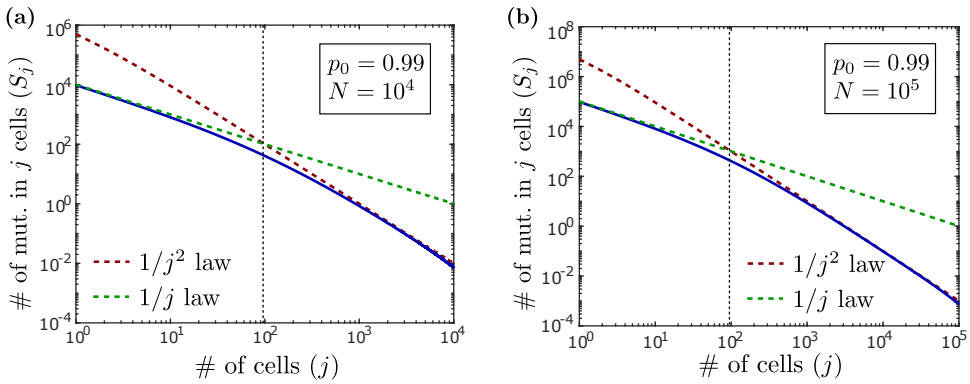

Figure 3.7: For a fixed, large value of the extinction probability p_0, the expected fixed-time spectrum (3.20) of Proposition 3.2 (solid blue line) transitions from the $1/j^2$ power law of (3.24) (dotted red line) at the large-frequency end to the $1/j$ power law of (3.25) (dotted green line) at the small-frequency end. The dotted vertical line shows the intersecting point $j = 1/q_0 - 1$ of the two power laws in (3.24) and (3.25). In **(a)**, the parameters are $N = 10^4$, $p_0 = 0.99$ and $w = 1$, and in **(b)**, we increase the expected tumor size to $N = 10^5$. Note that the transition between power laws occurs around the same value of j in both cases, $j = 1/q_0 - 1$.

than the scale of the other two. Also note that the $1/j$ law in (3.25) is independent of p_0, so it is the same curve in all panels. For small p_0, the fixed-time spectrum (3.20) is well-approximated by the $1/j^2$ law (3.24), and each lies below the $1/j$ law (3.25). To see why, note that the $1/j$ law gives wN as the number of mutations found in one cell, while the $1/j^2$ law gives $(1/2)(w/q_0)N$, which is $(1/2)wN$ for $p_0 = 0$. As p_0 increases to 1, the fixed-time spectrum (3.20) transitions to the $1/j$ law, while the $1/j^2$ law increases according to $(w/q_0)N \cdot 1/(j(j+1))$, diverging further and further from the fixed-time spectrum. In Figure 3.7a, we show how for a fixed, large value of p_0 ($p_0 = 0.99$), the fixed-time spectrum (3.20) transitions from the $1/j^2$ law at the large-frequency end to the $1/j$ law at the small-frequency end. Note that the power laws of (3.24) and (3.25) intersect at

$$j = 1/q_0 - 1,$$

which gives an indication of the frequency at which the transition occurs. This intersecting point is independent of N, as is illustrated in Figure 3.7b (dotted vertical line).

To understand this transition between power laws, note that mutations that occur

early in the evolution of a large tumor are only detected if they occur on the skeleton, which is why mutations at large frequencies in a large tumor follow the $1/j^2$ skeleton spectrum. For mutations that occur late, we need to consider both skeleton cells and finite-family cells. When p_0 is large, most cells are finite-family cells (their long-run proportion is p_0), and as p_0 approaches 1, finite-family clones start to behave like a critical branching process with net growth rate 0. This is why late mutations follow the $1/j$ law of a constant-sized population. Thus, even though the branching process dynamics of cell division and cell death are the same throughout the evolution of the tumor, from the perspective of mutation accumulation, the tumor effectively behaves like a pure-birth exponential growth process initially, and more like a constant-sized process at the end, assuming a large extinction probability p_0. The transition between power laws is the transition between these two growth regimes.

Note finally that as we observed for the skeleton spectrum in Section 3.3.2, the fixed-time spectrum (3.20) deviates from the $1/j^2$ law at the very largest frequencies (for j of order N) in Figure 3.7, due to the variability in tumor size at time t_N. It is therefore more correct to say that for fixed large p_0, the fixed-time spectrum transitions from the skeleton spectrum at the large-frequency end to the constant-sized spectrum at the small-frequency end, without reference to the power laws. It remains true, however, that as p_0 increases from 0 to 1, the small-frequency end of the spectrum transitions between the two power laws.

3.4.3 Total mutational burden of the tumor

We next wish to quantify how p_0 affects overall mutation accumulation. To this end, we derive in Proposition 3.3 the expected total mutational burden of the tumor, both under the fixed-time and fixed-size spectrum. This quantity indicates the genetic diversity of the tumor, which has implications e.g. for its adaptability under treatment. It also enables us to compute a normalized version of the SFS, which can be useful for parameter estimation, as is discussed further in Section 3.5 below. Before stating the proposition, we define $M_j(t) := \sum_{k \geq j} S_k(t)$ as the number of mutations found in $\geq j$ cells at time t.

Proposition 3.3. *(1) For $0 < p_0 < 1$, the expected total number of mutations in the*

fixed-time spectrum is given by

$$\begin{aligned}\mathbb{E}[M_1(t_N)|Z_0(t_N)>0] &= -wN\cdot(1/p_0)\log(q_0+p_0/N)\\ &\sim -wN\cdot\log(q_0)/p_0, \quad N\to\infty.\end{aligned} \tag{3.26}$$

For $p_0=0$, $\mathbb{E}[M_1(t_N)]=w(N-1)\sim wN$ *as* $N\to\infty$ *by* (3.16).

(2) *Define* $\mathcal{S}$ *and* A *as in Proposition 3.2. The expected total number of mutations in the fixed-size spectrum is given by*

$$\mathbb{E}[M_1(\tau_N)|\tau_N<\infty] = (w/q_0)\cdot\textstyle\sum_{k=1}^{N-1}(1-p_0^{N-k})\big(1-h_{(1,k)}^{(0,N)}-h_{(1,k)}^{(0,0)}\big), \tag{3.27}$$

where for each $(r,s)\in A$, *the vector* $\big(h_{(\ell,m)}^{(r,s)}\big)_{(\ell,m)\in\mathcal{S}}$ *solves the linear system* (3.23) *of Proposition 3.2. For* $p_0=0$, $\mathbb{E}[M_1(\tau_N)]=w(N-1)\sim wN$ *as* $N\to\infty$ *by* (3.16).

Proof. Appendix B.4. □

Now, for ease of notation, write $\overline{M}_1 := \mathbb{E}[M_1(t_N)|Z_0(t_N)>0]$ for the expected total mutational burden under the fixed-time spectrum. We are interested in comparing $\overline{M}_1$ with the expected number of mutations under the $1/j^2$ skeleton law of (3.24). To this end, define

$$\widehat{M}_1 := (w/q_0)N, \tag{3.28}$$

following (3.16). This simple estimate has been used e.g. in Ling et al. [108], where the authors estimate the total number of mutations in a hepatocellular carcinoma (HCC) tumor under a few different assumptions on tumor evolution. The ratio between $\overline{M}_1$ and $\widehat{M}_1$ is, for $0<p_0<1$,

$$\begin{aligned}\overline{M}_1/\widehat{M}_1 &= -(q_0/p_0)\cdot\log(q_0+p_0/N)\\ &\sim -(q_0/p_0)\cdot\log(q_0), \quad N\to\infty.\end{aligned} \tag{3.29}$$

In Figure 3.8a, we show this ratio as a function of p_0 for $N=1000$. The ratio is decreasing in p_0, it converges to 1 as $p_0\to 0$, and it converges to 0 as $p_0\to 1$. To give some examples, in the $N\to\infty$ limit, $\overline{M}_1/\widehat{M}_1=0.46$ for $p_0=0.75$ and $\overline{M}_1/\widehat{M}_1=0.047$ for $p_0=0.99$. Thus, if one is interested in estimating the total number of mutations

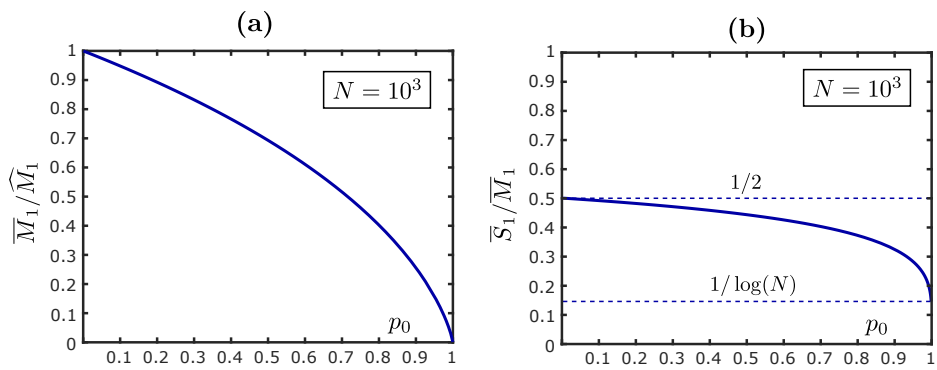

Figure 3.8: Graphs of the metrics $\overline{M}_1/\widehat{M}_1$ and $\overline{S}_1/\overline{M}_1$ given by (3.29) and (3.31), as a function of p_0. **(a)** Ratio between the expected total number of mutations in the fixed-time spectrum, $\overline{M}_1 := \mathbb{E}[M_1(t_N)|Z_0(t_N) > 0]$ given by (3.26), and the simple estimate $\widehat{M}_1$ of (3.28), derived from the $1/j^2$ skeleton law of (3.24), as a function of p_0 for $N = 10^3$. The two estimates $\overline{M}_1$ and $\widehat{M}_1$ agree for $p_0 = 0$, but $\widehat{M}_1$ becomes a significant overestimate of $\overline{M}_1$ as p_0 increases. **(b)** The proportion of mutations found in one cell, $\overline{S}_1/\overline{M}_1 = \varphi_N(p_0)$ as given by (3.31), quantifies the transition between the $1/j^2$ and $1/j$ power laws at the small-frequency end of the spectrum. Note that the transition between power laws accelerates as p_0 increases.

in an exponentially growing tumor, using the simple expression (3.28) implied by the skeleton spectrum will result in a significant overestimate when p_0 is large. Indeed, as p_0 increases, finite-family cells start to dominate the population, and they accumulate mutations less efficiently than skeleton cells.

3.4.4 Proportion of mutations found in one cell

At the extremes $p_0 = 0$ and $p_0 = 1$, the small-frequency end of the SFS is characterized by the $1/j^2$ skeleton law and the $1/j$ constant-sized law, respectively. To better understand how the small-frequency end behaves for intermediate values of p_0, we next determine the relative proportion of mutations found at the very smallest frequency, i.e. in one cell. This metric quantifies the transition between the $1/j^2$ and $1/j$ power laws at the small-frequency end, and it enables us to propose a simple estimator for p_0 in Section 3.5 below.

In Appendix B.7, we show that in the fixed-time spectrum, the expected number of

mutations found in one cell is given by, for $0 < p_0 < 1$,

$$\begin{aligned}\mathbb{E}[S_1(t_N)|Z_0(t_N) > 0] &= wN \cdot (1/p_0)\big(1 - 1/N + (q_0/p_0)\log(q_0 + p_0/N)\big) \\ &\sim wN \cdot (1/p_0)(1 + (q_0/p_0)\log(q_0)), \quad N \to \infty.\end{aligned} \tag{3.30}$$

As in Section 3.4.3, we write $\overline{S}_1 := \mathbb{E}[S_1(t_N)|Z_0(t_N) > 0]$ for ease of notation. We then define $\varphi_N(p_0) := \overline{S}_1/\overline{M}_1$ as the proportion of mutations found in one cell. By (3.26) and (3.30),

$$\begin{aligned}\varphi_N(p_0) &= -\big(1 - 1/N + (q_0/p_0)\log(q_0 + p_0/N)\big) \,/\, \log(q_0 + p_0/N) \\ &\sim -(1 + (q_0/p_0)\log(q_0))/\log(q_0), \quad N \to \infty.\end{aligned} \tag{3.31}$$

In Figure 3.8b, we show $\varphi_N = \overline{S}_1/\overline{M}_1$ as a function of p_0 for $N = 1000$. The function is strictly decreasing in p_0, it converges to 0.50 as $p_0 \to 0$, and it converges to $(1 - 1/N)/\log(N)$ as $p_0 \to 1$, which is of order $1/\log(N)$ for N large. In the $p_0 \to 0$ regime, the SFS of the total population $(Z_0(t))_{t\geq 0}$ is the SFS of the skeleton $(\tilde{Z}_0(t))_{t\geq 0}$, in which case half the mutations are found in one cell by Section 3.3.3. In the $p_0 \to 1$ regime, the SFS of $(Z_0(t))_{t\geq 0}$ is the SFS of the constant-sized Moran model, in which case the expected number of mutations is of order $wN\sum_{j=1}^{N} 1/j \sim wN\log(N)$ for N large, and the proportion of mutations found in one cell is of order $1/\log(N)$. Note that the rate of change of $\varphi_N(p_0)$ increases as p_0 increases (Figure 3.8b). This implies that the deviation from the $1/j^2$ skeleton law at the small-frequency end is initially slow for small values of p_0, but it accelerates as p_0 increases and transitions quickly to the $1/j$ law for large values of p_0. It also implies that $\overline{S}_1/\overline{M}_1$ is more useful for distinguishing larger values of p_0 than smaller values, as will become more apparent in Section 3.5 below.

3.4.5 Spectra of individual large tumors (laws of large numbers)

The results of Proposition 3.2 hold in expectation, meaning that they apply to an average SFS computed over a large number of tumors. If we want to use these results to understand the evolutionary history of individual tumors, we need to know more about how well they apply on a tumor-by-tumor basis. It is well-known that conditional on the nonextinction event Ω_∞, $Z_0(t) \sim Ye^{\lambda_0 t}$ as $t \to \infty$ almost surely, where Y follows the exponential distribution with mean $1/q_0$ (Theorem 1 of Durrett [50]). In other words,

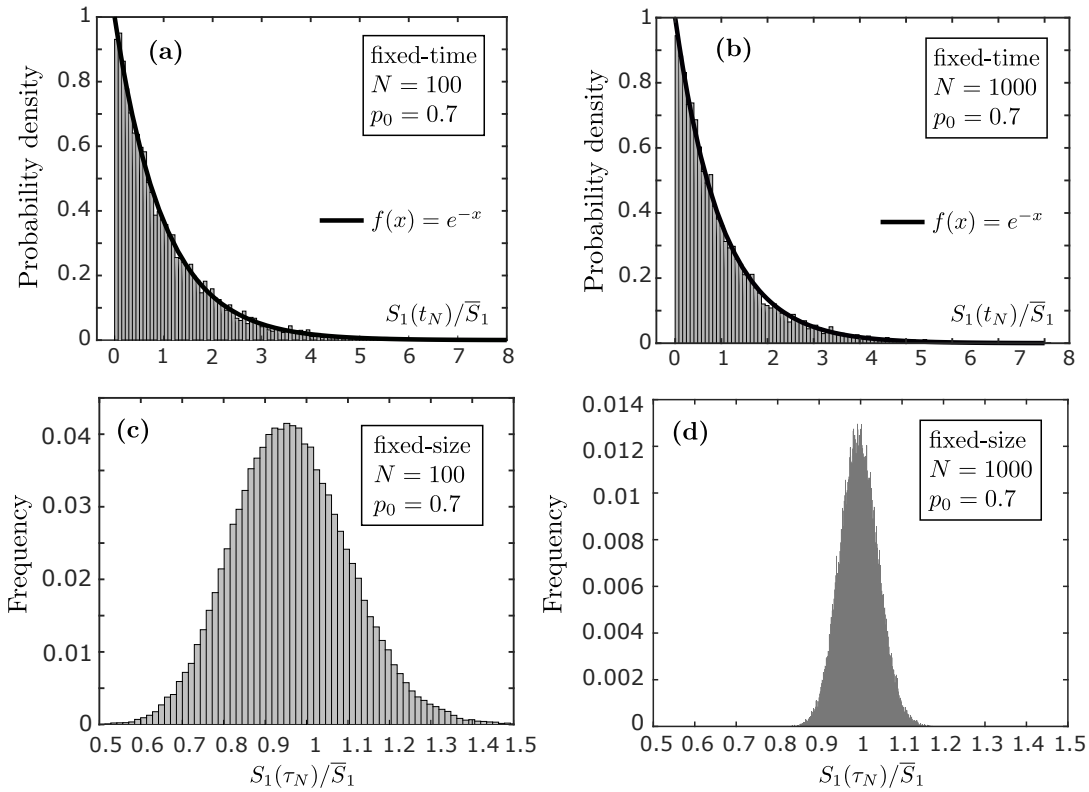

Figure 3.9: Simulation results support the conjectured laws of large numbers (3.32) and (3.33) for the fixed-time and fixed-size spectrum respectively. **(a)** Histogram of $S_1(t_N)/\overline{S}_1$ over 10^4 simulation runs with $N = 100$, $p = 0.7$ and $w = 1$, where t_N is defined by (3.18), and $\overline{S}_1 := \mathbb{E}[S_1(t_N)|Z_0(t_N) > 0]$ as given by (3.30). The y-axis is normalized so as to approximate the density of the underlying probability distribution. By comparison with $x \mapsto e^{-x}$, we see that $S_1(t_N)/\overline{S}_1$ appears to be a mean-1 exponential random variable, which is consistent with the conjectured law of large numbers (3.32). **(b)** When the population size is increased to $N = 1000$, $S_1(t_N)/\overline{S}_1$ retains the mean-1 exponential distribution, consistent with (3.32). **(c)** Histogram of $S_1(\tau_N)/\overline{S}_1$ over 10^4 simulation runs with $N = 100$, $p = 0.7$ and $w = 1$, where τ_N is defined by (3.19). **(d)** Same as in (c), except now $N = 1000$. Together, (c) and (d) indicate that the ratio $S_1(\tau_N)/\overline{S}_1$ concentrates around 1 as N increases, which is consistent with the conjectured law of large numbers (3.33).

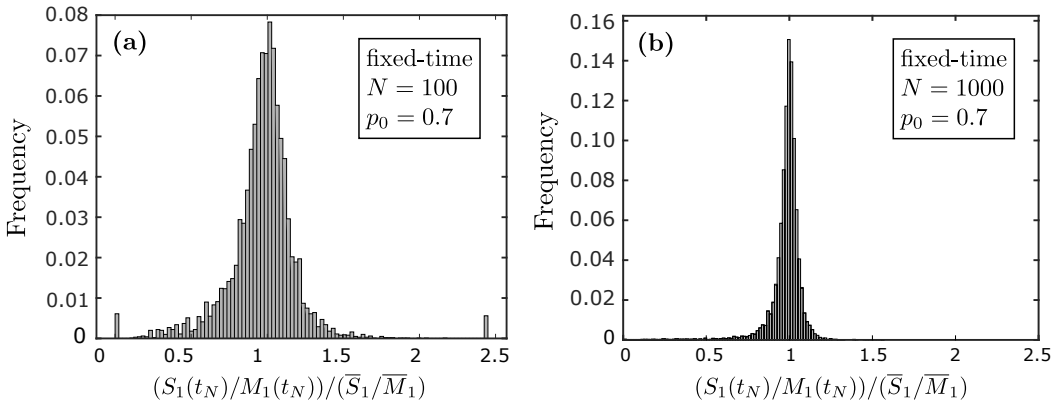

Figure 3.10: Simulation results indicate that if the SFS is normalized by the total number of mutations, the fixed-time and fixed-size spectrum obey the same law of large numbers. **(a)** Histogram of $(S_1(t_N)/M_1(t_N))/(\overline{S}_1/\overline{M}_1)$ over 10^4 simulation runs with $N = 100$, $p = 0.7$ and $w = 1$, where t_N is defined by (3.18), and the expected ratio $\overline{S}_1/\overline{M}_1$ is given by (3.31). Note the point masses at 0 and 2.44, which represent simulation runs where $S_1(t_N)/M_1(t_N) = 0$ and $S_1(t_N)/M_1(t_N) = 1$, respectively. **(b)** Same as in (a), except now, $N = 1000$. Together, (a) and (b) indicate that as N increases, $S_1(t_N)/M_1(t_N)$ concentrates around $\overline{S}_1/\overline{M}_1$. This in turn suggests that if the fixed-time spectrum is normalized by the total number of mutations $M_1(t_N)$, it obeys the same law of large numbers as the normalized version of the fixed-size spectrum.

the tumor population $Z_0(t)$ eventually grows at exponential rate λ_0, but the initial value of the exponential growth function is random and depends on the individual tumor. We can use this fact to formulate laws of large numbers for the fixed-time and fixed-size spectrum, which we state formally as conjectures. In Appendix B.8, we present simple calculations in support of these conjectures, and we also prove an analogous result (B.19) for a simplified, semideterministic version of our model.

Conjecture. (1) Define t_N as in (3.18). Then, there exists an exponential random variable X with mean 1 so that for fixed $j \geq 1$, conditional on Ω_∞,

$$S_j(t_N) \sim X \cdot wN \cdot \int_0^1 (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy \tag{3.32}$$

as $N \to \infty$ almost surely.

(2) Define τ_N as in (3.19). Then, for fixed $j \geq 1$, conditional on Ω_∞,

$$S_j(\tau_N) \sim wN \cdot \int_0^1 (1 - p_0 y)^{-1}(1 - y) y^{j-1} dy \tag{3.33}$$

as $N \to \infty$ almost surely.

Both the fixed-time conjecture (3.32) and the fixed-size conjecture (3.33) agree with simulation results, see Figures 3.9 and 3.10. The main difference between (3.32) and (3.33) is that the right-hand-side of (3.32) is stochastic, while the right-hand side of (3.33) is a constant. The former expression has a random scaling factor X, which captures the variability in tumor size at time t_N, whereas the fact that the tumor size is always N at time τ_N eliminates this variability in the latter expression. Note that since $\mathbb{E}[X] = 1$, the right-hand sides of (3.32) and (3.33) agree in the mean, and this mean agrees with the right-hand side of the asymptotic spectrum (3.21) of Proposition 3.2. Importantly, the scaling factor X in (3.32) is independent of j, which indicates that the proportion of mutations found in j cells is the same in (3.21), (3.32) and (3.33). In other words, according to these conjectures, if we normalize the SFS with the total number of mutations, $S_j(t)/M_1(t)$, the fixed-time and fixed-size spectrum of an individual large tumor will be completely characterized by the asymptotic expected spectrum (3.21) of Proposition 3.2 (Figure 3.10). This can be useful for parameter estimation, as we discuss next.

3.5 Signatures of cell viability

In this section, we use our theoretical results to propose a simple estimator for the extinction probability p_0, based on extracting one or more spatially separated subclones from a tumor. By a subclone, we mean all currently living descendants of a given common ancestor, i.e. all leaves of the branching tree started by a given tumor cell. Since every clone or subclone derived from a single tumor cell obeys the same branching process dynamics as the overall tumor, all of our previous results can be applied to individual subclones.

We make the strong assumption that each cell in each sampled subclone can be

single-cell sequenced so that all its mutations are captured, even at the smallest frequencies, which is beyond current sequencing technology. Our main purpose with this section is to show how the information contained in the small-frequency end of the SFS can in principle be used to decouple the mutation rate w and the extinction probability p_0. We discuss practical considerations and potential alternative approaches in Section 5.9 below.

Say that we sample a subclone of size n. For $1 \leq j \leq n-1$, let s_j be the number of mutations found in j cells of the subclone, and let $m_1 := \sum_{j=1}^{n-1} s_j$ be the total number of mutations. Here, we ignore mutations found in all cells of the subclone, since they include (i) mutations that accumulate prior to tumor initiation, (ii) mutations that occur post-tumor-initiation but prior to initiation of the subclone, and (iii) mutations that occur post-subclone-initiation but still end up in all subclone cells. Let $\overline{s}_1$ be the expected number of mutations found in one subclone cell under the fixed-time spectrum, and let $\overline{m}_1$ be the expected total number of mutations under the fixed-time spectrum. By (3.31), we can write

$$\overline{s}_1/\overline{m}_1 = \varphi_n(p_0),$$

where $\varphi_n(p_0)$ is continuous and strictly decreasing in p_0. In particular, $\varphi_n(p_0)$ is invertible. This implies that given $\overline{s}_1$ and $\overline{m}_1$, p_0 can be recovered from this expression via

$$p_0 = \varphi_n^{-1}(\overline{s}_1/\overline{m}_1).$$

For the sampled values s_1 and m_1, this suggests the following estimator for p_0:

$$\widehat{p}_0(n) := \varphi_n^{-1}(s_1/m_1). \tag{3.34}$$

Of course, if the subclone is sampled at a certain size, it makes more sense to use the fixed-size spectrum than the fixed-time spectrum. In addition, m_1 excludes mutations found in all subclone cells, while $\overline{m}_1$ includes some of these mutations. These potential sources of error are minor and can easily be resolved if necessary, as we discuss in more detail below.

Note that the ratio of expected values $\overline{s}_1/\overline{m}_1$ takes values in $[(1-1/n)/\log(n), 1/2]$ by Section 3.4.4, whereas due to stochasticity, the sampled ratio s_1/m_1 can take any

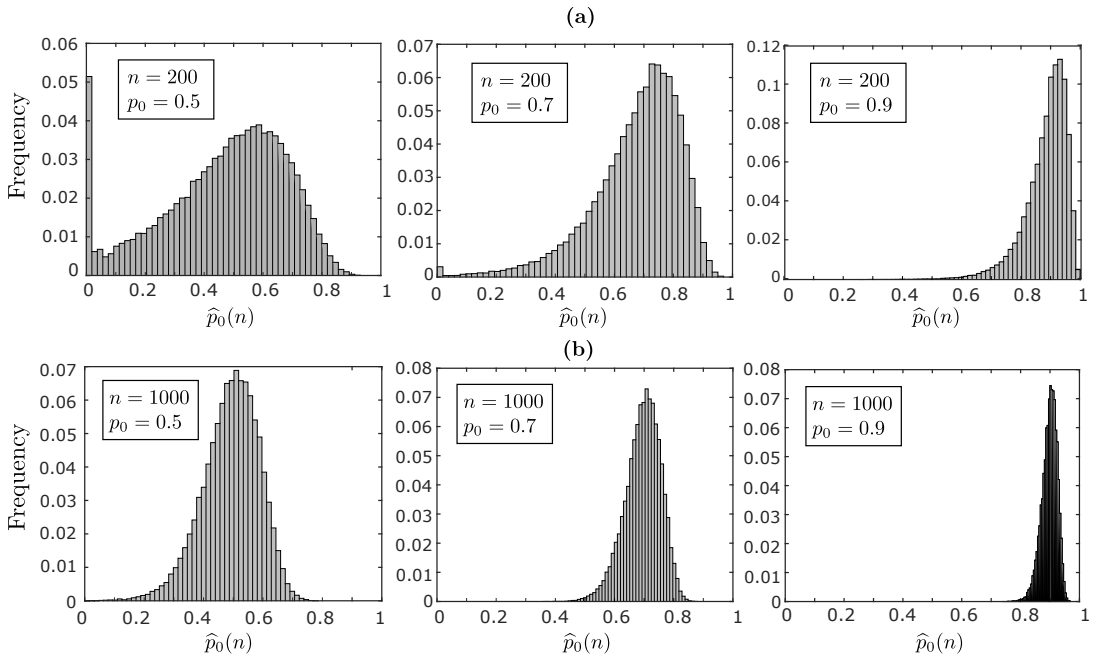

Figure 3.11: Histogram of the estimator $\widehat{p}_0(n)$ defined in (3.34) computed across 10^5 synthetic subclone samples of size n, given true values $p_0 \in \{0.5, 0.7, 0.9\}$ and $w = 1$. In **(a)**, the subclone size is $n = 200$, and in **(b)**, the size is $n = 1000$. Together, (a) and (b) indicate that as the size of the subclone increases, the estimator $\widehat{p}_0(n)$ concentrates around the true value of p_0, which in turn indicates statistical consistency of the estimator.

value in $[0, 1]$. To complete the estimator in (3.34), we therefore extend the definition of φ_n^{-1} by setting

$$\varphi_n^{-1}(x) := \begin{cases} 0, & 1/2 \leq x \leq 1, \\ 1, & 0 \leq x \leq (1 - 1/n)/\log(n). \end{cases} \tag{3.35}$$

For example, if we observe a ratio s_1/m_1 larger than $1/2$, we default to the estimate $\widehat{p}_0(n) = 0$, since the expected ratio $\overline{s}_1/\overline{m}_1$ is largest (and equal to $1/2$) for $p_0 = 0$. Once p_0 has been estimated, an estimate for the mutation rate w can be obtained from (3.30) or (3.26).

n	p_0	mean	median	st. dev.	defaults to 0 (%)	defaults to 1 (%)
200	0.5	0.4694	0.5041	0.2099	0.0459	0
	0.7	0.6755	0.7045	0.1505	0.0026	0
	0.9	0.8839	0.9001	0.0725	0	0.00002
1000	0.5	0.4921	0.5000	0.0948	0.00002	0
	0.7	0.6957	0.7010	0.0603	0	0
	0.9	0.8979	0.9009	0.0266	0	0

Table 3.2: Performance metrics for the estimator $\widehat{p}_0(n)$ computed from the data that underlies Figure 3.11. Both for $n = 200$ and $n = 1000$, the median of $\widehat{p}_0(n)$ accurately recovers the true value of p_0. In addition, the estimator improves in terms of standard deviation both as n increases and as p_0 increases.

The estimator $\widehat{p}_0(n)$ has the benefit of being simple to define and to compute. However, as was mentioned above, it may make more sense to use the fixed-size spectrum than the fixed-time spectrum, and $\overline{m}_1$ includes clonal mutations that arise post-subclone-initiation, whereas m_1 excludes these mutations. The second potential source of error is minor, and it can easily be removed simply by subtracting from $\overline{m}_1$ the contribution from mutations shared by all subclone cells. The first potential source of error is also likely to be insignificant when n is large, since by the conjectured laws of large numbers in Section 3.4.5, the normalized spectrum of an individual large subclone is robust to whether it is observed at a fixed time or a fixed size. For smaller values of n, one can replace $\overline{s}_1$ and $\overline{m}_1$ by the corresponding quantities (3.22) and (3.27) for the fixed-size spectrum, which we denote here by $\overline{\overline{s}}_1$ and $\overline{\overline{m}}_1$. It remains true that we can write $\overline{\overline{s}}_1/\overline{\overline{m}}_1 = \psi_n(p_0)$ for some function ψ_n of p_0, which allows us to define an estimator for p_0 as before. However, the fixed-size estimator has to be obtained numerically, e.g. by precomputing $\psi_n(p_0)$ over a grid of values for p_0, and minimizing the error between the observed ratio s_1/m_1 and the expected ratio $\overline{\overline{s}}_1/\overline{\overline{m}}_1$ over the grid.

To evaluate $\widehat{p}_0(n)$, we use computer simulations to generate multiple independent subclones of size n with true extinction probability p_0, and for each generated subclone, we compute the estimate $\widehat{p}_0(n)$. In Figure 3.11a, we show a histogram for $\widehat{p}_0(n)$ across 10^5 synthetic subclone samples of size $n = 200$ with true extinction probabilities $p_0 \in \{0.5, 0.7, 0.9\}$. In Table 3.2, we show performance metrics for $\widehat{p}_0(n)$ computed across the 10^5 samples. For $p_0 = 0.5$, the estimator defaults to $\widehat{p}_0(n) = 0$ for 4.6% of the subclone

samples, for $p_0 = 0.7$, it defaults to 0 in 0.3% of cases, and for $p_0 = 0.9$, it never defaults to 0. For all values of p_0, the median estimate of $\widehat{p}_0(n)$ accurately recovers the true value, and as p_0 increases, the quality of the estimate improves in terms of standard deviation. In Figure 3.11b, we increase the subclone size to $n = 1000$ and observe a marked improvement in the quality of $\widehat{p}_0(n)$. This indicates statistical consistency of the estimator, meaning that $\widehat{p}_0(n) \to p_0$ in probability as $n \to \infty$, which would also be a direct consequence of the conjectured laws of large numbers (3.32)-(3.33) and the continuous mapping theorem. In other words, the estimator appears to recover the true value of p_0 with arbitrarily high precision given a sufficiently large subclone. Note that the size of n required to return a high-precision estimate becomes smaller as p_0 increases, making $\widehat{p}_0(n)$ especially useful when p_0 is large. Indeed, as we remarked in Section 3.4.4, the rate of change of the expected ratio $\overline{s}_1/\overline{m}_1$ increases as p_0 increases, making it more useful for distinguishing between larger values of p_0 than smaller values.

Whenever it is possible to do multi-region sampling, there may be benefits to extracting multiple small, spatially separated subclones over a single large one. In this more general setting, we sample $K \geq 1$ subclones of size n. For $1 \leq j \leq n-1$, let s_j^k be the number of mutations found in j cells of subclone number k, and let $m_1^k := \sum_{j=1}^{n-1} s_j^k$ be the total number of mutations in subclone k. We replace s_1 and m_1 in the definition of $\widehat{p}_0(n)$ in (3.34) by the sums $\sum_{k=1}^{K} s_1^k$ and $\sum_{k=1}^{K} m_1^k$ to obtain the estimator

$$\widehat{p}_0(n, K) := \varphi_n^{-1}\big(\textstyle\sum_{k=1}^{K} s_1^k / \sum_{k=1}^{K} m_1^k\big). \tag{3.36}$$

Of course, $\widehat{p}_0(n, 1) = \widehat{p}_0(n)$. In Figure 3.12, we show a histogram for $\widehat{p}_0(n, K)$ evaluated across 10^5 synthetic samples, each sample consisting of $K = 5$ independent subclones of size $n = 200$. In Table 3.3, we show performance metrics computed across the 10^5 samples. Qualitatively, the histograms in Figure 3.12 are very similar to the histograms of Figure 3.11b, and quantitatively, the performance metrics in Table 3.3 mimic those for the $n = 1000$ case in Table 3.2. In other words, the quality of the estimate of p_0 obtained from sampling one subclone of size 1000 is comparable to the one obtained from sampling five subclones of size 200. In this scenario, it may make more sense to extract multiple small subclones than one large one, since it is impossible to tell from a single subclone sample alone whether the tumor as a whole can be considered as evolving

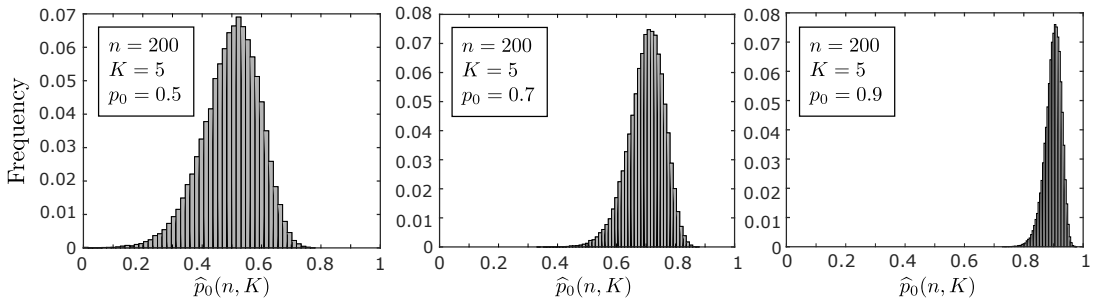

Figure 3.12: Histogram of the estimator $\widehat{p}_0(n, K)$ of (3.36) computed across 10^5 synthetic samples of $K = 5$ subclones of size $n = 200$, given true values $p_0 \in \{0.5, 0.7, 0.9\}$ and $w = 1$. Note the similarity between these histograms and the histograms in (b) of Figure 3.11, which indicates that sampling five subclones of size 200 gives a comparable estimate of p_0 to sampling a single subclone of size 1000 in these examples.

neutrally. Should there be differences in the subclone dynamics, a multiregion sample may tease this out, and should the dynamics be the same, the estimate one obtains for p_0 will be of comparable quality to the single large subclone case.

3.6 Discussion

In this work, we have established exact expressions for the expected site frequency spectrum of a tumor, or more generally any population, that evolves according to a branching process with neutral mutations under the infinite-sites assumption of population genetics. We first considered the skeleton subpopulation, consisting of cells with an infinite line of descent, and obtained explicit expressions for the SFS of the skeleton evaluated both at a fixed time and a fixed size. We then examined the total population, deriving an explicit expression for the fixed-time spectrum and a computational

K	n	p_0	mean	median	st. dev.	defaults to 0 (%)	defaults to 1 (%)
5	200	0.5	0.4965	0.5050	0.0957	0.00013	0
		0.7	0.7000	0.7055	0.0612	0	0
		0.9	0.8983	0.9014	0.0279	0	0

Table 3.3: Performance metrics for the the estimator $\widehat{p}_0(n, K)$ computed from the data that underlies Figure 3.12. Note the similarity between these metrics and the lower half of Table 3.2.

expression for the fixed-size spectrum. Our results apply to mutations at small and large frequencies, to tumor tissue samples and tumor subclones of any size, and to all values of the extinction probability p_0, even values as large as $p_0 = 0.90$ and above, which are broadly relevant for cancer. We now discuss in detail how our results relate to results previously obtained in the literature.

We begin by stating skeleton results established by Durrett [50, 53], Bozic et al. [16] and Williams et al. [160]. The result in Bozic et al. [16] says that for fixed $0 < f < 1$, the expected number of *subclonal* mutations found in a *proportion* $\geq f$ of cells at time t is, as $t \to \infty$,

$$(w/q_0)(1/f - 1). \tag{3.37}$$

Note that this is a *cumulative* spectrum of mutations with frequency at least f. A similar result appears in Williams et al. [160] under a deterministic growth model. Durrett's result [50, 53], which preceded the other two, is given under continuous mutation accumulation, and it includes clonal mutations, which by (B.15) of Appendix B.3 requires adding ν/λ_0 $(= w/q_0)$ mutations to (3.37). This yields $(\nu/\lambda_0)(1/f)$ as the cumulative spectrum, see Theorem 1 of Durrett [53] and Theorem 2 of Durrett [50].[3] Under discrete mutation accumulation, the number of clonal mutations is wp_0/q_0 by (3.14) of Proposition 3.1, which yields the result $(w/q_0)(1/f) - w$ including clonal mutations. The difference of w reflects the difference in the number of clonal mutations between the discrete and continuous model, see Appendix B.3.

To compare the $1/f$ law in (3.37) with our results of Proposition 3.1, recall that by (3.14), the fixed-size skeleton spectrum can be written as $\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] = (w/q_0)N \cdot 1/(j(j+1)) - (wp_0/q_0)\delta_{1,j}$ for $j = 1, \ldots, N-1$. Under this spectrum, it is easy to compute the expected number of subclonal mutations found in a proportion $\geq f$ of skeleton cells as

$$\begin{aligned}
\mathbb{E}\big[\textstyle\sum_{j=\lceil Nf \rceil}^{N-1} \tilde{S}_j(\tilde{\tau}_N)\big] &= (w/q_0)N \cdot \big(\textstyle\sum_{j=\lceil Nf \rceil}^{N-1} 1/(j(j+1))\big) - (wp_0/q_0)\delta_{1,\lceil Nf \rceil} \\
&= (w/q_0)N \cdot \big(1/\lceil Nf \rceil - 1/N\big) - (wp_0/q_0)\delta_{1,\lceil Nf \rceil} \\
&\sim (w/q_0)(1/f - 1), \quad N \to \infty.
\end{aligned} \tag{3.38}$$

[3] There is an apparent typo in Theorem 1 of Durrett [53]. The result is written as $(\nu/q_0)(1/f)$ in our notation, but should be $(\nu/\lambda_0)(1/f)$. This is corrected in Theorem 2 of the later notes Durrett [50] by the same author.

The $1/f$ law in (3.37) can therefore be viewed as a cumulative version of our $1/j^2$ law. Note that (3.37) is established in the fixed-time regime in the above cited works, while the calculations in (3.38) show that the $1/f$ law also holds in the fixed-size regime.

To summarize, in the fixed-size regime, we have established exact adherence to the $1/j^2$ law on $j = 2, \ldots, N-1$ for any skeleton size $N \geq 2$, and our result implies the cumulative $1/f$ law as $N \to \infty$. In the fixed-time regime, we have established the exact expression (3.12), which converges to the $1/j^2$ law for $j \ll N$ as $N \to \infty$. The $1/f$ result in (3.37) complements our fixed-time result, since it confirms that if we compute the spectrum of mutations found in a certain *proportion* of cells, rather than in a certain *number* of cells, the fixed-time spectrum converges to the same $1/f$ law as the fixed-size spectrum as $N \to \infty$, including at the very largest frequencies. It should be emphasized that the $1/f$ law, which has been extensively cited in the literature, is an asymptotic result established for the skeleton. For the total population, including finite-family cells, the $1/f$ law is applicable only to mutations at large frequencies in a large tumor. This is particularly the case for low-viability tumors.

The expressions (3.20) and (3.21) for the expected fixed-time spectrum of Proposition 3.2 have been previously obtained by Ohtsuki and Innan [124] under the assumption of deterministic growth of the tumor bulk and stochastic growth of mutant subclones, which is a reasonable approximation when p_0 is small. We have established these expressions for the fully stochastic model and for all values of p_0, by evaluating the spectrum at the fixed time t_N at which a surviving tumor has expected size N, defined in (3.18), and conditioning on the survival event $\{Z_0(t_N) > 0\}$. We have also conjectured the laws of large numbers (3.32) and (3.33) for the fixed-time and fixed-size spectrum, supported by heuristic calculations and simulations. We have proved a law of large numbers (B.19) for the semideterministic model of Ohtsuki and Innan [124], see Appendix B.8, and we plan to prove the fully stochastic results (3.32) and (3.33) in a future work. Lambert [101] has proved a similar result in the context of a coalescent point process (CPP), a framework under which an extant population is endowed with a coalescent structure that specifies how lineages coalesce when traced backwards in time. Lambert's result, see his Theorem 2.3, deals with a ranked sample of individuals from a CPP in the large-sample limit, and it has the same form as our conjectured fixed-size law of large numbers (3.33), the latter applying to our forwards-in-time branching process stopped

at a certain size. The expected fixed-size spectrum (3.22) of Proposition 3.2 is new as far as we know, as well as expressions (3.26) and (3.27) of Proposition 3.3 for the expected total mutational burden of the tumor.

All of the above results hold under the infinite-sites model of population genetics. Cheek and Antal [35] have recently examined the SFS of an exponentially growing population without this assumption, citing single-cell sequencing results of Kuipers et al. [100] as motivation. They observe that if recurrent mutations are allowed (but no back mutations), and there are S sites in the genome, the expected SFS at time t can be computed as $S \cdot \mathbb{P}(Y(t) = j)$, where $Y(t)$ is the number of mutants at time t in a two-type model of wild-type and mutant cells, each growing at the same rate. Then, to compute the SFS, one needs the distribution of $Y(t)$, which has been obtained under various assumptions e.g. by Antal and Krapivsky [3], Kessler and Levine [89, 90], and Keller and Antal [88]. Cheek and Antal obtain SFS results under limits of large time/size and small mutation rate, and their results obey $1/j^2$ power laws at large frequencies. However, when the mutation rate is sufficiently large compared to the population size, their small-frequency behavior diverges from ours, see e.g. their Figure 1. Yet other authors have substituted the infinite-sites model with the *infinite-alleles* model, under which each new mutation creates a new type of individual, see e.g. Champagnat et al. [31], Griffiths and Pakes [74], Pakes [125], Wu and Kimmel [166]. Under this model, the site frequency spectrum is usually replaced by an *allele frequency spectrum*, which tracks frequencies of genetically distinct individuals, known as *haplotypes* [101].

Our complete theoretical results give rise to several important insights. First of all, whereas the fixed-time and fixed-size skeleton spectrum depends on the mutation rate w and the extinction probability p_0 only through the effective mutation rate w/q_0, the two parameters decouple in the total population spectrum. The mutation rate w scales the spectrum linearly, whereas the extinction probability p_0 changes its shape at the small-frequency end. In fact, as p_0 increases from 0 to 1, the small-frequency end of the spectrum transitions from the $1/j^2$ power law characteristic of pure-birth exponential growth to the $1/j$ law characteristic of constant-sized populations. We examined the simple metrics $\overline{M}_1/\widehat{M}_1$ and $\overline{S}_1/\overline{M}_1$ that quantify this transition, where $\widehat{M}_1$ is the expected total mutational burden under the $1/j^2$ power law spectrum. We saw that $\overline{M}_1/\widehat{M}_1 \to 0$ as $p_0 \to 1$, which suggests that the simple estimate $\widehat{M}_1$ of the

total number of mutations, applied e.g. in Ling et al. [108], is a significant overestimate of the actual expected number of mutations $\overline{M}_1$ when p_0 is large. We finally used the metric $\overline{S}_1/\overline{M}_1$ to propose a simple estimator for p_0, based on sampling one or more spatially separated subclones from a tumor. This estimator accurately recovers the true value of p_0 from synthetic single-cell sequencing data, and it is most accurate when p_0 is large.

Our proposed estimator is currently of more theoretical than practical significance. It assumes that complete subclones of a given size can be reliably extracted from a tumor sample, and that each cell in each subclone can be single-cell sequenced so that all of its mutations are captured. We have proposed the estimator mainly to emphasize the information contained in the small-frequency end of the spectrum, and to show how it can in principle be used to decouple w and p_0 using the SFS alone. Whether and how this decoupling can be achieved under current and foreseeable limitations of genomic data warrants further investigation. For example, one can derive an estimator based on s_j/m_j for $j > 1$, where $m_j := \sum_{k=j}^{n-1} s_k$, which excludes mutations found at the smallest frequencies. One can also design a more elaborate estimating procedure, which retains some of the smallest frequencies, but explicitly models the sequencing error. As it becomes easier to distinguish small-frequency mutations from sequencing errors, our simple estimator may at the very least provide quick and easy identification of low-viability tumors, and complement other more involved techniques.

We finally note that the estimation of w and p_0 relies on several assumptions on tumor evolution. First of all, our model assumes that all mutations are selectively neutral. While this assumption is likely reasonable for smaller tumor subclones, it may be hard to verify that a subclone estimate is representative for the tumor as a whole. Our model also assumes exponential growth throughout tumor evolution, whereas e.g. due to spatial constraints and nutrient availability, the growth may be subexponential both during the early and late stages. Finally, our model assumes that the mutation rate is constant over time, and that the infinite-sites assumption holds for the tumor sample being analyzed. We note that the SFS-based estimates of w/q_0 obtained by Williams et al. [160] and Bozic et al. [16] are derived from the large-frequency end of the SFS, which reflects early tumor dynamics. The same is true of the Werner et al. [158] method of decoupling w and p_0 by tracing genealogies of spatially separated bulk samples,

since the common ancestors of these samples are likely to have existed early in tumor evolution. Conversely, our suggested approach of utilizing the small-frequency end of the SFS will more reflect late tumor evolution. On the one hand, in the evolving discussion of tumor evolutionary history inference, it is important to acknowledge that any given estimation procedure may only give a temporally or spatially constrained picture of the dynamics. On the other hand, utilizing different parts of the SFS, or combining SFS estimates with other estimates, may allow one to glean insights into the dynamics at different stages of tumor evolution, and to possibly assess the validity of any modeling assumptions.

Chapter 4

Understanding the role of phenotypic switching in cancer drug resistance

4.1 Introduction

While cancer has traditionally been considered a genetic disease driven by Darwinian evolution at the somatic level, it is now increasingly recognized that non-genetic sources of phenotypic variation may play an important role in tumor initiation, tumor progression and the evolution of drug resistance [21, 22, 56, 62, 84, 109, 112, 163]. Common sources of non-genetic heterogeneity include DNA methylation, histone modifications and other epigenetic mechanisms that alter gene expression, without changing the genetic code, by controlling DNA accessibility during transcription, replication and repair [165]. Since these mechanisms frequently operate at a significantly faster rate than genetic mutations, they can serve as a substrate for natural selection and permanently influence tumor evolution in the absence of any genetic events [21, 22, 84]. Another common source of variation in gene expression is the inherent stochasticity of intracellular biochemical reactions, which includes transcriptional noise. This stochasticity may give rise to heritable expression levels, albeit on the short time scale of one to a few cell divisions, whereas retention of epigenetic modifications has been estimated on the order

of 10-10^5 cell divisions [37, 119, 140]. As we will find, even such short-term phenotypic states can dramatically impact the course of tumor evolution.

Here, we are primarily interested in the role of non-genetic mechanisms in conferring acquired resistance to anti-cancer treatment, as has been explored in several recent experimental works. In Sharma et al. [139], for example, the authors investigate the acute response of several cancer cell lines[1] to targeted anti-cancer agents, and they consistently observe the emergence of a drug-tolerant phenotype (DTP) that is 'transiently acquired and relinquished by individual cells within the population at a low frequency'. The authors draw an analogy between DTP's and slowly-proliferating antibiotic-tolerant 'persisters' commonly observed in microbial populations [8, 43, 139], whose survival within a more rapidly proliferating population represents an evolutionary means of 'bet-hedging' against potential environmental stresses [68, 83, 115, 118]. Liau et al. [106] and Roesch et al. [133, 134] similarly describe slow-cycling DTP's in glioblastoma and melanoma, respectively, and Knoechel et al. [95] identify a reversible drug-tolerant state in leukemia which appears to be epigenetically mediated. For even further examples of experimental studies describing (often stem-like) non-genetic phenotypces associated with tumorigenic potential or drug resistance in cancer, we refer to e.g. [29, 33, 127, 137, 145], as well as recent reviews by Reyes and Lahav [131] and Salgia and Kulkarni [136]. We now turn our attention to such studies that also incorporate a modeling component.

In Gupta et al. [77], the authors employ a mathematical model of stochastic switching between three cell types to infer the rates at which breast cancer cells transition between a selectively resistant stem-like state and two non-stem-like states. Su et al. [144] show that phenotypic switching between a drug-sensitive and drug-resistant state in melanoma is well-captured by a similar model, and their study reveals the critical role played by drug-induced adoption of the resistant state, relative to selection of preexisting cells in this state. Goldman et al. [72] further provide evidence of chemotherapy-induced switching to a resistant $CD44^{Hi}CD24^{Hi}$ expression status in breast cancer, and Pisco et al. [126] show that vincristine resistance in leukemia, mediated by overexpression of the multidrug resistance protein 1 (MDR1), is primarily due to therapy-accelerated adoption of the overexpressed state. Thus, while reversible drug-tolerant cells may

[1] Non-small cell lung cancer, melanoma, colorectal cancer, breast cancer and gastric cancer.

arise naturally in drug-naïve cell populations, their emergence can also be directly influenced by anti-cancer treatment. Further complicating the picture, Craig et al. [38] have hypothesized that genotypically and phenotypically distinct cells can cooperate to induce the adoption of drug tolerance, although the potential mechanism behind such cooperation remains unclear.

Whereas transiently resistant cells serve the immediate function of protecting the tumor population from extinction, their ultimate role appears to be to set the stage for the evolution of more permanent resistance mechanisms, both of the genetic and epigenetic kind. In Sharma et al. [139], for example, the authors report that during prolonged exposure of *EGFR*-mutant non-small cell lung cancer (PC9) to erlotinib, a fraction of drug-tolerant cells become capable of proliferating normally in drug, and that this more aggressive phenotype reverts less readily to sensitivity once removed from drug. During even more prolonged exposure, Ramirez et al. [128] find that PC9 persister cells give rise to diverse genetic resistance mechanisms, including *de novo* adoption of the $EGFR^{\mathrm{T790M}}$ gatekeeper mutation commonly observed in the clinic (see also Hata et al. [78]). Shaffer et al. [138] report findings conceptually similar to Sharma et al. for melanoma cells treated with vemurafenib, and they further report that prolonged drug exposure induces epigenetic reprogramming of the drug-tolerant state, ultimately leading to a stable drug-resistant phenotype.

The above studies indicate that tumor cells in a wide variety of cancer types have the ability to adopt reversible drug-resistant phenotypes, which can in turn facilitate the eventual acquisition of bona fide resistance mechanisms. Such phenotypes can both preexist anti-cancer treatment and be specifically induced or accelerated by therapy. The relatively fast rate at which non-genetic phenotypes can be adopted, as compared to resistance-conferring mutations, poses a major challenge for clinical strategies. A deeper understanding of these dynamics is crucial to furthering our understanding of cancer and to informing novel therapeutic efforts. Here, we develop a mathematical model to investigate the evolutionary dynamics of a cell population that is able to employ transiently resistant states as a survival strategy. Our goal is to gain quantitative insights into resistance evolution in this setting, to highlight some of its unique characteristics, and to discuss how an understanding of these dynamics can inform the design of treatment strategies.

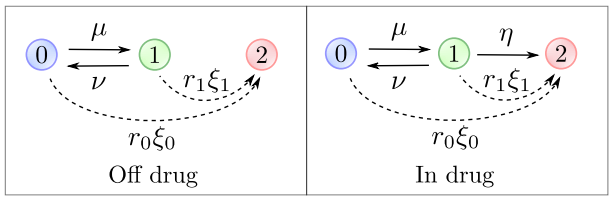

Figure 4.1: Graphical representation of the model. Type-0 (drug-sensitive) cells transition to type-1 (transiently resistant) cells at rate μ and type-1 cells transition back at rate ν. Each phenotype has distinct growth characteristics, with r_0 and r_1 denoting the rates of cell division and d_0 and d_1 denoting the rates of cell death. In the presence of an anti-cancer agent, the transiently resistant type-1 phenotype undergoes epigenetic reprogramming to stable drug resistance at rate η. Additionally, each cell type can acquire a resistance-conferring mutation at rate ξ_0 and ξ_1 per cell division, respectively, on and off drug. Stably resistant (type-2) cells divide at rate r_2 and die at rate d_2.

4.2 Model description

We consider a multi-type branching process model [6], in which cells switch stochastically between two distinct phenotypes, the drug-sensitive type-0 phenotype and the drug-resistant type-1 phenotype. A type-0 cell divides into two type-0 cells at (Poisson) rate r_0, it dies at rate d_0 and adopts the type-1 phenotype at rate $\mu > 0$, with $\lambda_0 := r_0 - d_0$ the net birth rate[2]. A type-1 cell divides at rate r_1, dies at rate d_1 and reverts to type-0 at rate $\nu > 0$, with $\lambda_1 := r_1 - d_1$ the net birth rate (Fig 4.1). To capture the drug-sensitivity of type-0 cells and drug-resistance of type-1 cells, we assume that $\lambda_0 < 0$ and $\lambda_1 > 0$ in the presence of an anti-cancer agent. Although this model assumes that switching between phenotypes can occur at any time during the cell cycle, it can easily be adjusted to allow switching to only occur at cell divisions (Appendix C.5).

These general two-type switching dynamics enable description of a variety of sources of phenotypic heterogeneity; for example, short-term drug-tolerant states conferred by transcriptional noise can be captured by a large reversion rate ν, while longer-term states induced by epigenetic phenomena are captured by a smaller ν. We also note that

[2]This means that each type-0 cell waits an exponential amount of time with rate $a_0 := r_0 + d_0 + \mu$ (i.e. the waiting time is exponentially distributed with mean $1/a_0$) before it either divides with probability r_0/a_0, dies with probability d_0/a_0, or adopts the type-1 phenotype with probability μ/a_0.

the switching rates μ and ν are in general distinct, since the mechanism underlying phenotypic switching is in general asymmetric. A simple example is DNA methylation/demethylation, where *de novo* methylation is carried out by the DNA methyltransferases DNMT3a and DNMT3b, while demethylation usually occurs due to a failure of the *maintenance* methyltransferase DNMT1 to faithfully preserve methylation patterns during cell division [57]. In what follows, we will usually refer to the transition from type-0 to type-1 as an *epimutation* and the transition back as a *reversion*, while keeping more general sources of non-genetic variation (e.g. transcriptional noise) in mind.

To capture the evolution of more permanent resistance mechanisms, we assume that under anti-cancer treatment, the type-1 phenotype undergoes epigenetic reprogramming to a stably resistant phenotype at rate η. Alternatively, type-0 and type-1 cells can acquire a resistance-conferring mutation at rate ξ_0 and ξ_1 per cell division, respectively (i.e. the mutation rate per time unit is $r_0\xi_0$ for type-0 cells and $r_1\xi_1$ for type-1 cells), both in the presence and absence of the anti-cancer agent (Fig 4.1). Stably resistant (type-2) cells divide at rate r_2 and die at rate d_2, with $\lambda_2 := r_2 - d_2 > 0$ the net birth rate. We assume throughout that no type-2 cell is present at detection, focusing on how acquired resistance develops during anti-cancer treatment.

We note that the model outlined above assumes that epigenetic reprogramming from type-1 to type-2 occurs in a single stochastic event. Single-stage reprogramming is consistent e.g. with a model of epigenetic gene silencing under recruitment of chromatin regulators suggested in a recent paper by Bintu et al. [11], and with Brown et al.'s [22] notion of 'epigenetically poised' states that evolve to fixed acquired-resistant states via DNA methylation. Conversely, a multi-stage (or more continuous) model appears more consistent e.g. with the findings of Sharma et al. [139] and Shaffer et al. [138] described above, although the data in these works are insufficient to infer an exact model. The distinction between single-stage and multi-stage reprogramming will not be important to much of our investigation, since in Sections 4.4.1-4.4.2, we focus on behavior in the absence of permanent resistance mechanisms ($\eta = 0$ and $\xi_0 = \xi_1 = 0$). In Sections 4.4.3-4.4.4, we work with the minimal single-stage model, as this allows us to gain theoretical insights into the role played by transiently resistant phenotypes in facilitating the evolution of more permanent resistance mechanisms in the simplest possible setting, in terms of a single reprogramming parameter η. Our results in these sections will

still be meaningful for the multi-stage case if we interpret the type-2 phenotype more generally as a stabler and more aggressive form of epigenetic resistance, not necessarily representing a permanent acquired-resistant state. Moreover, once multi-stage models are better understood mechanistically and quantitatively, our analysis can be easily extended to capture these more complex dynamics, using the same multi-type branching process framework as employed below.

4.2.1 Parametrization

For demonstration purposes, we adopt a baseline parameter regime chosen to mimic *in vitro* behavior of *EGFR*-mutant non-small cell lung cancer (PC9) reported by Sharma et al. [139] and Hata et al. [78]. We estimate the birth and death rates of all phenotypes and the epimutation rate μ using these works, but refer to other works for estimation of the reversion rate ν and the rates η, ξ_0 and ξ_1 of stable resistance acquisition. Details are provided in Appendix C.1. The exact parameter values are not important to our investigation, but rather the qualitative setting encoded in the regime: Type-0 cells proliferate rapidly in the absence of an anti-cancer agent and die rapidly in its presence, while type-1 cells are able to maintain slow net proliferation under the anti-cancer agent. Epigenetic reprogramming is further assumed to occur at a faster rate than resistance-conferring mutations. To ensure that our main insights are not particular to our chosen regime, but that they apply more generally across cancer cell populations capable of adopting slow-cycling, transiently resistant phenotypes, we will usually examine a range of possible switching dynamics. We furthermore discuss the sensitivity of our results to main model parameters, and perform robustness analysis where appropriate. We finally note that the rapid *in vitro* dynamics that underlie Figures 4.2 can be translated into slower *in vivo* dynamics through appropriate rescaling of time, as is discussed in Section 4.4.

4.3 Mean behavior and survival probability

We begin by deriving expressions for the average number of type-0 and type-1 cells alive at any time t, assuming the absence of permanent resistance mechanisms (i.e. $\eta = 0$ and $\xi_0 = \xi_1 = 0$). These expressions will be useful both for analyzing long-term tumor

evolution and for estimating the time at which permanent resistance first arises (Section 4.4.4).

Assume that the tumor consists of n type-0 cells and m type-1 cells at the start of anti-cancer treatment. We will both be interested in the case $m = 0$, where resistance is mediated by drug-induced adoption of the type-1 phenotype, and $m \gg 0$, where transiently resistant cells preexist treatment (in significant numbers). The switching dynamics of tumor cells are encoded in the *infinitesimal generator* for the process,

$$\mathbf{A} = \begin{bmatrix} \lambda_0 - \mu & \mu \\ \nu & \lambda_1 - \nu \end{bmatrix}, \tag{4.1}$$

where $\lambda_0 - \mu$ (resp. $\lambda_1 - \nu$) is the net rate at which a type-0 (type-1) cell produces another type-0 cell, and μ (resp. ν) is the transition rate from type-0 to type-1 (type-1 to type-0). If we let $\phi_0(t)$ (resp. $\phi_1(t)$) denote the mean number of type-0 (type-1) cells alive at time t, we can calculate these means as

$$[\phi_0(t) \;\; \phi_1(t)] = [n \;\; m] \;\; \exp(\mathbf{A}t), \tag{4.2}$$

which allows use to derive the following expressions:

$$\begin{aligned} \phi_0(t) &= \frac{n\delta + m}{\delta + \gamma} e^{\sigma t} + \frac{n\gamma - m}{\delta + \gamma} e^{\rho t}, \\ \phi_1(t) &= \frac{\gamma(n\delta + m)}{\delta + \gamma} e^{\sigma t} - \frac{\delta(n\gamma - m)}{\delta + \gamma} e^{\rho t}. \end{aligned} \tag{4.3}$$

Details of the derivation are provided in Appendix C.2. The rate constants σ and ρ, with $\sigma > \rho$, are the (real) eigenvalues of the infinitesimal generator $\mathbf{A}$, given by

$$\frac{(\lambda_0 - \mu) + (\lambda_1 - \nu) \pm \sqrt{((\lambda_0 - \mu) - (\lambda_1 - \nu))^2 + 4\mu\nu}}{2}. \tag{4.4}$$

In addition,

$$\gamma := (\sigma - (\lambda_0 - \mu))/\nu > 0 \tag{4.5}$$

is the long-run ratio between type-1 and type-0 cells in the tumor population, and

$$\delta := ((\lambda_0 - \mu) - \rho)/\nu > 0 \tag{4.6}$$

is the long-run ratio between the size of a clone derived from a single type-0 vs. a single type-1 cell. Note that the mean behavior of the process can either be expressed as a function of the fundamental constants λ_0, λ_1, μ and ν, which capture single-cell level dynamics, or as a function of the derived quantities γ, δ, σ and ρ, which capture long-run population-level behavior and may be more easily observable in an experimental setting. Also note that the above expressions only depend on the birth and death rates of each cell type through their net birth rates $\lambda_0 = r_0 - d_0$ and $\lambda_1 = r_1 - d_1$.

As expression (4.3) indicates, the long-run behavior of the tumor population is determined by the sign of the rate constant σ. The population survives with positive probability if and only if $\sigma > 0$, in which case it is said to be *supercritical*, while extinction is guaranteed for $\sigma < 0$, in which case it is *subcritical* (see e.g. [6] for further information). For a supercritical population, the survival probability can be computed by solving a system of two nonlinear equations, as is outlined in Appendix C.3.

4.4 Results

4.4.1 Resistance driven solely by phenotypic switching

We begin by investigating whether phenotypic switching can drive long-term resistance to continuous anti-cancer treatment, even in the absence of permanent resistance mechanisms (i.e. $\eta = 0$ and $\xi_0 = \xi_1 = 0$). By examining when $\sigma > 0$, we can show that tumor survival is possible (i.e. the tumor cell population is supercritical) if and only if

$$\nu\lambda_0 + \mu\lambda_1 > \lambda_0\lambda_1 \tag{4.7}$$

(see Appendix C.4 for details). By rewriting this condition as (recall that $\lambda_0 < 0$ under anti-cancer therapy)

$$\nu/\lambda_1 - \mu/|\lambda_0| < 1, \tag{4.8}$$

we see that the rates ν/λ_1 and $\mu/|\lambda_0|$ of phenotypic switching, relative to net growth (or net decay) of each phenotype, determine whether the tumor can persist under therapy. It is furthermore easy to see that a sufficient condition for (4.8) is

$$\nu/\lambda_1 \leq 1 \quad \text{or} \quad 1/\nu \geq 1/\lambda_1, \tag{4.9}$$

which can be interpreted as a simple condition on the time scale of 'phenotypic memory': The tumor population has a chance of surviving treatment whenever the average memory $1/\nu$ of the resistant state equals or exceeds $1/\lambda_1$.

If we assume that the growth characteristics of type-0 and type-1 cells (λ_0 and λ_1) are fixed, the condition (4.8) for supercriticality can be viewed as describing the region in the (μ, ν) plane that lies below the 'critical line'

$$\nu = \lambda_1/|\lambda_0| \cdot \mu + \lambda_1. \tag{4.10}$$

Since μ and ν are small, and they may differ by orders of magnitude, it is more instructive to visualize their relationship on a logarithmic scale. In Figure 4.2a, we show this log-scale 'critical curve' for two cases: (i) $\lambda_1 \ll |\lambda_0|$, i.e. type-0 cells die rapidly and type-1 cells proliferate slowly under anti-cancer treatment, as in our baseline parameter regime, and (ii) $\lambda_1 \sim |\lambda_0|$, i.e. the type-1 net birth rate is of the same order as the type-0 net death rate. In both cases, the tumor population is supercritical for $\nu \leq \lambda_1$ (region A), which is our sufficient condition from (4.9), but the population can still be supercritical for $\nu > \lambda_1$ (region B for $\lambda_1 \ll |\lambda_0|$ and regions B+C for $\lambda_1 \sim |\lambda_0|$), the degree to which is controlled by the slope $\lambda_1/|\lambda_0|$ in (4.10).

We conclude from the above that phenotypic switching can in fact drive long-term drug resistance in the absence of more permanent resistance mechanisms, and we identify two qualitatively distinct evolutionary pathways to such resistance:

1. $\nu \leq \lambda_1$: Population survival is driven by sufficiently long retention of the resistant phenotype, independently of the rate of epimutation (μ) and type-0 sensitivity to the anti-cancer agent (λ_0).

2. $\nu > \lambda_1$ and $\mu > \lambda_0(1-\nu/\lambda_1)$: Type-1 cells lose the resistant phenotype too quickly to sustain the tumor by themselves, but this loss is compensated by sufficiently

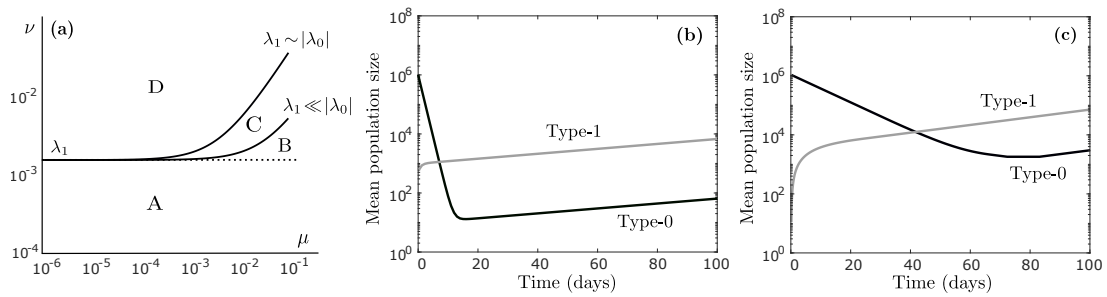

Figure 4.2: **(a)** Graphical depiction of the region in the (μ, ν) plane where the tumor cell population is supercritical (the region below the line in (4.10)), displayed here on a logarithmic scale for two cases ($\lambda_1 \ll |\lambda_0|$ and $\lambda_1 \sim |\lambda_0|$). The two curves are drawn assuming $r_1 = 0.0162$ (per hour) and $d_1 = 0.015$, and $r_0 = 0.04$ and $d_0 = 0.08$ for the $\lambda_1 \ll |\lambda_0|$ case (baseline parameter regime), and $r_0 = 0.004$ and $d_0 = 0.008$ for the $\lambda_1 \sim |\lambda_0|$ case. Since the slope of (4.10) is $\lambda_1/|\lambda_0| > 0$, and it intercepts the ν-axis at $\nu = \lambda_1$, the population is supercritical whenever $\nu \leq \lambda_1$, independently of the relationship between λ_0 and λ_1 (region A). When phenotypic memory is short ($\nu > \lambda_1$), the population can still be supercritical, the degree to which depends on the slope $\lambda_1/|\lambda_0|$ (region B for $\lambda_1 \ll |\lambda_0|$ and regions B+C for $\lambda_1 \sim |\lambda_0|$). **(b)** Time-evolution of the expected number (log-scale) of type-0 cells (dark curve) and type-1 cells (light curve) during continuous anti-cancer treatment, assuming no transiently resistant cell is present at detection ($m = 0$), calculated using (4.3). Parameter values are $r_0 = 0.04$ (per hour), $d_0 = 0.08$, $r_1 = 0.0162$, $d_1 = 0.015$, $\mu = 4 \cdot 10^{-5}$, $\nu = 4 \cdot 10^{-4}$, $\eta = 0$, $\xi_0 = \xi_1 = 0$, $n = 10^6$ and $m = 0$. **(c)** Same as (**b**), except now $r_0 = 0.004$ (per hour) and $d_0 = 0.008$.

fast adoption of the resistant state.

If we assume $d_1 = 0$, i.e. type-1 cells have a stem-like ability to proliferate indefinitely, the first condition implies that even single-generation phenotypic memory may be sufficient to confer long-term resistance, while the second condition implies that even non-heritable traits, e.g. stochastic variation in gene transcription, may be able to save the tumor from extinction.

To further elucidate the coevolutionary dynamics between type-0 and type-1 cells, we show in Figure 4.2b the long-term expected behavior of the population in the baseline parameter regime, assuming no transiently resistant cell is present at detection ($m = 0$). When the anti-cancer agent is applied, sensitive type-0 cells initially die at a fast rate, while a small fraction of them adopts the resistant type-1 phenotype. Eventually, the population settles into an equilibrium where back-and-forth epimutations of type-0 cells

and reversions of type-1 cells drive an expansion of both subpopulations, albeit at a slow rate. Once the type-1 population has reached a sufficient size, we can expect it to develop more permanent resistance mechanisms, as is discussed further in subsequent sections. In Figure 4.2c, we show long-term expected behavior under an alternative parameter regime, where $|\lambda_0|$ is reduced so that it is of the same order as λ_1 ($\lambda_1 \sim |\lambda_0|$). In this regime, type-0 decay is less rapid under anti-cancer treatment, which implies both that the type-1 population builds up more quickly, and that more type-0 cells remain at equilibrium. Note that the equilibrium proportion between type-1 and type-0 cells can in general be computed using expression (4.5) above.

By examining cell behavior at the individual level, we observe that each type-0 cell is almost guaranteed to go extinct in the baseline regime (it survives with probability 0.005%), while each type-1 cell survives with 4.9% probability (Equation C.9). Despite these odds, the anti-cancer agent is unsuccessful in eradicating the drug-sensitive population due to the dynamic switching between phenotypes. In fact, the population as a whole is guaranteed to survive treatment (Equation C.10), which is due to the substantial buffer of type-1 cells that accumulates through type-0 epimutations at the start of treatment and protects the somewhat fragile type-1 population against extinction. These survival probabilities are generally robust to significant variation in r_0, d_0, r_1 and d_1, as we display in Table C.3 in Appendix C.8. The main exception occurs when r_1 and d_1 are changed so that the population becomes subcritical, in which case there is no chance of tumor survival. This transition from guaranteed tumor survival to guaranteed extinction can be quite abrupt, as we explore further in Section 4.4.2.

The above example allows us to glean two important insights: First of all, the relatively rapid adoption of an epigenetically-mediated resistant phenotype places less restrictions on the robustness of such a phenotype than what is the case for a rare genetic variant arising through mutation. Thus, even a barely viable non-genetic phenotype may allow the tumor population to escape anti-cancer therapy with 100% probability, in the absence of any more permanent resistance mechanisms. Secondly, a tumor population that appears to be static or slow-growing at the population level may in fact be driven be rapid switching dynamics at the single cell level, and uncovering the exact dynamics may be crucial to understanding how the population responds to treatment, which is the subject of our next section.

4.4.2 Tumor survival when switching dynamics are perturbed

Targeted epigenetic agents, e.g. inhibitors of DNA methylation and histone deacetylation, have been considered both as a means of reversing the tumorigenic potential of cancer cells and of resensitizing resistant cells to anti-cancer therapy (see e.g. [10, 61, 73, 87, 114]). In this section, we examine how the probability of tumor survival under continuous anti-cancer treatment depends on the rate of epimutation (μ) and reversion (ν). We then extract insights into the potential benefits of a joint application of an anti-cancer agent, aimed at killing the tumor bulk, and an epigenetic drug, aimed at disrupting the phenotypic switching dynamics.

The probability of survival of the tumor cell population, derived in Appendix C.3, is shown in Figure 4.3a as a function of μ and ν. We consider first the case where no transiently resistant cell preexists treatment ($m = 0$) and permanent resistance mechanisms are absent ($\eta = 0$ and $\xi_0 = \xi_1 = 0$). We observe transitions in the dynamics around threshold values of $\mu' \sim 10^{-6}$ and $\nu' \sim 10^{-3}$ per hour. The latter threshold reflects a regime change from supercriticality to subcriticality, since the baseline value for the net birth rate λ_1 is of order 10^{-3} per hour (the population is supercritical below the 'critical curve' in Fig 4.3a; see Section 4.4.1). When the population is subcritical, the tumor goes extinct with 100% probability, since the high reversion rate ν to sensitivity makes it impossible for type-1 cells to expand under treatment. In the supercritical regime, there is always some positive probability that the tumor survives, although this probability will be small for low epimutation rates ($\mu \ll 10^{-6}$). For example, the tumor survival probability corresponding to $\mu = 10^{-10}$ and $\nu = 10^{-4}$ in Figure 4.3a is 0.017%, since in this case, epimutations are so infrequent that the type-1 state is unlikely to be adopted by any type-0 cell before the population goes extinct. For high epimutation rates ($\mu \gg 10^{-6}$), however, the type-1 buffer that accumulates at the start of treatment becomes so large that tumor survival is guaranteed whenever the population is supercritical, while extinction is guaranteed whenever the population is subcritical.

The threshold value $\mu' \sim 10^{-6}$ per hour represents the minimal epimutation rate at which tumor survival is certain, given a supercritical population. In Appendix C.6, we show that the epimutation rate at which the survival probability is at least $1 - u$

$(u \ll 1)$ is

$$\mu' \approx \frac{|\lambda_0| \log u}{n \log(d_1/r_1)}. \tag{4.11}$$

The threshold value μ' thus depends on the drug-sensitivity λ_0 of type-0 cells, the robustness of the resistant phenotype (d_1/r_1 is the extinction probability of a type-1 clone, assuming no epimutations or reversions, see e.g. [50]), and the initial population size n. An order of magnitude change in λ_0 or n will result in an order of magnitude change in μ', and in Table C.4 in Appendix C.8, we show some examples of sensitivity to r_1 and d_1. For the baseline parameter regime and $u = 0.001$, expression (4.11) yields $\mu' \approx 3.59 \cdot 10^{-6}$ per hour, compared to a type-0 birth rate of $r_0 = 0.04$ per hour. This implies that if the resistant phenotype is adopted once in every 10,000 cell divisions during treatment, the tumor is guaranteed to survive, even if no resistant cell preexists in the population. For larger tumor sizes, e.g. $n = 10^8$ or $n = 10^{10}$ cells, the required adoption rate lowers to once every 10^6 and once every 10^8 cell divisions, respectively, which are frequencies typical of resistance-conferring mutations (Appendix C.1). Since epigenetic modifications and other non-genetic mechanisms can operate much faster, the above discussion implies that the mere possibility of non-genetically conferred resistance can all but guarantee its emergence, especially when the tumor is large at detection.

It is worth noting that the transition from guaranteed tumor survival to guaranteed extinction in Figure 4.3a is much more gradual around the threshold value $\mu' \sim 10^{-6}$ per hour on the μ-axis than around the critical value $\nu' \sim 10^{-3}$ on the ν-axis. This reflects the asymmetric role of type-0 and type-1 cells, and of epimutations and reversions, in the evolutionary dynamics. Lowering the epimutation rate μ will reduce the type-1 buffer that accumulates at the start of treatment, which gradually impairs the collective survival prospects of type-1 cells. On the other hand, since any reversion from type-1 to type-0 effectively amounts to cell death in our setting, increasing the reversion rate ν will directly affect the survival prospects of individual type-1 cells. As long as ν is smaller than the critical value, each type-1 cell has some positive probability of persisting therapy, and given a sufficiently large type-1 buffer (i.e. sufficiently high epimutation rate μ), the tumor as a whole can be guaranteed to survive. Once ν increases above the critical value, however, each individual type-1 cell becomes certain to go extinct, and

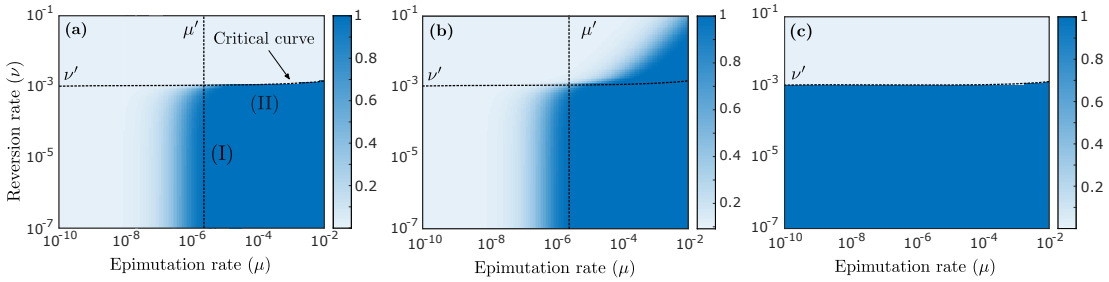

Figure 4.3: **(a)** Probability that the tumor survives continuous anti-cancer therapy, as a function of μ and ν, assuming $n = 10^6$ and $m = 0$ and the absence of permanent resistance mechanisms ($\eta = 0$ and $\xi_0 = \xi_1 = 0$), calculated using (C.10). The black dotted curve ('critical curve') separates the regions in the (μ, ν) plane where the population is subcritical (top) and supercritical (bottom). This curve can be extracted from expression (4.10) upon logarithmic scaling (see Fig 4.2a). Region (I) in the figure indicates a parameter regime where inhibiting epimutations (decreasing μ) with an epigenetic drug can be an effective treatment strategy, whereas inducing reversions (increasing ν) does little. The reverse is true in the parameter regime of region (II), where a slight perturbation to the reversion rate ν can guarantee eradication of the tumor cell population. **(b)** Same as (a), now assuming $\eta = 4 \cdot 10^{-7}$ (per hour) and $\xi_0 = \xi_1 = 10^{-7}$ (per cell division), calculated using (C.18). **(c)** Same as (a), now assuming $n = 10^6 \cdot 0.999$ and $m = 10^6 \cdot 0.001$, calculated using (C.10). Other parameter values are $r_0 = 0.04$ (per hour), $d_0 = 0.08$, $r_1 = 0.0162$, $d_1 = 0.015$, $\mu = 4 \cdot 10^{-5}$ and $\nu = 4 \cdot 10^{-4}$.

the same goes for the tumor as a whole, no matter how large the buffer is. This explains why for high epimutation rates ($\mu \gg 10^{-6}$), we observe such a sharp transition between guaranteed tumor survival and guaranteed extinction across the critical curve for ν.

Identifying where a particular cancer cell population falls within the (μ, ν) parameter space can yield important insights into the relative attractiveness of targeting μ and ν with an epigenetic drug, and the degree to which these parameters should be perturbed. When $\mu \sim 10^{-6}$ and $\nu \ll 10^{-3}$ per hour, for example (region (I) in Fig 4.3a), inhibiting epimutations (reducing μ) may significantly reduce the tumor survival probability, while inducing reversions (increasing ν) may accomplish little. When $\mu \gg 10^{-6}$ and $\nu \sim 10^{-3}$ per hour, however (region (II) in Fig 4.3a), a slight perturbation to the reversion rate may be the difference between certain tumor survival and certain extinction. This suggests that even if no resistant cell preexists treatment, it may be more effective to revert resistant cells created during the initial stages of therapy than to prevent their

emergence. We also note the importance of identifying the relationship between ν and λ_1 for therapeutic considerations. Indeed, recognizing that a small perturbation to the average retention of the resistant phenotype may yield significant treatment benefits can help minimize the risk of any unwanted side effects of the epigenetic treatment.

For the case where permanent resistance mechanisms are assumed ($\eta > 0$ and $\xi_0, \xi_1 > 0$), the probability of tumor survival can be derived by solving a system of nonlinear equations as shown in Appendix C.7. Figure 4.3b shows that in this case, the tumor can survive even if the population of type-0 and type-1 cells is subcritical. Indeed, if the epimutation rate μ is sufficiently high, the large type-1 buffer created at the start of treatment may allow even a subcritical type-1 phenotype to hold off extinction long enough for bona fide resistance to develop.

If we assume that a significant number of resistant cells preexists therapy ($m \gg 0$), the survival probability becomes insensitive to changes in μ under treatment, since the type-1 buffer needed to save the tumor from extinction will already be present at the onset (Fig 4.3c). It remains true, however, that a small perturbation to the reversion rate can turn certain therapy failure into certain success.

4.4.3 Pathway to stable resistance and rate of acquisition

We now examine the mode and speed of stable resistance acquisition ($\eta > 0$ and $\xi_0, \xi_1 > 0$) during continuous anti-cancer treatment. In Figure 4.4a, we show the average time at which the first stably resistant (type-2) cell arises in the population, as a function of the epimutation rate μ, given that a type-2 cell emerges before extinction of the population. In Figure 4.4b, we show the probability of the conditioning event, and the probability that stable resistance is conferred through mutation as opposed to epigenetic reprogramming.

We first assume that no transiently resistant (type-1) cell is present at the onset ($m = 0$; solid curves). We again note a transition in the dynamics around a threshold value of $\mu \sim 10^{-6}$ per hour, and interestingly, the average time to resistance is both non-monotonic and highly variable in μ (Fig 4.4a; solid curve). When μ is small ($\mu \ll 10^{-6}$), epimutations are so infrequent that the tumor can only survive via mutation in the type-0 population. Since type-0 cells decay at a fast rate, any such mutation has to occur early if it is to occur at all. As μ increases, the burden of saving the tumor from

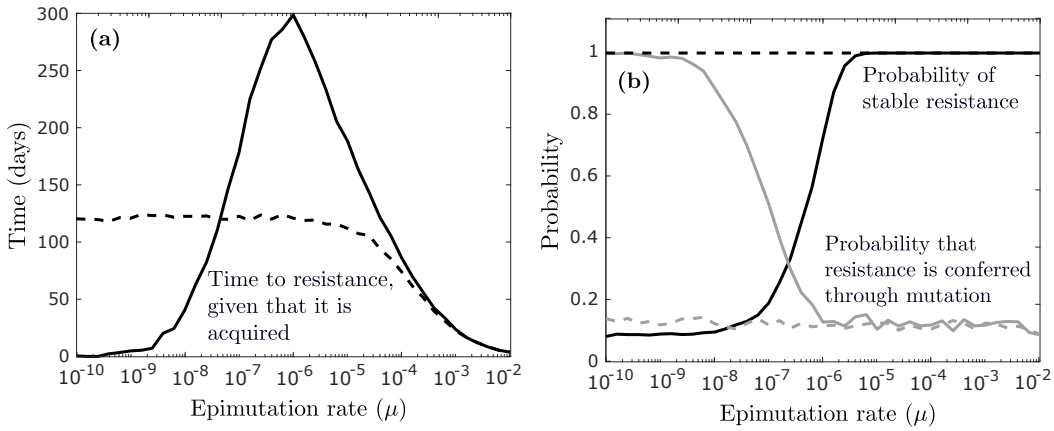

Figure 4.4: **(a)** Expected time at which the first type-2 cell emerges in the population as a function of the epimutation rate μ, conditioned on the event that a type-2 cell emerges before extinction, first assuming $n = 10^6$ and $m = 0$ (solid curve), and then $n = 10^6 \cdot 0.999$ and $m = 10^6 \cdot 0.001$ (dashed curve). Produced via simulation by running the process until a type-2 cell emerged on 1000 different occasions and calculating an average. **(b)** Probability of type-2 emergence (dark curves) and probability that the first type-2 cell arises through mutation (light curves). Other parameter values are $r_0 = 0.04$ (per hour), $d_0 = 0.08$, $r_1 = 0.0162$, $d_1 = 0.015$, $\nu = 4 \cdot 10^{-4}$, $\eta = 4 \cdot 10^{-8}$ and $\xi_0 = \xi_1 = 10^{-7}$ (per cell division).

extinction moves increasingly from type-0 to type-1 cells (Fig 4.4b; light solid curve), which creates a new pathway for resistance acquisition at a later time. Once the role of type-0 and type-1 cells in conferring resistance becomes stabilized, however, the mean acquisition time starts to decrease as μ increases, since the type-1 population is created earlier and in greater numbers. If a significant number of type-1 cells is present at the onset ($m \gg 0$), resistance will be guaranteed to form independently of the epimutation rate under treatment (Fig 4.4b; dark dashed curve), and the expected acquisition rate will be insensitive to μ for small μ (Fig 4.4a; dashed curve), since resistance will most likely develop through the type-1 population that preexists treatment (Fig 4.4b; light dashed curve).

In Figure 4.5a, we consider time to resistance as a function of the reversion rate ν instead. We again observe non-monotonicity and significant variability, with resistance development being slowest around the critical value $\nu \sim 10^{-3}$ per hour. In this case, however, the rate of resistance acquisition is generally insensitive to changes in ν for small ν, and the transition around the critical value $\nu' \sim 10^{-3}$ is much sharper than

around the threshold value $\mu' \sim 10^{-6}$ in Figures 4.4a and 4.4b. For small ν, stable resistance is mediated primarily through epigenetic reprogramming (Fig 4.5b; light solid curve), so time to resistance is governed by the size of the type-1 population. Since the net proliferation rate $\lambda_1 - \nu$ determines net growth of type-1 cells, the rate of resistance acquisition is not affected by ν as long as $\nu \ll \lambda_1$. As ν approaches λ_1, however, the net proliferation rate $\lambda_1 - \nu$ approaches zero, and epigenetic reprogramming slows down considerably. As ν increases above λ_1, the population becomes subcritical, and the probability of tumor survival decreases abruptly (Fig 4.5b; dark solid curve), similarly to what we observed in Section 4.4.2. During this sharp transition, the burden of saving the tumor from extinction moves from the type-1 to the type-0 population (Fig 4.5b; light solid curve), and since any mutation in the type-0 population must occur early if it is to occur at all, the time to resistance also decreases abruptly as ν enters the subcritical regime. Contrary to Figures 4.4a and 4.4b, we now observe the same qualitative behavior for the $m = 0$ (solid curves) and $m \gg 0$ (dashed curves) cases, which mirrors our discussion from Section 4.4.2.

The above observations add an interesting layer to our earlier discussion on the effect of manipulating μ and ν on treatment outcome. As an example, when $\mu \gg 10^{-6}$ per hour, inhibiting epimutations (reducing μ) may not do much to prevent acquired resistance, but it can slow it down considerably. Moreover, any perturbation to the reversion rate ν when it is around its critical value may dramatically affect the rate of resistance acquisition. Note that when μ is around its threshold value of $\mu \sim 10^{-6}$, resistance will be expected to arise extremely late, while still being guaranteed to develop. The former is a consequence of the slow-cycling nature of the resistant phenotype, while the latter is a consequence of its rapid adoption under treatment. This is an important feature of epigenetically-mediated resistance, and this behavior stands in stark contrast to more robust genetically-resistant phenotypes that are adopted less frequently.

4.4.4 Evaluation of combination treatment strategies

To further illustrate how our model can aid medical decision-making, we now examine a combination treatment of an anti-cancer agent and a hypothetical epigenetic drug that directly targets the rates of epimutation (μ) and reversion (ν). Our main questions are whether such a combination is likely to be effective, and whether the epigenetic drug

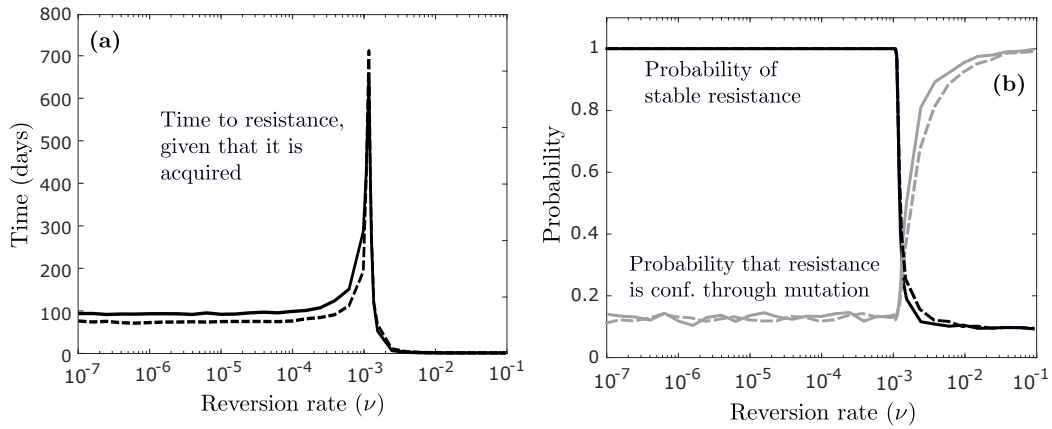

Figure 4.5: **(a)** Expected time at which the first type-2 cell emerges in the population as a function of the reversion rate ν, conditioned on the event that a type-2 cell emerges before extinction, first assuming $n = 10^6$ and $m = 0$ (solid curve), and then $n = 10^6 \cdot 0.999$ and $m = 10^6 \cdot 0.001$ (dashed curve). Produced via simulation by running the process until a type-2 cell emerged on 1000 different occasions and calculating an average. **(b)** Probability of type-2 emergence (dark curves) and probability that the first type-2 cell arises through mutation (light curves). Other parameter values are $r_0 = 0.04$ (per hour), $d_0 = 0.08$, $r_1 = 0.0162$, $d_1 = 0.015$, $\mu = 4 \cdot 10^{-5}$, $\eta = 4 \cdot 10^{-8}$ and $\xi_0 = \xi_1 = 10^{-7}$ (per cell division).

should be applied as pretreatment, posttreatment, or simultaneously with the anti-cancer agent. To evaluate treatment outcome, we use the following expression for the probability that a successful type-2 cell (i.e. a type-2 cell that gives rise to a clone that does not go extinct) has emerged by time t,

$$1 - \exp\left(-\frac{\lambda_2}{r_2}\int_0^t (\xi_0 r_0 \phi_0(s) + \xi_1 r_1 \phi_1(s) + \eta\phi_1(s))ds\right) \tag{4.12}$$

derived in Appendix C.7. Using (4.3), we can derive explicit expressions for the integrals in (4.12) as follows:

$$\begin{aligned} \int_0^t \phi_0(s)ds &= -\frac{\beta}{\sigma}(1 - e^{\sigma t}) - \frac{\alpha}{\rho}(1 - e^{\rho t}), \\ \int_0^t \phi_1(s)ds &= -\frac{\gamma\beta}{\sigma}(1 - e^{\sigma t}) + \frac{\delta\alpha}{\rho}(1 - e^{\rho t}), \end{aligned} \tag{4.13}$$

assuming $\sigma \neq 0$ and $\rho \neq 0$. These two integrals can be interpreted as the average 'total mass' of type-0 and type-1 cells, respectively, up until time t.

We assume that each treatment cycle consists of three two-day blocks and we examine four schedules. In Schedule A, the epigenetic drug is applied as pretreatment to the anti-cancer agent, whereas in Schedules B, C and D, the anti-cancer agent is applied during the first two blocks of each cycle, and the epigenetic drug is applied during the first, second or third block, respectively (Fig 4.6a). We show results assuming that the epigenetic drug increases the reversion rate ν by two orders of magnitude and decreases the epimutation rate μ to the same extent, while noting that our insights are robust to significant variation in this assumption (Appendix C.8). We also note that the relatively short duration of each treatment cycle is a function of the rapid *in vitro* dynamics of our baseline parameter regime. If each rate constant in the baseline regime is reduced by an order of magnitude, which may more accurately represent *in vivo* dynamics, the results shown below will continue to hold unchanged if we use 60-day treatment cycles instead of 6-day cycles.

Figure 4.6b indicates that whereas the anti-cancer agent alone is guaranteed to fail, combining it with a drug that disturbs the epigenetic switching dynamics has the potential to eradicate the tumor with high probability. Pretreatment with the epigenetic drug performs the worst by far (Schedule A; red solid curve), which is due to the fast proliferation of type-0 cells in the absence of the anti-cancer agent, and the relatively fast acquisition of the resistant phenotype in its presence. This both increases the risk of a resistance-conferring mutation during the initial pretreatment phase, and ensures that type-1 cells killed during this phase are rapidly replenished once the anti-cancer agent is applied (Fig 4.6c). Schedule A continues to perform the worst even if we assume that the tumor population cannot expand beyond its initial size due to spatial constraints, which we model by setting $r_0 = d_0 = 0.04$ per hour during the initial pretreatment phase (red dashed curves). This implies that pretreatment with an epigenetic drug may not be judicious when drug-sensitive type-0 cells proliferate too rapidly in the absence of the anti-cancer agent and/or when the resistant type-1 phenotype is adopted too quickly under treatment. This insight is consistent with experimental findings by Sharma et al. [139], where the authors observe that simultaneous application of erlotinib and the histone deacetylase inhibitor TSA to PC9 cells eliminates the emergence of resistant

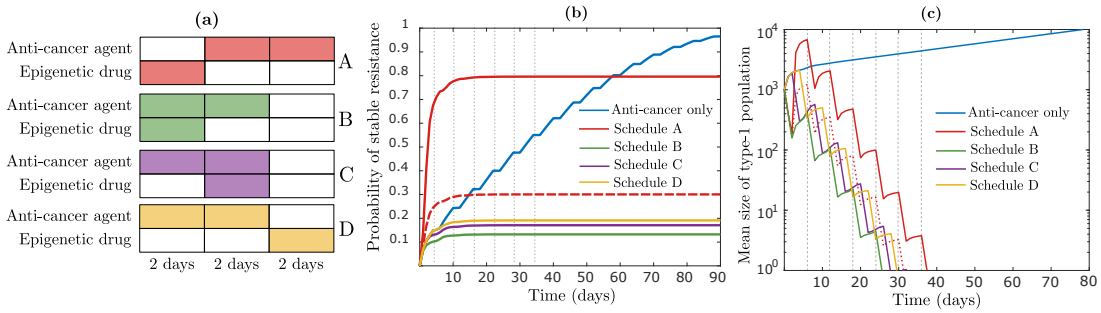

Figure 4.6: **(a)** Schematic representation of one treatment cycle in the schedules considered. Each schedule is identified by a distinct color which is also applied in (b) and (c). **(b)** Time-evolution of the probability of successful type-2 emergence, first when the anti-cancer agent is applied alone during the first two blocks of each cycle (blue curve), and then for the four schedules depicted in (a), calculated using (4.12). Under the anti-cancer agent, $r_0 = 0.04$ (per hour) and $d_0 = 0.08$, while off it, $r_0 = 0.04$ and $d_0 = 0.0015$. Under the epigenetic drug, $\mu = 4 \cdot 10^{-7}$ and $\nu = 4 \cdot 10^{-2}$, while off it, $\mu = 4 \cdot 10^{-5}$ and $\nu = 4 \cdot 10^{-4}$. The dashed red line shows a modification of Schedule A where $r_0 = d_0 = 0.04$ during the first 2-day block of the first treatment cycle. **(c)** Time-evolution of the expected number of type-1 cells for the same schedules as considered in (b), calculated using (4.3). Parameters not mentioned above are fixed at $r_1 = 0.0162$ (per hour), $d_1 = 0.015$, $\eta = 4 \cdot 10^{-7}$, $\xi_0 = \xi_1 = 10^{-7}$ (per cell division), $n = 10^6 \cdot 0.999$ and $m = 10^6 \cdot 0.001$. The vertical broken lines in (b) and (c) signify treatment cycles.

clones, whereas a substantial number of such clones arises when extended pretreatment with TSA is followed by exposure to erlotinib.

Focusing now on the other three schedules (B, C and D), we note that the number of type-1 cells remaining at the end of each treatment cycle is similar for each of them (Fig 4.6c). Schedule B performs the best, however, since it contains the type-1 population as early as possible in each cycle, which reduces the 'total mass' of type-1 cells in the presence of the anti-cancer agent and slows down epigenetic reprogramming by (4.12). This finding is robust to significant perturbation of main model parameters, and the relative attractiveness of Schedule B over the other schedules can be much more pronounced than shown here, e.g. when the rate of epimutation (μ) or epigenetic reprogramming (η) is higher, or the initial tumor size ($n+m$) is larger (Appendix C.8). As an example, when $n = 10^8 \cdot 0.999$ and $m = 10^8 \cdot 0.001$, with other parameters unchanged, resistance will be all but guaranteed to form in Schedules A, C and D, while Schedule B is capable of preventing resistance via epigenetic reprogrammming with high

probability (Fig C.2d).

Our results indicate both that combining an epigenetic drug with an anti-cancer agent can vastly outperform anti-cancer only treatment, and that the epigenetic drug should be applied simultaneously with the anti-cancer agent. These insights are consistent with experimental findings e.g. by Su et al. [144], who find that simultaneous *in vitro* inhibition of the NfκB p65 and MEK/ERK signaling axes in melanoma, in combination with the anti-cancer agent vemurafenib, significantly outperforms monotherapy, by arresting cell-state transitions to a drug-tolerant state. In Goldman et al. [72], the authors find that applying the SFK inhibitor dasatinib *in vivo* on days 8-11 of breast cancer treatment, following two maximum tolerated doses of the anti-cancer agent docetaxal (DTX) on days 2 and 5, is much more effective than DTX-alone treatment. They further show that this treatment is preferable to applying dasatinib on days 2-5 (Schedule I), which shows modest improvement over stand-alone DTX treatment, and applying dasatinib on days 14-17 (Schedule II), which shows no improvement over stand-alone treatment. Although the authors in [72] refer to Schedule I as a 'simultaneous schedule', the application of dasatinib on days 2-5 should not necessarily be expected to inhibit epimutations following the large DTX dose on day 5. Application on days 8-11 may be better timed both to revert already drug-tolerant cells to sensitivity and to "target the induction phase of DTX-induced cell state transition" as the authors put it. Nevertheless, this example serves as an important reminder that actual *in vivo* dynamics are likely to be more complex than captured by our simple abstraction. As an example, if the epimutation rate μ is influenced by anti-cancer treatment, but this effect is delayed, there may be reason to delay application of the epigenetic drug. Any substantial presence of drug-tolerant cells that preexist treatment will then create a trade-off between attacking these two temporally distinct sources of resistance. Indeed, a more complete mechanistic and quantitative understanding of the dynamics of phenotypic switching at the single-cell-level, and how these dynamics are influenced by anti-cancer treatment, will give rise to more complex mathematical questions, as we address further in the discussion section.

4.5 Discussion

The pervasiveness of acquired drug resistance remains one of the major challenges during clinical management of cancer patients. We have established through our evolutionary modeling that non-heritable stochastic fluctuations in gene expression and short-term epigenetic modifications can 'save' a tumor from extinction in the absence of any more permanent resistance mechanisms. We have also seen that the potential for rapid adoption of non-genetic resistance implies that such resistance can be all but guaranteed to develop, even if no resistant cell preexists treatment and the resistant phenotype is barely viable. This suggests ample opportunity for non-genetic mechanisms to induce failure of conventional anti-cancer therapy, and it may help explain why so many tumors develop resistance without acquiring mutations in drug targets or activated pathways [163].

Throughout, we have highlighted how thinking in terms of our mathematical model can aid medical decision-making. In Section 4.4.2, we saw how a quantitative understanding of the dynamics of epimutation and reversion at the single-cell level can yield important insights into which parameter to attack with an epigenetic drug and the extent to which it needs to be perturbed. As an example, when the average retention $1/\nu$ of the resistant phenotype is near a critical value of $1/\lambda_1$, a slight perturbation to ν can turn guaranteed treatment failure into guaranteed success. In Section 4.4.3, we further saw how a quantitative understanding of the underlying switching dynamics, and the rate at which any permanent resistance mechanisms are adopted, can allow us to predict the mode and time scale of resistance acquisition. We noted that when the resistant type-1 phenotype is slow-cycling, the expected time at which permanent resistance arises in the population can vary greatly depending on the epimutation rate μ, and that when μ is near a certain threshold value, resistance can be expected to develop extremely late while still being guaranteed to emerge.

In Section 4.4.4, we finally observed that combining an epigenetic drug with an anti-cancer agent can significantly outperform anti-cancer only treatment, and that the epigenetic drug should be applied simultaneously with the anti-cancer agent rather than as pretreatment or posttreatment. We also noted that pretreatment with the epigenetic drug is not advisable especially when drug-sensitive cells proliferate too rapidly in the

absence of the anti-cancer agent, or the resistant phenotype is adopted too quickly in its presence. It should of course be stressed that our model is highly simplified. As an example, we have not assumed any delay in the potential drug-induced adoption of transient resistance under anti-cancer treatment, and we have not assumed any interaction between the anti-cancer agent and the epigenetic drug. Also, whereas we have assumed that the drug-resistant phenotype is continuously slow-cycling, it may respond to drug pressure by entering a quiescent state before resuming proliferation. Finally, our model does not incorporate spatial effects or toxicity constraints, all of which may call for modifications or extensions to the model as our biological understanding accumulates.

Our analysis represents, as far as we know, a first attempt toward a deeper understanding of the evolutionary dynamics of a population that is able to employ transiently resistant cell states to escape anti-cancer therapy. The model presented is flexible in that it enables description of a variety of sources of phenotypic heterogeneity, and the analytical expressions we have derived allow for easy estimation and comparison of treatment outcomes under different regimens. Our investigation is theoretical in nature, however, and the true power of our mathematical model will only be realized through a more complete mechanistic understanding of the dynamics of phenotypic switching at the single-cell-level, and through new methods to infer these dynamics quantitatively. Our hope is that a combined biological and mathematical effort will pave the way toward the design of novel therapeutic strategies, as well as mathematical optimization of already available drug combinations and treatment schedules [36, 103].

Chapter 5

Statistical inference of the rates of cell proliferation and phenotypic switching in cancer

5.1 Introduction

Cancer evolution has long been understood to be a genetic process. However, recent evidence suggests an equally important role for non-genetic forces, including epigenetic mechanisms and the inherent stochasticity in gene transcription and translation [12, 21, 22, 62, 84, 136]. These mechanisms are heritable and reversible, and they can enable cells to dynamically switch between two or more phenotypic states. Such switching dynamics have been observed e.g. in lung cancer [78, 128, 139], melanoma [133, 138, 144], glioblastoma [106, 117], leukemia [95, 126], colon cancer [59, 70, 156, 168] and breast cancer [9, 72, 77, 86]. The different phenotypes commonly show differential responses to drug treatments, which enhances the adaptability of the cancer under treatment and significantly increases the probability of treatment resistance [75].

Unraveling how the cancer-specific rates of cell division, cell death and phenotypic switching shape tumor evolution over time is crucial to furthering our understanding of the disease and to informing new treatment strategies. For example, in a two-phenotype cancer where one type is drug-sensitive and the other is drug-tolerant, the change in

phenotypic proportions during the initial stages of treatment can be explained by a combination of sensitive cells dying, drug tolerant cells proliferating, and cells switching between sensitivity and tolerance. Disentangling the relative rates at which these events occur can help us to better understand how resistance arises, how it evolves over time, and how best to combat it [75].

Our current quantitative understanding of the rates of cell proliferation and phenotypic switching in cancer is largely derived from cell line experiments. In these experiments, live cells are commonly sorted into phenotypes, e.g. based on their gene expression profiles or morphologies, isolated subpopulations are expanded in culture, and phenotypic proportions are tracked over time (Fig. 5.1). These isolated subpopulations have often been observed to give rise to all other phenotypes over time, and to eventually reconstitute a stable distribution of phenotypes that is characteristic of the parental population [9, 70, 77, 86, 117, 156, 168].

To explain this behavior, simple mathematical models of phenotypic switching have been proposed, and these models have been used to estimate the rates at which cells switch between states [24, 42, 72, 77, 82, 105, 144, 171]. These works are reviewed in Appendix D.1. Previous estimation methods have been deterministic in nature, and they have generally derived their estimates from data on the fraction of cells in each state at each time point. If the total size of the cell population is measured at the same time points, as e.g. in [42], one obtains data on the number of cells in each state at each time point. We will show that when cell fraction data is used, the rates of phenotypic switching may be the only parameters that can be estimated accurately. In contrast, using cell number data enables accurate estimation of the net cell division rate for each phenotype, and it can even enable estimation of the state-dependent rates of cell division and cell death. Understanding how growth rates vary between types is as important as understanding the rates of phenotypic switching, especially in the context of treatment response. Not only do the growth rates influence the phenotypic composition of the population, they also control the evolution of the tumor burden over time.

Our goal in this work is to develop a statistically rigorous framework for estimating the rates of cell proliferation and phenotypic switching in cancer. In contrast to previous approaches, our framework explicitly models the stochastic dynamics of cell division, cell death and phenotypic switching, it provides likelihood-based confidence intervals for

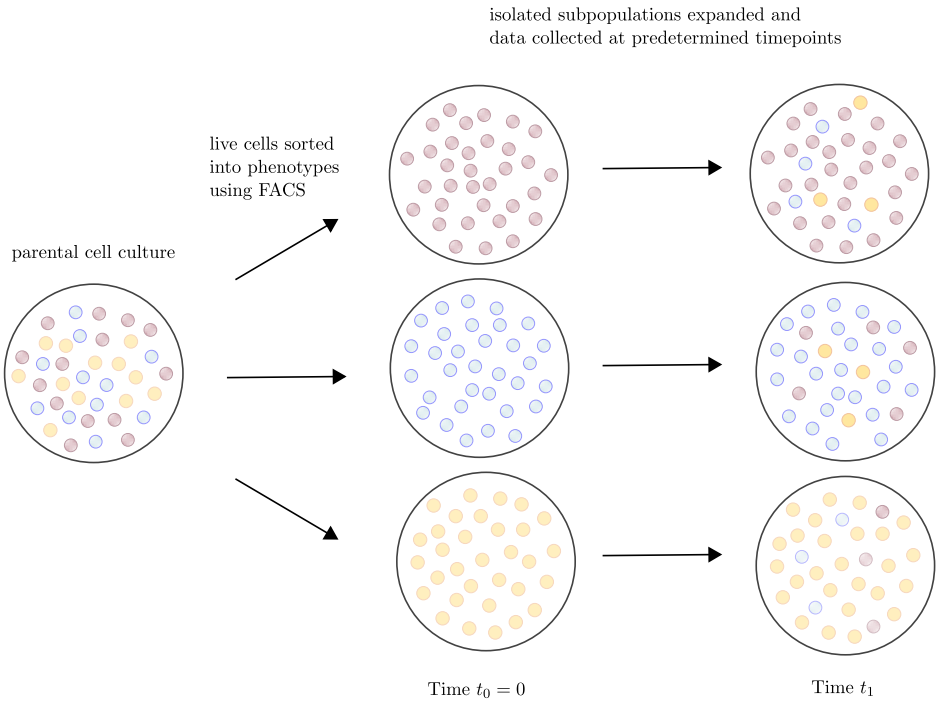

Figure 5.1: The population-level dynamics of phenotypic switching are commonly interrogated by sorting live cells into isolated phenotypic subpopulations and expanding these subpopulations in culture [9, 70, 77, 86, 117, 156, 168]. By tracking the evolution of phenotypic proportions over time and applying mathematical models of phenotypic switching, it becomes possible to estimate the quantitative parameters of phenotypic switching [24, 42, 72, 77, 82, 105, 144, 171].

the model parameters, and it enables estimation both from cell fraction and cell number data. We also use our framework to analyze the identifiability of model parameters and how it depends on the input data. This important topic has not been addressed by previous works.

The rest of the chapter is organized as follows. In Section 5.2, we introduce our stochastic model of cell division, cell death and phenotypic switching. In Section 5.3, we state our assumptions on the cell line experiments conducted and the data collected. In Sections 5.4 and 5.5, we propose statistical models for cell number and cell fraction data, respectively, and discuss how parameter estimates and confidence intervals are computed. In Section 5.6, we present theoretical analysis of the identifiability of parameters under each model. In Section 5.7, we conduct numerical experiments, and

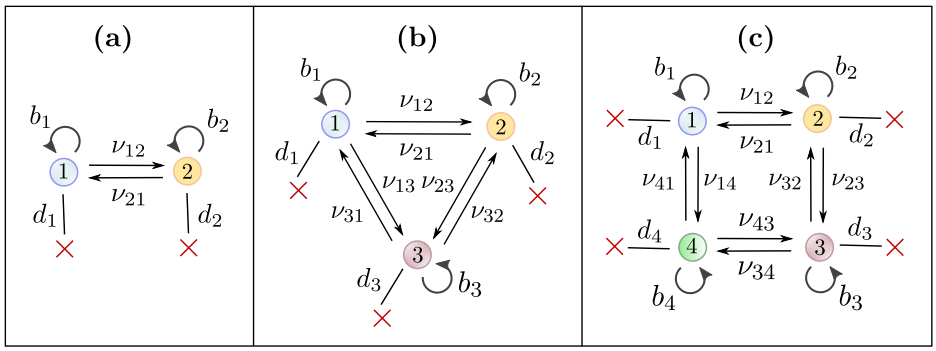

Figure 5.2: The multitype branching process model captures a variety of switching dynamics. **(a)** A two-type model captures e.g. the dynamics between HER2+ and HER2− cell states in Brx-82 and Brx-142 breast cancer cells [86]. **(b)** A three-type model captures e.g. the dynamics between stem-like, basal and luminal cell states in SUM149 and SUM159 breast cancer cells [77]. **(c)** A four-type model captures e.g. the dynamics between $\text{CD24}^{\text{Low}}/\text{ALDH}^{\text{High}}$, $\text{CD24}^{\text{Low}}/\text{ALDH}^{\text{Low}}$, $\text{CD24}^{\text{High}}/\text{ALDH}^{\text{High}}$ and $\text{CD24}^{\text{Hich}}/\text{ALDH}^{\text{Low}}$ cell states in GBC02, SCC029B, SCC070, GBC035 and SCC032 oral cancer cells [153].

in Section 5.8, we apply our framework to a published dataset. We conclude with a discussion section (Section 5.9).

5.2 Multitype branching process model

5.2.1 Model definition and model parameters

To model the cell population dynamics, we employ a multitype branching process model in continuous time, with $K \geq 2$ types [6]. In the model, a type-j cell divides into two cells at rate $b_j \geq 0$, it dies at rate $d_j \geq 0$, and it switches to type-k at rate $\nu_{jk} \geq 0$ for $k \neq j$, independently of all other cells. This means that in an infinitesimally short time interval of length $\Delta t > 0$, a type-j cell divides with probability $b_j \Delta t$, it dies with probability $d_j \Delta t$, and it switches to type k with probability $\nu_{jk} \Delta t$. The multitype branching process model captures a variety of switching dynamics previously observed in the literature (Fig. 5.2).

We allow $\nu_{jk} = 0$ for some j and k, which means that a type-j cell is not able to switch directly to type-k. However, in our exposition, we assume that the model is

irreducible, in that each cell type is accessible from any other cell type, possibly through intermediate types. In other words, for each $j, k = 1, \ldots, K$ with $k \neq j$, we assume that there exist $r \geq 0$ integers $m_1, \ldots, m_r \in \{1, \ldots, K\}$ so that $\nu_{jm_1}\nu_{m_1m_2}\cdots\nu_{m_rk} > 0$. Our estimation framework also applies to the reducible case, which is discussed in Appendix D.2 below.

For $j = 1, \ldots, K$, we define $\lambda_j := b_j - d_j$ as the net birth rate of a type-j cell. We collect the growth parameters into $1 \times K$ vectors $\mathbf{b} = (b_1, \ldots, b_K)$, $\mathbf{d} = (d_1, \ldots, d_K)$ and $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_K)$. We collect the switching rates into a $K \times (K-1)$ matrix $\boldsymbol{\nu}$, where the j-th row vector is $(\nu_{jk})_{k\neq j}$. We finally define $\boldsymbol{\lambda}^{[-j]} := \boldsymbol{\lambda} - \lambda_j\mathbf{1}$ as the vector of net birth rates relative to λ_j, with $\lambda_k^{[-j]} = \lambda_k - \lambda_j$ for $k \neq j$ and $\lambda_j^{[-j]} = 0$.

5.2.2 Stochastic processes

Let $\mathbf{n} = (n_1, \ldots, n_K)$ be the $1 \times K$ vector of starting cell numbers of each type. The state of the model at time $t \geq 0$ is encoded in the $1 \times K$ vector of cell numbers $\mathbf{Z}^{\mathbf{n}}(t) = \big(Z_1^{\mathbf{n}}(t), \ldots, Z_K^{\mathbf{n}}(t)\big)$, where $Z_k^{\mathbf{n}}(t)$ is the number of type-k cells at time t. We will commonly deal with the case $\mathbf{n} = \mathbf{e}_j$ for some $j = 1, \ldots, K$, where $\mathbf{e}_j$ is the j-th unit vector, meaning that the process is started by a single type-j cell. We write $\mathbf{Z}^{(j)}(t) := \mathbf{Z}^{\mathbf{e}_j}(t)$ for that case, and we let $\mathbf{m}^{(j)}(t)$ and $\boldsymbol{\Sigma}^{(j)}(t)$ denote the mean vector and covariance matrix for $\mathbf{Z}^{(j)}(t)$:

$$\begin{aligned}
&\mathbf{m}^{(j)}(t) := \mathbb{E}\big[\mathbf{Z}^{(j)}(t)\big],\\
&\boldsymbol{\Sigma}^{(j)}(t) := \mathbb{E}\big[\big(\mathbf{Z}^{(j)}(t) - \mathbf{m}^{(j)}(t)\big)^T\big(\mathbf{Z}^{(j)}(t) - \mathbf{m}^{(j)}(t)\big)\big], \quad t \geq 0.
\end{aligned}$$

On the event $\big\{\sum_{k=1}^K Z_k^{\mathbf{n}}(t) \neq 0\big\}$, we let $\boldsymbol{\Delta}^{\mathbf{n}}(t)$ denote the vector of cell fractions, i.e.

$$\Delta_j^{\mathbf{n}}(t) := Z_j^{\mathbf{n}}(t)/\big(\textstyle\sum_{k=1}^K Z_k^{\mathbf{n}}(t)\big), \quad j = 1, \ldots, K.$$

As above, we write $\boldsymbol{\Delta}^{(j)}(t) := \boldsymbol{\Delta}^{\mathbf{e}_j}(t)$ for $j = 1, \ldots, K$.

5.2.3 Infinitesimal generator and mean matrix

We define the $K \times K$ matrix $\mathbf{A}$ with $a_{jj} := \lambda_j - \sum_{k\neq j}\nu_{jk}$ for $j = 1, \ldots, K$ and $a_{jk} := \nu_{jk}$ for $k \neq j$ as the *infinitesimal generator* of the model, where a_{jk} is the net rate at which

Symbol	Dimension	Description
K	1	Number of types
b_j	1	Division rate of type-j cells
d_j	1	Death rate of type-j cells
ν_{jk}	1	Rate of switching from type-j to type-k
λ_j	1	Net birth rate of type-j cells, $\lambda_j = b_j - d_j$
$\lambda_k^{[-j]}$	1	Net birth rate relative to λ_j, $\lambda_k^{[-j]} = \lambda_k - \lambda_j$
$\mathbf{A}$	$K \times K$	Infinitesimal generator of the branching process model
$\mathbf{M}(t)$	$K \times K$	Mean matrix at time t, $\mathbf{M}(t) = \exp(t\mathbf{A})$
$\overline{\gamma}$	$1 \times K$	equilibrium proportions between cell types
$\mathbf{Z}^{\mathbf{n}}(t)$	$1 \times K$	Vector of cell numbers at time t, initial condition $\mathbf{n}$
$\mathbf{Z}^{(j)}(t)$	$1 \times K$	Vector of cell numbers at time t, initial condition $\mathbf{e}_j$
$\mathbf{m}^{\boldsymbol{\alpha}}(t)$	$1 \times K$	$\mathbf{m}^{\boldsymbol{\alpha}}(t) = \boldsymbol{\alpha}\mathbf{M}(t) = \sum_{j=1}^{K} \alpha_j \mathbf{m}^{(j)}(t)$
$\mathbf{m}^{(j)}(t)$	$1 \times K$	$\mathbf{m}^{(j)}(t) = \mathbb{E}\big[\mathbf{Z}^{(j)}(t)\big] = \mathbf{e}_j\mathbf{M}(t)$
$\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)$	$K \times K$	$\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t) = \sum_{j=1}^{K} \alpha_j \boldsymbol{\Sigma}^{(j)}(t)$
$\boldsymbol{\Sigma}^{(j)}(t)$	$K \times K$	$\boldsymbol{\Sigma}^{(j)}(t) = \mathbb{E}\big[\big(\mathbf{Z}^{(j)}(t) - \mathbf{m}^{(j)}(t)\big)^T \big(\mathbf{Z}^{(j)}(t) - \mathbf{m}^{(j)}(t)\big)\big]$
$\boldsymbol{\Delta}^{\mathbf{n}}(t)$	$1 \times K$	Vector of cell fractions at time t, initial condition $\mathbf{n}$
$\mathbf{p}^{\boldsymbol{\alpha}}(t)$	$1 \times K$	$\mathbf{p}^{\boldsymbol{\alpha}}(t) = \big(\boldsymbol{\alpha}\mathbf{M}(t)\mathbf{1}^T\big)^{-1}\boldsymbol{\alpha}\mathbf{M}(t)$
$\mathbf{Q}^{\boldsymbol{\alpha}}(t)$	$K \times K$	$\mathbf{Q}^{\boldsymbol{\alpha}}(t) = \mathbf{I} - \mathbf{1}^T\mathbf{p}^{\boldsymbol{\alpha}}(t)$
$\mathbf{S}^{\boldsymbol{\alpha}}(t)$	$K \times K$	$\mathbf{S}^{\boldsymbol{\alpha}}(t) = \big(\mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T\big)^{-2}\big(\mathbf{Q}^{\boldsymbol{\alpha}}(t)\big)^T \boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)\, \mathbf{Q}^{\boldsymbol{\alpha}}(t)$

Table 5.1: Notation for the stochastic model of Section 2.

a cell of type j produces a cell of type k. We next define the matrix exponential

$$\mathbf{M}(t) = \exp(t\mathbf{A}) = \textstyle\sum_{k=0}^{\infty}(t^k/k!)\mathbf{A}^k,$$

which is referred to as the *mean matrix*. It can be shown that the j-th row vector of $\mathbf{M}(t)$ is $\mathbf{m}^{(j)}(t) = \mathbb{E}\big[\mathbf{Z}^{(j)}(t)\big]$, i.e. the mean vector for the process started by a single type-j cell [6]. Note that $\mathbf{A}$ and $\mathbf{M}(t)$ depend only on the switching rates $\boldsymbol{\nu}$ and the net birth rates $\boldsymbol{\lambda}$.

5.2.4 Long-run behavior

In the absence of phenotypic switching, the subpopulation of type-j cells grows at exponential rate λ_j on average. However, in the presence of phenotypic switching, and under the assumption of irreducible switching dynamics, all subpopulations eventually

grow at the same rate σ. More precisely, there exists a real number σ and positive $1 \times K$ vectors $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_K)$ and $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_K)$ so that

$$\lim_{t\to\infty} e^{-\sigma t}\mathbf{M}(t) = \boldsymbol{\beta}^T\boldsymbol{\gamma}. \tag{5.1}$$

This means that if the process is started by a single type-j cell, the mean number of type-k cells at time t is approximately $\beta_j\gamma_k e^{\sigma t}$ when t is large. It follows that if we define

$$\overline{\gamma}_k := \gamma_k / \big(\textstyle\sum_{m=1}^K \gamma_m\big), \quad k = 1, \ldots, K, \tag{5.2}$$

then $\overline{\gamma}_k$ is the long-run proportion of type-k cells in the population, independently of the initial condition. Thus, in the long run, cell proportions tend towards an equilibrium distribution given by $\overline{\boldsymbol{\gamma}}$, which is consistent with the experimental observations discussed in the introduction. For the mathematical details, including how to compute $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, see e.g. Section V.7 of [6].

5.3 Experiments and data

We assume that each experiment is started with a known initial condition, encoded by the $1 \times K$ vector $\mathbf{n} = (n_1, \ldots, n_K)$ of starting cell numbers for each type. We let $I \geq 1$ denote the number of distinct initial conditions and $\mathbf{n}_i = (n_{i1}, \ldots, n_{iK})$ denote the i-th initial condition. We assume that for each $i = 1, \ldots, I$ and $j = 1, \ldots, K$, either $n_{ij} = 0$ or n_{ij} is large, which is generally the case for the cell line experiments discussed in the introduction.

We define $N_i := \sum_{k=1}^K n_{ik}$ as the total number of starting cells in the i-th condition and $\mathbf{f}_i = (f_{i1}, \ldots, f_{iK})$ as the vector of starting cell fractions, with $f_{ij} := n_{ij}/N_i$. We let $L \geq 1$ be the number of time points at which data is collected and we let $0 < t_1 < t_2 < \cdots < t_L$ denote the time points. Finally, we let $R \geq 1$ be the number of experimental replicates performed.

In the development of our estimation framework, we assume that each experiment returns measurements from a single time point only, meaning that the sample is discarded once measurements are taken (endpoint data). This is usually the case when

FACS or flow cytometry is used to identify phenotypes. Sometimes, the data collected is sequential, meaning that a single experiment returns measurements from multiple time points. This is e.g. the case when live cell imaging is used to identify phenotypes. While our framework is developed for endpoint data, we show in Section 5.7 that it yields reasonable estimates for sequential data.

The data collected in each experiment is either a vector $\mathbf{n}_{i,\ell,r} = (n_{i,\ell,r,1}, \ldots, n_{i,\ell,r,K})$ of cell numbers or $\mathbf{f}_{i,\ell,r} = (f_{i,\ell,r,1}, \ldots, f_{i,\ell,r,K})$ of cell fractions. Here, $n_{i,\ell,r,j}$ is the number of type-j cells in the r-th replicate of the experiment started by the i-th initial condition and ended at the ℓ-th timepoint, and $f_{i,\ell,r,j}$ is the corresponding cell fraction.

5.4 Estimation for cell number data

In this section, we propose a maximum likelihood estimation framework for cell number data. Our framework is rooted in a central limit theorem for the vector $\mathbf{Z}^{\mathbf{n}}(t)$ of cell numbers at time t, which approximates the distribution of $\mathbf{Z}^{\mathbf{n}}(t)$ by a normal distribution for a large starting population. Using the central limit theorem, we propose a statistical model and derive a likelihood function for cell number data, which we then use to compute maximum likelihood estimates and confidence intervals for all parameters.

5.4.1 Central limit theorem

We begin by establishing a central limit theorem (CLT) for the vector $\mathbf{Z}^{\mathbf{n}}(t)$ of cell numbers, which involves decomposing the branching process into i.i.d. processes started by single cells. By fixing the vector of starting cell proportions, and sending the starting population size to infinity, we can apply the standard (multivariate) CLT to obtain the following result.

Proposition 5.1. *Let $\boldsymbol{\alpha}$ be $1 \times K$ with $\alpha_i \geq 0$ for $i = 1, \ldots, K$ and $\sum_{i=1}^{K} \alpha_i = 1$. Define*

$$\begin{aligned} \mathbf{m}^{\boldsymbol{\alpha}}(t) &:= \boldsymbol{\alpha}\mathbf{M}(t) = \textstyle\sum_{j=1}^{K} \alpha_j \mathbf{m}^{(j)}(t), \\ \boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t) &:= \textstyle\sum_{j=1}^{K} \alpha_j \boldsymbol{\Sigma}^{(j)}(t). \end{aligned} \tag{5.3}$$

Let $J \geq 1$ be any integer. For any $K \times J$ matrix $\mathbf{C}$, then as $N \to \infty$,

$$N^{-1/2}\big(\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)\,\mathbf{C} - N\mathbf{m}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\big) \xrightarrow{d} \mathcal{N}\big(\mathbf{0}, \mathbf{C}^T\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\big).$$

Here, the covariance matrix $\boldsymbol{\Sigma}^{(j)}(t)$ is given by

$$\begin{aligned}\boldsymbol{\Sigma}^{(j)}(t) = 2\textstyle\int_0^t (\mathbf{M}(t-\tau))^T \mathrm{diag}\big(\mathbf{b}\odot\mathbf{m}^{(j)}(\tau)\big)(\mathbf{M}(t-\tau))d\tau \\ + \mathrm{diag}\big(\mathbf{m}^{(j)}(t)\big) - (\mathbf{m}^{(j)}(t))^T\mathbf{m}^{(j)}(t).\end{aligned} \tag{5.4}$$

Proof. Appendix D.6. □

Note that in Proposition 5.1, some coordinates of the vector $\boldsymbol{\alpha}$ of starting cell proportions are allowed to be 0. In the $N \to \infty$ regime, the starting condition $\lfloor N\boldsymbol{\alpha}\rfloor$ will therefore either include no cell or a large number of cells of any given type. This is consistent with our assumptions on the vectors $\mathbf{n}_1, \ldots, \mathbf{n}_I$ of experimental starting conditions (Section 5.3).

Also note that Proposition 5.1 is established for linear transformations $\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)\,\mathbf{C}$ of $\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)$. This allows us to obtain a CLT for cases where we do not observe the full vector $\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)$. For example, if we set $\mathbf{C} := \mathbf{1}^T$, then $\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)\,\mathbf{C} = \sum_{i=1}^K Z_i^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)$ is the total number of cells at time t. This more general CLT also becomes useful when estimating from models with reducible switching dynamics, as we discuss in Appendix D.2.

5.4.2 Statistical model

Based on Proposition 5.1, we propose the following statistical model for the data $\mathbf{n}_{i,\ell,r}$:

$$\mathbf{n}_{i,\ell,r} \sim \underbrace{N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)}_{\substack{\text{mean}\\ \text{behavior}}} + \underbrace{\mathcal{N}\big(\mathbf{0}, N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell)\big)}_{\substack{\text{variability in}\\ \text{population dynamics}}} + \underbrace{\mathcal{N}\big(\mathbf{0}, \mathbf{E}^{\text{num}}_{i,\ell}\big)}_{\substack{\text{measurement}\\ \text{error}}}. \tag{5.5}$$

The first two terms capture the mean and variance of the population dynamics, in accordance with Proposition 5.1, while the final term captures experimental measurement error.

The vectors $\mathbf{n}_{i,\ell,r}$ and $\mathbf{n}_{j,m,s}$ are assumed independent for $(i,\ell,r) \neq (j,m,s)$, and

they are assumed i.i.d. for $(i,\ell) = (j,m)$ and $r \neq s$. This assumes that data from distinct time points come from distinct experiments (endpoint data). We make this assumption since the CLT of Proposition 5.1 holds for the distribution of $\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)$ at a fixed time point t. Developing an analogous statistical model for sequential data requires extending the CLT in Proposition 5.1 to a process-level or functional CLT. As stated above, we will show in Section 5.7 that the statistical model in (5.5) yields reasonable estimates for sequential data.

Note that the mean behavior $N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)$ of the model depends only on $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}$, while the variance term $N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell)$ depends on $\boldsymbol{\nu}$, $\boldsymbol{\lambda}$ and $\mathbf{b}$ by Proposition 5.1. It is therefore natural to parametrize the first two terms in (5.5) by $\mathbf{b}, \boldsymbol{\lambda}, \boldsymbol{\nu}$ instead of the primary parameters $\mathbf{b}, \mathbf{d}, \boldsymbol{\nu}$. We assume that the $K \times K$ covariance matrix $\mathbf{E}^{\text{num}}_{i,\ell}$ associated with measurement error can be written as a function of $\mathbf{b}, \boldsymbol{\lambda}, \boldsymbol{\nu}$ and additional error parameters $\boldsymbol{\omega}_{\text{num}} = (\omega_1, \ldots, \omega_{M_{\text{num}}})$ for some $M_{\text{num}} \geq 0$. A simple example is $\mathbf{E}^{\text{num}}_{i,\ell} = \omega^2\big(\text{diag}\big(N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)\big)\big)^2$ for some $\omega > 0$, where the measurement error is assumed to be uncorrelated between types, and to scale with the mean experimental outcomes. We let $\boldsymbol{\theta}_{\text{num}} = \big[\mathbf{b}, \boldsymbol{\lambda}, \boldsymbol{\nu}, \boldsymbol{\omega}_{\text{num}}\big]$ be the complete $1 \times \big(K(K+1) + M_{\text{num}}\big)$ vector of model parameters including the error parameters.

5.4.3 Maximum likelihood estimates and confidence intervals

From the statistical model (5.5), it is straightforward to derive a likelihood function, i.e. the probability of observing the experimental data as a function of the model parameters:

$$\begin{aligned}
&\mathcal{L}_{\text{num}}\big(\boldsymbol{\theta}_{\text{num}}\big|(\mathbf{n}_{i,\ell,r})\big) \\
&= \textstyle\prod_{i=1}^{I}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^K \det\big(N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell) + \mathbf{E}^{\text{num}}_{i,\ell}\big)\Big)^{-1/2} \\
&\quad \cdot \exp\big(-\tfrac{1}{2}\big(\mathbf{n}_{i,\ell,r} - N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)\big)\big(N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell) + \mathbf{E}^{\text{num}}_{i,\ell}\big)^{-1}\big(\mathbf{n}_{i,\ell,r} - N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)\big)^T\big). \qquad (5.6)
\end{aligned}$$

We next define the negative double log-likelihood,

$$l_{\text{num}}\big(\boldsymbol{\theta}_{\text{num}}\big) := -2\log\mathcal{L}_{\text{num}}\big(\boldsymbol{\theta}_{\text{num}}\big|(\mathbf{n}_{i,\ell,r})\big). \qquad (5.7)$$

The maximum likelihood estimate $\widehat{\boldsymbol{\theta}}_{\text{num}}$ for the parameter vector $\boldsymbol{\theta}_{\text{num}}$ is obtained by minimizing $l_{\text{num}}(\boldsymbol{\theta}_{\text{num}})$ over a set of feasible parameters $\boldsymbol{\Theta}_{\text{num}}$:

$$\widehat{\boldsymbol{\theta}}_{\text{num}} := \text{argmin}_{\boldsymbol{\theta}_{\text{num}} \in \boldsymbol{\Theta}_{\text{num}}} \, l_{\text{num}}(\boldsymbol{\theta}_{\text{num}}). \tag{5.8}$$

In the feasible set $\boldsymbol{\Theta}_{\text{num}}$, we restrict the parameter values so that $\boldsymbol{\nu} \geq \mathbf{0}$, $\mathbf{b} \geq \mathbf{0}$ and $\boldsymbol{\lambda} \leq \mathbf{b}$. Further restrictions can be made depending on the context, see e.g. Section 5.8 and Appendix D.2. A $1-\alpha$ likelihood-based confidence interval $\left[\theta^{-}_{\text{num},i}, \theta^{+}_{\text{num},i}\right]$ for the i-th model parameter can be obtained by solving the following two constrained optimization problems:

$$\begin{aligned} \theta^{-}_{\text{num},i} &= \min_{\boldsymbol{\theta}_{\text{num}} \in \boldsymbol{\Theta}_{\text{num}}} \left\{\theta_{\text{num},i} : l_{\text{num}}(\boldsymbol{\theta}_{\text{num}}) \leq l_{\text{num}}(\widehat{\boldsymbol{\theta}}_{\text{num}}) + \chi^2_{1,1-\alpha}\right\}, \\ \theta^{+}_{\text{num},i} &= \max_{\boldsymbol{\theta}_{\text{num}} \in \boldsymbol{\Theta}_{\text{num}}} \left\{\theta_{\text{num},i} : l_{\text{num}}(\boldsymbol{\theta}_{\text{num}}) \leq l_{\text{num}}(\widehat{\boldsymbol{\theta}}_{\text{num}}) + \chi^2_{1,1-\alpha}\right\}, \end{aligned} \tag{5.9}$$

where $\widehat{\boldsymbol{\theta}}_{\text{num}}$ is the MLE estimator defined by (5.8) and $\chi^2_{1,1-\alpha}$ is the $(1-\alpha)$-th quantile of the χ^2-distribution, see e.g. [13, 60, 116, 129, 152]. Our estimation framework is based on solving the optimization problems in (5.8) and (5.9) using the sqp solver in MATLAB.

5.5 Estimation for cell fraction data

In this section, we propose a maximum likelihood estimation framework for cell fraction data. As for cell number data, the framework is rooted in a central limit theorem for the vector of cell fractions $\boldsymbol{\Delta}^{\mathbf{n}}(t)$, which we use to propose a statistical model and derive a likelihood function for cell fraction data.

5.5.1 Central limit theorem

We begin by establishing a central limit theorem for the vector of cell fractions $\boldsymbol{\Delta}^{\mathbf{n}}(t)$. This CLT has already been established for the case of an isolated large starting population by Yakovlev and Yanev [167]. We extend their argument to more general starting conditions by fixing the vector $\boldsymbol{\alpha}$ of starting cell proportions and sending the total population size N to infinity. We provide a simplified expression for the covariance matrix $\mathbf{S}^{\boldsymbol{\alpha}}(t)$ and show that the mean function $\mathbf{p}^{\boldsymbol{\alpha}}(t)$ can be written solely in terms of $\boldsymbol{\nu}$ and

$\boldsymbol{\lambda}^{[-j]}$ for any $j = 1, \ldots, K$.

Proposition 5.2. *Let $\boldsymbol{\alpha}$ be $1 \times K$ with $\alpha_i \geq 0$ for $i = 1, \ldots, K$ and $\sum_{i=1}^{K} \alpha_i = 1$. Define*

$$\begin{aligned} \mathbf{p}^{\boldsymbol{\alpha}}(t) &:= \left(\mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T\right)^{-1}\mathbf{m}^{\boldsymbol{\alpha}}(t), \\ \mathbf{Q}^{\boldsymbol{\alpha}}(t) &:= \mathbf{I} - \mathbf{1}^T\mathbf{p}^{\boldsymbol{\alpha}}(t), \\ \mathbf{S}^{\boldsymbol{\alpha}}(t) &:= \left(\mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T\right)^{-2}\left(\mathbf{Q}^{\boldsymbol{\alpha}}(t)\right)^T \boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)\, \mathbf{Q}^{\boldsymbol{\alpha}}(t), \end{aligned} \tag{5.10}$$

$\mathbf{p}^{(j)}(t) := \mathbf{p}^{\mathbf{e}_j}(t)$ and $\mathbf{S}^{(j)}(t) := \mathbf{S}^{\mathbf{e}_j}(t)$, where $\mathbf{m}^{\boldsymbol{\alpha}}(t)$ and $\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)$ are defined as in (5.3). Let $J \geq 1$ be any integer. For any $K \times J$ matrix $\mathbf{C}$, then as $N \to \infty$,

$$N^{1/2}\left(\boldsymbol{\Delta}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)\,\mathbf{C} - \mathbf{p}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\right) \xrightarrow{d} \mathcal{N}\left(\mathbf{0}, \mathbf{C}^T\mathbf{S}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\right).$$

Here, the mean function $\mathbf{p}^{\boldsymbol{\alpha}}(t)$ can be written solely as a function of the switching rates $\boldsymbol{\nu}$ and the relative net birth rates $\boldsymbol{\lambda}^{[-j]}$ for any $j = 1, \ldots, K$.

Proof. Appendix D.7. □

5.5.2 Statistical model and maximum likelihood estimates

Based on Proposition 5.2, we propose the following statistical model for the data $\mathbf{f}_{i,\ell,r}$:

$$\mathbf{f}_{i,\ell,r} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\left(\mathbf{0}, N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)\right) + \mathcal{N}\left(\mathbf{0}, \mathbf{E}_{i,\ell}^{\text{frac}}\right). \tag{5.11}$$

Note that the mean behavior $\mathbf{p}^{\mathbf{f}_i}(t_\ell)$ depends only on $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}^{[-1]}$, while the variance term $N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)$ depends on all model parameters $\boldsymbol{\nu}, \boldsymbol{\lambda}^{[-1]}, \lambda_1, \mathbf{d}$. The choice of type-1 as a reference phenotype is arbitrary, and we use $\mathbf{d}$ as opposed to $\mathbf{b}$ as we have found it to perform well numerically. We assume that the $K \times K$ covariance matrix $\mathbf{E}_{i,\ell}^{\text{frac}}$ associated with measurement error can be written as a function of $\mathbf{d}, \lambda_1, \boldsymbol{\lambda}^{[-1]}, \boldsymbol{\nu}$ and added error parameters $\boldsymbol{\omega}_{\text{frac}} = (\omega_1, \ldots, \omega_{M_{\text{frac}}})$ for some $M_{\text{frac}} \geq 0$. A simple example is $\mathbf{E}_{i,\ell}^{\text{frac}} = \omega^2\mathbf{I}$ for some $\omega > 0$, where the measurement error is assumed uncorrelated between cell types, and of the same magnitude for all initial conditions and all time points. We let $\boldsymbol{\theta}_{\text{frac}} = \left[\mathbf{d}, \lambda_1, \boldsymbol{\lambda}^{[-1]}, \boldsymbol{\nu}, \boldsymbol{\omega}_{\text{frac}}\right]$ denote the complete vector of model parameters including the error parameters.

When deriving a likelihood function for the statistical model (5.11), we note that

the last coordinate of $\mathbf{f}_{i,\ell,r}$ provides no new information over the first $K-1$ coordinates, since the coordinates always sum to one. In the likelihood function, we therefore only consider the first $K-1$ coordinates, which we can accomplish by multiplying $\mathbf{f}_{i,\ell,r}$ by the $K \times (K-1)$ matrix $\mathbf{B}$ with 1 on the diagonal and 0 off it. In this way, we obtain the following likelihood:

$$\begin{aligned}
&\mathcal{L}_{\text{frac}}\big(\boldsymbol{\theta}_{\text{frac}}\big|(\mathbf{f}_{i,\ell,r})\big) \\
&= \textstyle\prod_{i=1}^{I}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^{K-1}\det\big(\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\mathbf{B}\big)\Big)^{-1/2} \\
&\quad\cdot \exp\Big(-\tfrac{1}{2}\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{\mathbf{f}_i}(t_\ell)\big)\,\mathbf{B}\,\big(\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\,\mathbf{B}\big)^{-1}\,\mathbf{B}^T\,\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{\mathbf{f}_i}(t_\ell)\big)^T\Big).
\end{aligned} \tag{5.12}$$

As for cell number data, we define the negative double log-likelihood,

$$l_{\text{frac}}\big(\boldsymbol{\theta}_{\text{frac}}\big) := -2\log\mathcal{L}_{\text{frac}}\big(\boldsymbol{\theta}_{\text{frac}}\big|(\mathbf{f}_{i,\ell,r})\big), \tag{5.13}$$

and obtain the maximum likelihood estimate for $\boldsymbol{\theta}_{\text{frac}}$ by solving

$$\widehat{\boldsymbol{\theta}}_{\text{frac}} := \operatorname{argmin}_{\boldsymbol{\theta}_{\text{frac}}\in\boldsymbol{\Theta}_{\text{frac}}} l_{\text{frac}}(\boldsymbol{\theta}_{\text{frac}}). \tag{5.14}$$

The computation of confidence intervals proceeds as described in Section 5.4.3.

5.5.3 Comparison with least squares estimation

It is instructive to compare our statistical model (5.11) with the following simpler model, where the data is assumed to follow the mean behavior of (5.11) with an i.i.d. Gaussian error:

$$\mathbf{f}_{i,\ell,r} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0},\omega^2\mathbf{I}\big). \tag{5.15}$$

In this case, the model only involves the parameters $\boldsymbol{\lambda}^{[-1]}$ and $\boldsymbol{\nu}$ in addition to ω, and the maximum likelihood estimates of $\boldsymbol{\lambda}^{[-1]}$ and $\boldsymbol{\nu}$ can be computed independently of ω

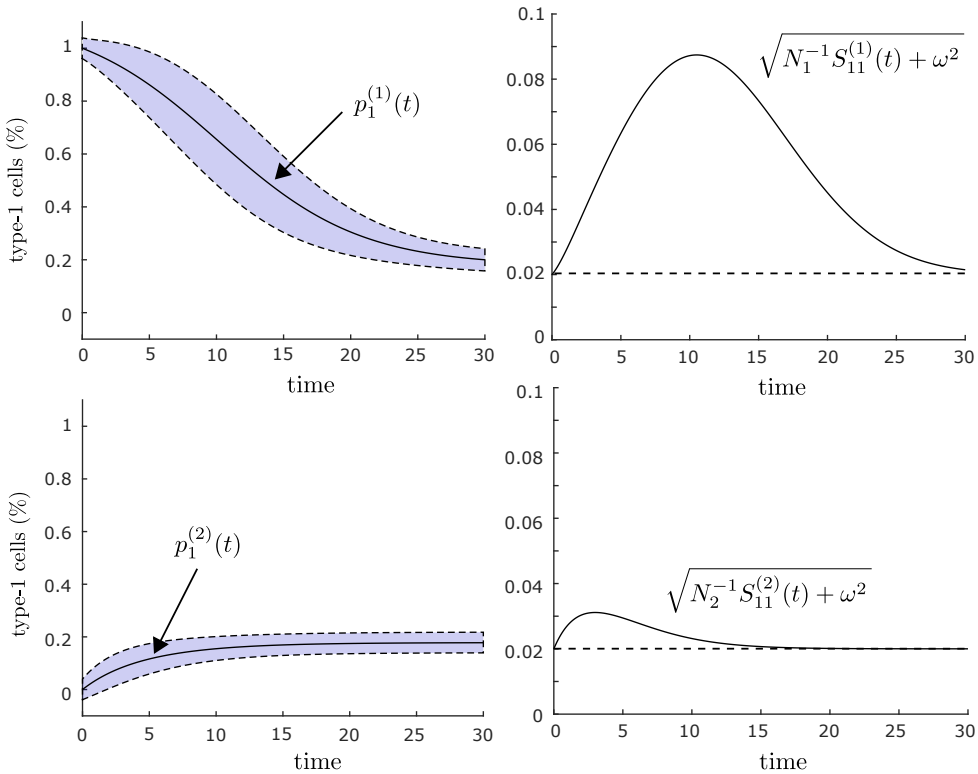

Figure 5.3: Behavior of the model $\mathbf{f}_{i,\ell,r} \sim \mathbf{p}^{(i)}(t_\ell) + \mathcal{N}\big(\mathbf{0}, N_i^{-1}\mathbf{S}^{(i)}(t_\ell)\big) + \mathcal{N}\big(\mathbf{0}, \omega^2\mathbf{I}\big)$ for $K = 2$, $b_1 = 0.6$, $d_1 = 0.3$, $b_2 = 1.0$, $d_2 = 0.5$, $\nu_{12} = 0.02$, $\nu_{21} = 0.04$ and $\omega = 0.02$. The top left panel shows the time-evolution of the fraction of type-1 cells in the population, starting from $N_1 = 100$ type-1 cells. The solid line indicates the mean cell fraction, while the shaded regions indicate 95% confidence intervals. The top right panel shows the time-evolution of the standard deviation of the type-1 fraction. The bottom panels show the same quantities, starting from $N_2 = 100$ type-2 cells.

as

$$\{\widehat{\boldsymbol{\lambda}}^{[-1]}, \widehat{\boldsymbol{\nu}}\} = \text{argmin}_{\boldsymbol{\lambda}^{[-1]},\boldsymbol{\nu}} \sum_{i=1}^{I}\sum_{\ell=1}^{L}\sum_{r=1}^{R} \left\|\mathbf{f}_{i,\ell,r} - \mathbf{p}^{\mathbf{f}_i}(t_\ell)\right\|^2. \tag{5.16}$$

Note that (5.16) simply minimizes the sum of squared errors between the mean prediction of our stochastic model and the data. This is a common approach to parameter estimation in systems biology, and it can be viewed as the continuous-time version of the TRANSCOMPP estimation problem [82], which is stated as expression (D.2) in Appendix D.1.

To illustrate the difference between our statistical model (5.11) and the simple model (5.15), we show in Figure 5.3 an example of how the error terms in (5.11) evolve over time. Note that for both initial conditions, the error increases during the initial moments, as the population moves from known initial conditions towards more uncertainty. Then, as cell fractions approach their steady state, the error decreases again. Also note that the maximum error is significantly different depending on the initial condition. In contrast, in the simple model (5.15), the statistical noise is the same for all initial conditions and all time points. In addition, the simple model does not account for correlations between cell types, which (5.11) does. When encountering real data, one can evaluate whether the data displays these properties by comparing the likelihoods of the two models, as we show in Section 5.8 below.

5.6 Structural identifiability

In this section, we analyze the structural identifiability of the statistical models of Sections 5.4 and 5.5. For a stochastic model, a parameter is structurally identifiable if complete knowledge of the data distribution uniquely determines the value of the parameter, in the absence of any measurement noise [23, 135]. Here, we will assume that we know the behavior of the mean functions $\mathbf{m}^{(j)}(t)$ and $\mathbf{p}^{(j)}(t)$ and the covariance functions $\mathbf{\Sigma}^{(j)}(t)$ and $\mathbf{S}^{(j)}(t)$ close to time 0, and we will analyze to what extent the model parameters can be extracted from this information. In other words, we are interested in the following question: If we conduct experiments started from isolated subpopulations, and perfect observations are obtained for the first two statistical moments of the data close to time 0, is it possible to learn the model parameters?

The analysis in this section serves two purposes. First of all, it ascertains whether in this idealized setting, the model parameters can be extracted uniquely from short-term observations of the population dynamics. Second, the analysis indicates how much information is required to estimate each model parameter accurately, which yields valuable insights into how comparatively difficult it is to estimate the parameters from more limited data.

5.6.1 Cell number data

In the following proposition, we show that for cell number data, the switching rates $\boldsymbol{\nu}$ and the net birth rates $\boldsymbol{\lambda}$ can be recovered uniquely from knowledge of the mean functions $\mathbf{m}^{(j)}(t)$ close to time 0, while the birth rates $\mathbf{b}$ can be recovered from the covariance matrices $\boldsymbol{\Sigma}^{(j)}(t)$.

Proposition 5.3. *(1) For each $j = 1, \ldots, K$, the switching rates ν_{jk}, $k \neq j$, and the net birth rate λ_j are uniquely determined by $\frac{d}{dt}\mathbf{m}^{(j)}(t)\big|_{t=0}$.*

(2) For each $j = 1, \ldots, K$, if the switching rates ν_{jk}, $k \neq j$, and the net birth rate λ_j are known, the birth rate b_j is uniquely determined by $\left(\frac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\big|_{t=0}\right)_{jj}$.

Proof. Appendix D.8. □

Proposition 5.3 establishes the structural identifiability of all model parameters for cell number data. The process of extracting the parameters as suggested by Proposition 5.3 can be thought of as follows: If we want to know ν_{jk} for some $k \neq j$, we can simply plot the mean function $M_{jk}(t) = \mathbb{E}\big[Z_k^{(j)}(t)\big]$ and compute its slope at 0. If we want to know the birth rate b_j, we can plot the variance function $\big(\boldsymbol{\Sigma}^{(j)}(t)\big)_{jj} = \mathrm{Var}\big[Z_j^{(j)}(t)\big]$ and compute its slope at 0.

It is important to note that we are not suggesting to use this approach to estimate parameters from real data. Instead, we are establishing theoretically that there is sufficient information in the distribution of the data close to time 0 to determine all model parameters uniquely. In particular, we can in theory predict the entire evolutionary trajectory of the population from short-term observations of the initial population dynamics.

5.6.2 Cell fraction data

In the following proposition, we show that for cell fraction data, only the switching rates $\boldsymbol{\nu}$ can be recovered from the slopes of the mean functions $\mathbf{p}^{(j)}(t)$ at time 0. The net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ can be recovered from the curvatures of the mean functions at time 0 or from the equilibrium proportions $\overline{\gamma}$ between cell types if they are known. We are not able to learn any more parameters from the mean functions, since $\mathbf{p}^{(j)}(t)$ can be written solely as a function of $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}^{[-1]}$ by Proposition 5.2. The slopes of

Moment	Derivative	Cell number data	Cell fraction data
1	1	$\boldsymbol{\nu}, \boldsymbol{\lambda}$	$\boldsymbol{\nu}$
	2	–	$\boldsymbol{\lambda}^{[-1]}$
2	1	$\mathbf{b}$	$\boldsymbol{\nu}$

Table 5.2: Summary of the structural identifiability analysis of Propositions 5.3 and 5.4. For cell number data, the switching rates $\boldsymbol{\nu}$ and the net birth rates $\boldsymbol{\lambda}$ are identifiable from the slopes (first derivatives) of the mean functions $\mathbf{m}^{(j)}(t)$ (first moments) at time 0. The birth rates $\mathbf{b}$ are identifiable from the slopes of the covariance functions $\boldsymbol{\Sigma}^{(j)}(t)$ (second moments). For cell fraction data, only the switching rates $\boldsymbol{\nu}$ are identifiable from the slopes of the mean functions $\mathbf{p}^{(j)}(t)$, while the net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ can be determined from their curvatures (second derivatives). In contrast to cell number data, the slopes of the covariance functions $\mathbf{S}^{(j)}(t)$ for cell fraction data provide no extra information on the model parameters.

the covariance functions $\mathbf{S}^{(j)}(t)$ depend only on $\boldsymbol{\nu}$, meaning that they provide no extra information on the model parameters.

Proposition 5.4. *(1) For $j = 1, \ldots, K$, the switching rates ν_{jk}, $k \neq j$, are uniquely determined by $\frac{d}{dt}\mathbf{p}^{(j)}(t)\big|_{t=0}$.*

(2) If the switching rates $\boldsymbol{\nu}$ are known, the net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ are uniquely determined by (i) $\frac{d^2}{dt^2}\mathbf{p}^{(j)}(t)\big|_{t=0}$ for $j = 1, \ldots, K$ or (ii) the equilibrium proportions $\overline{\boldsymbol{\gamma}}$.

(3) For $j = 1, \ldots, K$, $\frac{d}{dt}\mathbf{S}^{(j)}(t)\big|_{t=0}$ only depends on the switching rates ν_{jk} for $k \neq j$.

Proof. Appendix D.9. □

As for the remaining model parameters, the net birth rate λ_1 and the birth rates $\mathbf{b}$, they require information on the curvatures of the covariance functions $\mathbf{S}^{(j)}(t)$ at time 0 at the least. We will not analyze the structural identifiability of these parameters further. Proposition 5.4 indicates that one should not expect to be able to estimate these parameters accurately from cell fraction data, which is confirmed by numerical experiments in Section 5.7.3.

5.6.3 Comparison

The results of our identifiability analysis are summarized in Table 5.2. The birth rates $\mathbf{b}$ are harder to estimate, since they require second moment information, but they may still be obtainable with sufficient data, as we discuss further in Section 5.7. For cell fraction data, the switching rates $\boldsymbol{\nu}$ are easy to estimate using the mean behavior of the population. The net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ can also be estimated from the mean, but they require more information. The remaining model parameters are unlikely to be obtainable for cell fraction data.

5.7 Numerical experiments

In this section, we apply our maximum likelihood framework to computer-generated data, and we explore how well the model parameters can be estimated depending on what data is collected. In all cases, we assume that experiments are conducted from isolated initial conditions, and we assume no measurement noise, i.e. $\mathbf{E}_{i,\ell}^{\text{num}} = \mathbf{0}$ and $\mathbf{E}_{i,\ell}^{\text{frac}} = \mathbf{0}$.

5.7.1 Implementation in MATLAB

Our estimation framework is implemented in MATLAB. The framework returns (i) a maximum likelihood estimate and (ii) a likelihood-based confidence interval for each parameter, using the sequential quadratic programming (sqp) solver in MATLAB. Before solving the maximum likelihood problem, we compute initial parameter estimates from a simpler model, which we use to initialize the optimization and to rescale the model parameters so that they are of similar magnitude. In most cases, we have found it sufficient to solve the maximum likelihood problem once, starting from the simple estimates. However, our MATLAB codes provide the option to solve the problem several times starting using different initial guesses. Details of the implementation are provided in Appendix D.3.

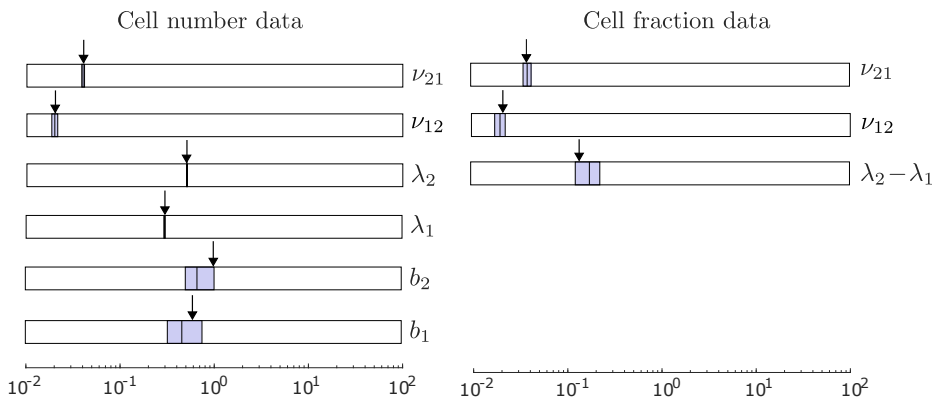

Figure 5.4: Graphical depiction of the output of our estimation framework. We first generated artificial cell-number and cell-fraction data by simulating the branching process model of Section 5.2 for $b_1 = 0.6$, $d_1 = 0.3$, $b_2 = 1.0$, $d_2 = 0.5$, $\nu_{12} = 0.02$, $\nu_{21} = 0.04$, $\mathbf{f}_1 = [1, 0]$, $\mathbf{f}_2 = [0, 1]$ and $N_1 = N_2 = 1,000$. Using this data, we computed maximum likelihood estimates and likelihood-based 95% confidence intervals (CIs) for the model parameters. For each parameter, the shaded region indicates the CI, the vertical bar inside the interval indicates the MLE estimate, and the arrow points to the true value of the parameter.

5.7.2 Illustrative example

For illustrative purposes, we first show a graphical depiction of the output of our estimation framework for a single dataset. We generated artificial cell number and cell fraction data by performing a stochastic simulation of the branching process model from Section 5.2. We then used the data to compute MLE estimates and confidence intervals for the model parameters. The data was generated assuming $K = 2$ cell types, $I = 2$ isolated initial conditions, $L = 6$ time points, and $R = 3$ replicates. Estimation results are shown in Figure 5.4.

Note first the difference in scale between the switching rates and the rates involving cell division and death. This is typically the case, since epigenetic modifications can generally be retained for 10–10^5 cell divisions [22, 119]. Also note that all model parameters are estimated more accurately for cell number data than cell fraction data, in that their confidence intervals are narrower for cell number data. Otherwise, the relative accuracy with which different model parameters can be estimated is in line with our identifiability analysis in Section 5.6.

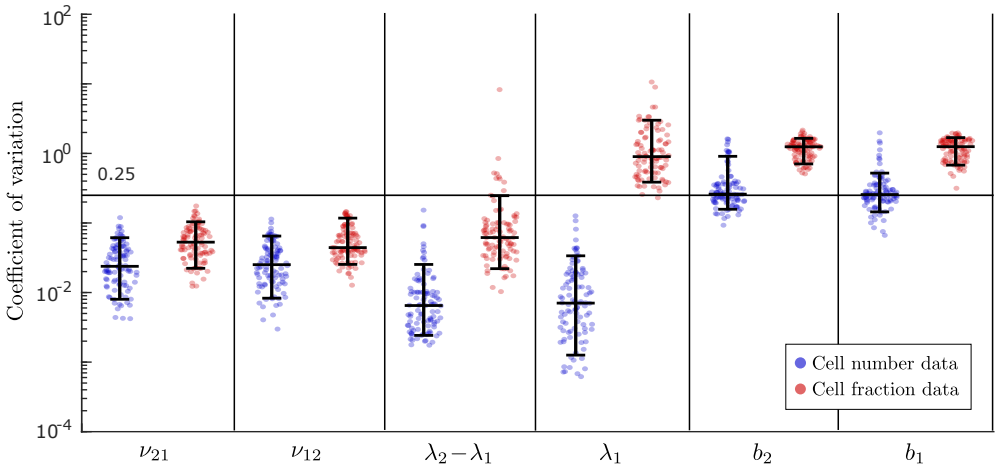

CV	ν_{21}		ν_{12}		$\lambda_2 - \lambda_1$		λ_1		b_2		b_1	
Median	0.0238	0.0530	0.0252	0.0443	0.0065	0.0617	0.0070	0.8987	0.2595	1.2452	0.2560	1.2496
10th pct	0.0080	0.0223	0.0083	0.0254	0.0024	0.0220	0.0013	0.3862	0.1575	0.7057	0.1443	0.6787
90th pct	0.0614	0.1039	0.0649	0.1175	0.0253	0.2466	0.0337	2.9983	0.9082	1.6460	0.5210	1.6868

Figure 5.5: Assessment of estimation error across a wide range of biologically realistic parameter regimes. We first generated 100 different parameter regimes, then generated 100 artifical datasets for each regime, and finally computed parameter estimates for each dataset. To generate the parameter regimes, we sampled birth and death rates uniformly between 0 and 1, and sampled switching rates log-uniformly between 10^{-1} and 10^{-3} (Appendix D.4). For each parameter and each parameter regime, we used the 100 estimates to compute the coefficient of variation (CV) for the estimates, which measures the error in the estimation. Each dot in the figure represents the CV for a single parameter under a single regime, with the blue dots (resp. red dots) representing estimates from cell number data (resp. cell fraction data). Collectively, the dots enable comparison of estimation error between different model parameters and between cell number and cell fraction data. The horizontal bars represent the 10th percentile, median and 90th percentile of the CVs, bottom to top.

5.7.3 Estimation across a wide range of biologically realistic regimes

For a more thorough evaluation of estimation accuracy, we generated 10,000 artificial datasets for $K = 2$ cell types. We first generated 100 biologically realistic parameter regimes and then generated 100 datasets for each regime. To generate the parameter regimes, we sampled birth and death rates uniformly between 0 and 1, and sampled switching rates log-uniformly between 10^{-1} and 10^{-3}. We considered both regimes where the two phenotypes have positive net birth rates ($\lambda_1, \lambda_2 > 0$) and regimes where

one phenotype has a negative net birth rate ($\lambda_1 < 0, \lambda_2 > 0$). The latter regimes are relevant to the dynamics of anti-cancer treatment response, where one phenotype is drug-sensitive and the other is drug-tolerant. As in Section 5.7.2, we assumed $I = 2$ isolated initial conditions, $L = 6$ time points and $R = 3$ replicates. Further details of the data generation are provided in Appendix D.4.

For each dataset, we used our framework to compute MLE estimates for all model parameters. In this way, we obtained 100 estimates of each parameter under each parameter regime, which we used to compute the coefficient of variation (CV) for the MLE estimator of the parameter. The CV is the sample standard deviation of the MLE estimator as a proportion of its sample mean, and it measures the percentage error in the estimation.

The results are shown in Figure 5.5. A horizontal line is drawn at 25% CV to indicate whether parameters can be estimated with reasonable accuracy. Note that for the switching rates $\boldsymbol{\nu}$, the median CV for cell fraction data is about twice as large as for cell number data. The median CV for the net birth rate difference $\lambda_2 - \lambda_1$ is an order of magnitude larger for cell fraction data than cell number data, and it is two orders of magnitude larger for the net birth rate λ_1. The birth rates $\mathbf{b}$ can in many cases be estimated reasonably well for cell number data, whereas they are never estimated accurately for cell fraction data. These results are very much in line with our identifiability analysis in Section 5.6.

Note that for cell fraction data, the estimation error for the net birth rate difference $\lambda_2 - \lambda_1$ exceeds the 25% threshold CV for several parameter regimes. This occurs when $\lambda_2 - \lambda_1$ is small in magnitude, i.e. when it is smaller than 0.1 in regimes where the birth rates lie between 0.1 and 1. Note in contrast that for cell number data, the estimation error for $\lambda_2 - \lambda_1$ never exceeds the 25% threshold. This indicates that for cell fraction data, it may be difficult to distinguish the net birth rate difference $\lambda_2 - \lambda_1$ from 0 unless it is relatively pronounced. We discuss this point further in Section 5.8 below.

5.7.4 Experimental design: Adding replicates vs. adding time points

Our framework can be used to evaluate to what extent additional data can improve parameter estimates and to identify experimental designs that best accomplish this goal. To illustrate this point, we compared the effect of (i) doubling the number of

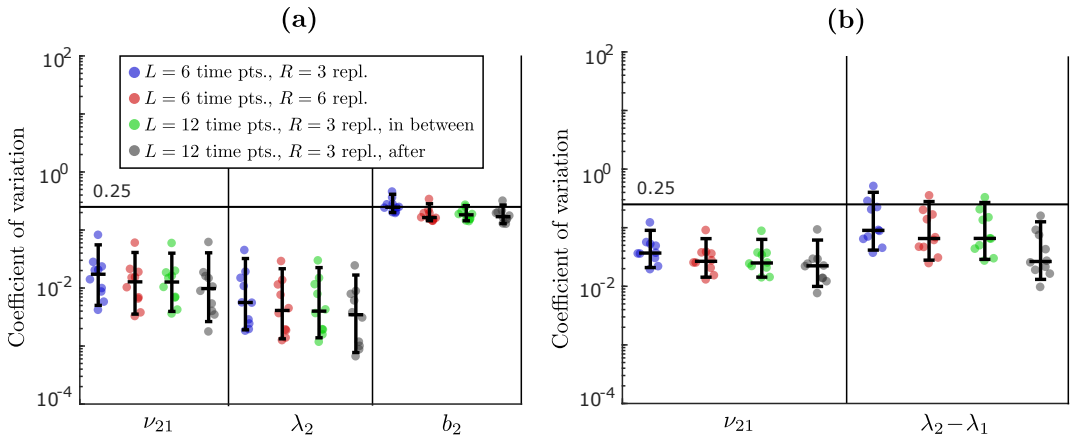

Cell number data	ν_{21}				λ_2				b_2			
Median	0.0173	0.0128	0.0127	0.0098	0.0056	0.0041	0.0040	0.0035	0.2465	0.1638	0.1826	0.1697
10th pct	0.0050	0.0035	0.0040	0.0026	0.0019	0.0013	0.0014	0.0008	0.1999	0.1447	0.1441	0.1289
90th pct	0.0555	0.0410	0.0398	0.0405	0.0322	0.0215	0.0225	0.0167	0.4129	0.2834	0.2624	0.2692

Cell fraction data	ν_{21}				$\lambda_2 - \lambda_1$			
Median	0.0369	0.0267	0.0250	0.0223	0.0900	0.0653	0.0656	0.0265
10th pct	0.0209	0.0143	0.0143	0.0099	0.0414	0.0279	0.0288	0.0131
90th pct	0.0901	0.0650	0.0632	0.0615	0.4015	0.2792	0.2688	0.1260

Figure 5.6: Comparison of estimation error for different experimental designs when the number of data points is doubled. We generated 10 parameter regimes and 100 datasets for each regime. The blue dots represent estimation from datasets with $L = 6$ time points and $R = 3$ replicates. The red dots represent estimation from $L = 6$ time points and $R = 6$ replicates. The green and grey dots represent estimation from $L = 12$ time points and $R = 3$ replicates, where the extra time points are added in between and after the previous time points, respectively. Panel **(a)** shows estimation from cell number data and panel **(b)** shows estimation from cell fraction data.

replicates from $R = 3$ to $R = 6$ (design 1), (ii) doubling the number of time points from $T = 6$ to $T = 12$, adding time points in between the previous time points (design 2), and (iii) doubling the number of time points, adding time points after the previous points (design 3) (Appendix D.4). Here, we generated 10 parameter regimes and 100 datasets for each regime. The results are shown in Figure 5.6.

For cell number data, the median CV for the switching rate ν_{21} and the net birth rate λ_2 reduces by 26% and 27%, respectively, when the number of replicates is doubled (design 1) (Fig. 5.6a). This is consistent with the fact that the standard deviation of an MLE estimator can be expected to decrease with $1/\sqrt{n}$, where n is the number

of datapoints ($1 - 1/\sqrt{2} = 0.29$) [28]. Adding data from time points in between the previous time points (design 2) has a similar effect on the median CV. However, adding time points after the previous points (design 3) reduces the median CV of ν_{21} and λ_2 by 23% and 16%, respectively, over adding replicates (design 1). We also note that the 10th percentile of the CV for ν_{21} and λ_2 reduces by 26% and 42%, respectively, between design 1 and design 3, which indicates that the degree of improvement between design 1 and design 3 depends very much on the parameter regime.

For cell fraction data, the relative attractiveness of the three experimental designs is similar (Fig. 5.6b). However, in this case, the estimate for the net birth rate difference $\lambda_2 - \lambda_1$ benefits significantly more from using design 3 than the estimate for the switching rate ν_{21}. For example, the median CV for ν_{21} reduces by 16% and the 10th percentile by 30% between design 1 and design 3, while the analogous reduction for $\lambda_2 - \lambda_1$ is 59% and 53%, respectively.

In our structural identifiability analysis for cell fraction data (Section 5.6.2), we observed that it is more difficult to estimate $\lambda_2 - \lambda_1$ than ν_{21} from the initial population dynamics, and that $\lambda_2 - \lambda_1$ can be identified from the equilibrium proportions $\overline{\boldsymbol{\gamma}}$ if the switching rates $\boldsymbol{\nu}$ are known. The fact that adding more information on the long-run behavior of the population benefits the estimation of $\lambda_2 - \lambda_1$ more than ν_{21} is consistent with these insights. Of course, the results of the previous Section 5.7.3 indicate that the estimation of $\lambda_2 - \lambda_1$ can be improved even further by using cell number data as opposed to cell fraction data. In general, Sections 5.7.3 and 5.7.4 show how our framework can be used to evaluate the estimation accuracy that can be achieved by different experimental designs, depending e.g. on what data is collected, when it is collected, how many replicates are performed, etc.

5.7.5 Improving identifiability of the rates of cell division and cell death

For cell number data, even though the birth rates $\mathbf{b}$ can be estimated reasonably well in many cases by Section 5.7.3, they are estimated much less accurately than the net birth rates $\boldsymbol{\lambda}$ and the switching rates $\boldsymbol{\nu}$. In Figure 5.8a, we show that as the number of replicates is increased from 3 to 20 or above, the accuracy in the estimation becomes more acceptable. However, even with 100 replicates, the birth rates $\mathbf{b}$ are still estimated

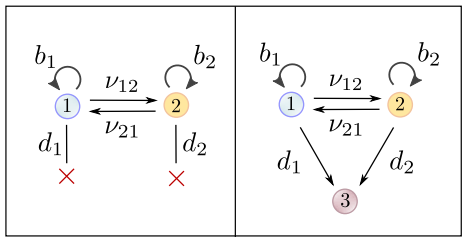

Figure 5.7: Augmentation of the mathematical model when data is available on the number of dead cells at each time point. In that case, in stead of cells being lost from the model upon dying (left panel), they transition into a new state (right panel).

less accurately than the net birth rates $\boldsymbol{\lambda}$ with 3 replicates, see Figure 5.5. As we mentioned in the introduction (Section 5.1), data on the number of cells in each state at each time point can be obtained by separately measuring the fraction of cells in each state and the total number of cells. In addition, it is often possible to estimate the number of dead cells at each time point, see e.g. [42]. If this data is obtained, we can augment our mathematical model by introducing a new cell state, which cells transition into upon death. In Figure 5.8b, we show that if we apply our estimation framework to this model, the birth rates $\mathbf{b}$ become as easy to estimate as the net birth rates $\boldsymbol{\lambda}$. Thus, if data is collected on the number of live and dead cells at each time point, it becomes possible to estimate all model parameters accurately using our framework. It should be noted that data collection on the number of dead cells is confounded by the fact that dead cells will eventually dissolve and be cleared from the system. This can potentially be addressed by introducing a clearance rate for cells in the dead cell state.

5.7.6 Estimation using endpoint data vs. sequential data

We conclude by examining how well our estimation framework applies to sequential data, when data is collected at multiple time points in the same experiment (Section 5.3). In Figure 5.9, we see that for cell number data, the CV for each parameter approximately doubles when applying our framework to sequential data vs. endpoint data. However, it remains true that the switching rates $\boldsymbol{\nu}$ and net birth rates $\boldsymbol{\lambda}$ can be estimated with good accuracy. For cell fraction data, the difference in the estimation error for $\boldsymbol{\nu}$ and $\lambda_2 - \lambda_1$ is even smaller. Together, these results indicate that our framework can yield reasonable estimates for sequential data. At the same time, for cell number data in particular,

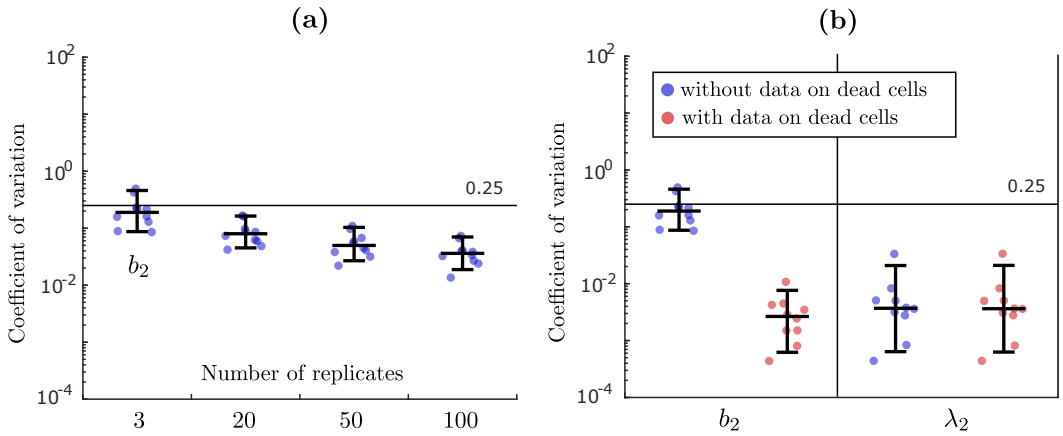

CV for b_2	Number of replicates			
	3	20	50	100
Median	0.1892	0.0798	0.0498	0.0361
10th pct	0.0870	0.0451	0.0268	0.0188
90th pct	0.4582	0.1632	0.1033	0.0701

CV	b_1		λ_1	
	●	●	●	●
Median	0.1892	0.0027	0.0037	0.0036
10th pct	0.0870	0.0006	0.0006	0.0006
90th pct	0.4582	0.0076	0.0209	0.0210

Figure 5.8: Two ways of improving the estimation accuracy for the birth rates **b** when cell number data is used. In **(a)**, we show how the estimation accuracy for the birth rate b_2 improves as the number of experimental replicates is increased. In **(b)**, we compare the estimation accuracy for the birth rate b_2 and the net birth rate λ_2 depending on whether data on the number of dead cells at each time point is included in the estimation or not.

there can be a significant benefit to developing a method tailored to sequential data.

5.8 Application: Transition between stem and non-stem cell states in SW620 colon cancer

To demonstrate the applicability of our framework, we next use it to analyze a publicly available cell fraction dataset. We use data collected by Yang et al. [168] on the dynamics between stem-like (type-1) and non-stem (type-2) cells (type-2) in SW620 colon cancer. In Yang et al. [168], the two cell types were sorted based on expression of the CD133 cell-surface antigen marker. Isolated subpopulations were expanded and phenotypic proportions were tracked for 24 days, with data collected every other day. This dataset has previously been analyzed using the CellTrans estimation method [24], which we use as a basis for comparison.

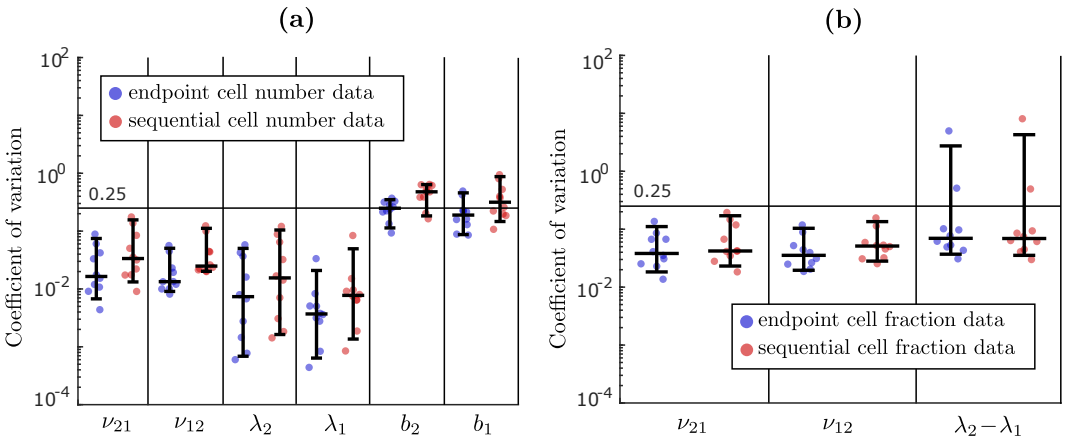

Cell number data	ν_{21}		ν_{12}		λ_2		λ_1		b_2		b_1	
Median	0.0164	0.0336	0.0134	0.0248	0.0074	0.0155	0.0037	0.0077	0.2486	0.4779	0.1892	0.3160
10th pct	0.0067	0.0132	0.0091	0.0202	0.0007	0.0016	0.0006	0.0014	0.1134	0.1831	0.0870	0.1466
90th pct	0.0745	0.1564	0.0505	0.1117	0.0503	0.1045	0.0209	0.0495	0.3497	0.6373	0.4582	0.8749

Cell fraction data	ν_{21}		ν_{12}		$\lambda_2 - \lambda_1$	
Median	0.0381	0.0420	0.0354	0.0512	0.0696	0.0691
10th pct	0.0183	0.0231	0.0195	0.0281	0.0370	0.0354
90th pct	0.1110	0.1707	0.1036	0.1351	2.7398	4.2787

Figure 5.9: Comparison of estimation error depending on whether our framework is applied to endpoint data or sequential data. The blue dots show the estimation error when endpoint data is used, i.e. when experiments from different time points are independent, and the red dots show the error when sequential data is used, i.e. when data is collected at multiple time points in the same experiment. Panel **(a)** shows the comparison for cell number data and panel **(b)** for cell fraction data. Even though our framework is derived for endpoint data, it provides reasonable estimation accuracy for sequential data.

Since data on individual experimental replicates is not available, we use data on the mean cell fraction across replicates as input to our estimation framework. We begin our analysis by considering the following two statistical models:

- **Model I:** $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, \omega^2 \mathbf{I}\big)$,
- **Model II:** $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, N_i^{-1} \mathbf{S}^{\mathbf{f}_i}(t_\ell)\big) + \mathcal{N}\big(\mathbf{0}, \omega^2 \mathbf{I}\big)$.

Model I is the simple model from (5.15), which has four parameters. As discussed in Section 5.5.3, the MLE problem for this model is a least squares problem, which can be viewed as the continuous-time version of the TRANSCOMPP estimation problem [82]

Model	w(AIC)	w(BIC)	ν_{21}	ν_{12}	$\lambda_2 - \lambda_1$
I	0.005	0.027	0.1567	0.0573	0.0841
II	0.995	0.973	0.1540	0.0570	0.0804

Table 5.3: Comparison of point estimates and model fit quality for the statistical models $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, \omega^2\mathbf{I}\big)$ (Model I) and $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)\big) + \mathcal{N}\big(\mathbf{0}, \omega^2\mathbf{I}\big)$ (Model II) applied to publicly available cell fraction data from Yang et al. [168].

(Appendix D.1). Model II is our statistical model (5.11) with $\mathbf{E}_{i,\ell}^{\text{frac}} = \omega^2\mathbf{I}$ for all i, ℓ, which has seven parameters. We note that in the implementation of our framework, the estimates for Model I are always computed as part of the estimation for Model II, since they are used both as initial guesses and for parameter scaling when solving the MLE problem (5.14) (Appendix D.3).

In Table 5.3, we show the results of the parameter estimation for the two models. To assess the quality of model fit relative to model complexity, we compute model weights based on the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) [155]. The weights are normalized to take values between 0 and 1, and they indicate the strength of evidence for each model relative to the other (Appendix D.5). We observe that whereas the point estimates for ν_{21}, ν_{12} and $\lambda_2 - \lambda_1$ are very similar for the two models, both the AIC and BIC weights indicate that Model II explains the data significantly better than Model I.

In Figure 5.10, we show 95% confidence intervals (CIs) for ν_{21}, ν_{12} and $\lambda_2 - \lambda_1$ under Model II. The CIs show that while the point estimates for ν_{21} and ν_{12} are 0.1540 and 0.0570, the true value of ν_{21} may range between 0.1110 and 0.2119, and the true value of ν_{12} may range between 0.0361 and 0.0872. Since the two CIs do not overlap, we can be confident that $\nu_{21} > \nu_{12}$, but there is significant uncertainty as to the true values of these parameters. The CI for $\lambda_2 - \lambda_1$ is even wider, which is in line with our earlier observations that this parameter is more difficult to estimate from cell fraction data than the switching rates, especially when $\lambda_2 - \lambda_1$ is relatively small in magnitude (Sections 5.6.2 and 5.7.3). In fact, the CI for $\lambda_2 - \lambda_1$ includes zero, meaning that it is plausible that $\lambda_1 = \lambda_2$.

According to Wang et al. [156], the cell cycles of the two phenotypes in the SW620 cell line are both approximately one day. As a result, the CellTrans paper [24] assumes

that the two phenotypes have the same growth rate, which is a prerequisite for using the method (Appendix D.1). We can build this assumption into our estimation by solving the MLE problem for Model II under the constraint $\lambda_2 - \lambda_1 = 0$. We refer to this as Model IIa:

- **Model IIa**: $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)\big) + \mathcal{N}\big(\mathbf{0}, \omega^2\mathbf{I}\big)$, $\lambda_2 - \lambda_1 = 0$.

Estimation results for Model IIa are shown in Figure 5.10. The assumption $\lambda_1 = \lambda_2$ has a noticeable effect on both the point estimates of ν_{21} and ν_{12} and their confidence intervals. For example, the ratio ν_{21}/ν_{12} is 2.70 under Model II, while it is 1.85 under Model IIa. In other words, switching from type-2 to type-1 happens about three times as often as switching from type-1 to type-2 under Model II, while it happens about two times as often under Model IIa. Furthermore, under Model IIa, the CI for ν_{21} is reduced by a half compared to Model II, meaning that Model IIa significantly restricts the plausible values of ν_{21}.

In the CellTrans paper [24], the same dataset is used to estimate switching probabilities of $p_{21} = 0.1030$ and $p_{12} = 0.0545$, based on a discrete-time Markov model with a time step of $\Delta t =$ one day. In the CellTrans model, type switches are synchronized between all cells in the population, and they occur at discrete time steps. In our continuous-time model, the time steps are infinitesimally small, and each cell has a certain probability of switching, proliferating and dying during each step, independently of other cells (Section 5.2.1). The CellTrans estimates become more comparable to our estimates if we shorten the time step e.g. to $\Delta t = 1/10$ day. In that case, the switching probabilities become 0.01111 and 0.00588, which translate to continuous-time rates of $\widetilde{p}_{21} = 0.1111$ and $\widetilde{p}_{12} = 0.0588$. These estimates fall at the lower limits of our CIs for ν_{21} and ν_{12} under Model IIa (Figure 5.10).

According to our estimation framework, the CellTrans method yields plausible point estimates for the switching rates, in that they fall within our 95% confidence intervals when the time step is taken to be sufficiently small. The same can be said about the estimates from Model I, which can be viewed as the continuous-time version of TRANSCOMPP. However, the preceding analysis highlights several advantages of using our framework. First of all, estimates from discrete-time models are sensitive to the length of the time step chosen, which makes their interpretation less clear. Second,

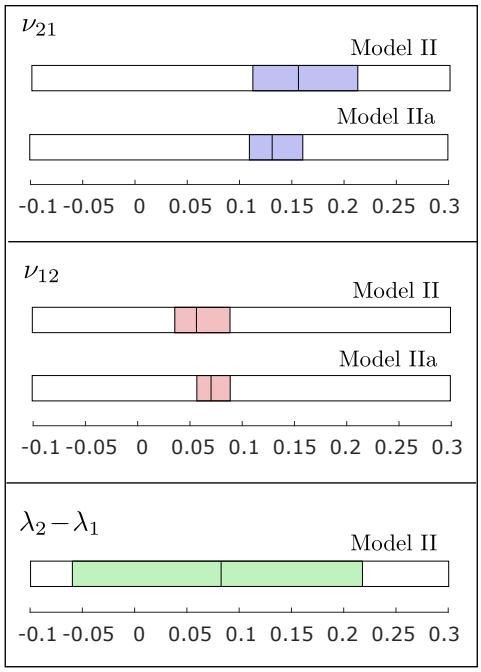

Parameter	Model	Estimate
ν_{21}	II	0.1540 CI: [0.1110,0.2119]
	IIa	0.1314 CI: [0.1100,0.1607]
ν_{12}	II	0.0570 CI: [0.0361,0.0872]
	IIa	0.0709 CI: [0.0574,0.0891]
$\lambda_2 - \lambda_1$	II	0.0804 CI: [-0.0582,0.2188]
	IIa	0

Figure 5.10: Comparison of point estimates and 95% confidence intervals for the statistical model $\mathbf{f}_{i,\ell} \sim \mathbf{p}^{\mathbf{f}_i}(t_\ell) + \mathcal{N}\big(\mathbf{0}, N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)\big) + \mathcal{N}\big(\mathbf{0}, \omega^2\mathbf{I}\big)$ (Model II) and the same model with $\lambda_2 - \lambda_1 = 0$ (Model IIa) applied to publicly available cell fraction data from Yang et al. [168].

our framework provides likelihood-based confidence intervals, which encompass all of the plausible values for the parameters. For example, the CIs for Model II reveal how uncertain the value of $\lambda_2 - \lambda_1$ is compared to ν_{21} and ν_{12}. Furthermore, the CIs for Model IIa show that even under the assumption $\lambda_1 = \lambda_2$, the true values of ν_{21} and ν_{12} may be around 50% larger than what is predicted by CellTrans. Third, our framework clearly shows the effect of simplifying assumptions such as $\lambda_1 = \lambda_2$ on the estimates of ν_{21} and ν_{12} and their CIs. As we have observed, making this assumption significantly restricts the plausible values of ν_{21} and ν_{12}, which may underestimate the true uncertainty in the estimation depending on the certainty of the claim $\lambda_1 = \lambda_2$. Finally, if data is collected on the total size of the population at each time point, our framework can easily make use of this data to produce improved parameter estimates.

5.9 Discussion

In this work, we have proposed a maximum likelihood framework for estimating the rates of cell proliferation and phenotypic switching in cancer. Our framework is novel in that it explicitly models the stochastic dynamics of cell division, cell death and phenotypic switching, it provides likelihood-based confidence intervals for the model parameters, and it enables estimation from data on the fraction of cells or the number of cells in each state.

Through a combination of theoretical and numerical investigation and application to real data, we have seen that when cell fraction data is used, the switching rates $\boldsymbol{\nu}$ may be the only parameters that can be estimated accurately, while the net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ can also be obtained when they are sufficiently large. Including information on the total size of the population at each time point yields significantly better estimates of $\boldsymbol{\lambda}^{[-1]}$, and it also enables accurate estimation of the net birth rates $\boldsymbol{\lambda}$. Finally, if enough experimental replicates are performed or if data is collected on the number of dead cells at each time point, it even becomes possible to estimates the birth rates $\mathbf{b}$ and death rates $\mathbf{d}$ accurately.

In a previous work, we discussed how knowledge of the model parameters $\boldsymbol{\nu}$, $\boldsymbol{\lambda}$, $\mathbf{b}$ can enhance our understanding of resistance evolution in cancer and inform the design of combination treatments of anti-cancer agents and epigenetic drugs [75]. Together, these parameters shape the evolution of phenotypic proportions and the tumor burden over time, each of which is relevant to the dynamics of tumor recurrence. Our current work shows that it is not possible to estimate the net birth rates $\boldsymbol{\lambda}$ or the birth rates $\mathbf{b}$ accurately from cell fraction data, it indicates what data is required to obtain these parameters, and it offers a rigorous approach to parameter estimation and uncertainty quantification once the data has been acquired.

There are several avenues for future development of the framework. First, our multitype branching process model assumes that cells are allowed to grow uninterrupted for the duration of the experiments. This does not address the effect of passaging in longer-duration experiments. One possible solution is to keep track of cell proportions and seeding densities for each passage, and treat each new passage as a new experiment. However, our framework currently assumes that initial conditions are completely

known, while uncertainty is assigned to all subsequent time points. In reality, the initial conditions are subject to measurement error, and it may become important to model this error for the case of repeated passaging.

Second, our framework currently models measurement error as an additive Gaussian noise with a general covariance matrix. We have suggested simple ways of choosing the covariance matrix both for cell number and cell fraction data, but further exploration of appropriate choices is warranted. Depending on the application, it may also be necessary to develop a more sophisticated model for the measurement error. For example, for cell number data, if the measurement error is proportional to the population size, it may be necessary to model it as a multiplicative term rather than an additive term, or to build the experimental cell counting procedure more explicitly into the model. We plan to address this in future work.

Third, we have focused on estimation from experiments started with isolated subpopulations of each phenotype, as this is a common experimental design, and we have analyzed parameter identifiability in this setting. Understanding to what extent the model parameters, or some combinations of the parameters, can be estimated from more limited data is an interesting avenue for future investigation. For example, if we are interested in estimating parameters from clinical data, the data will likely contain much less information than we have assumed here, and it will become necessary to analyze what parameters are identifiable and how identifiability can be improved, e.g. by combining data from similar patients.

Finally, we believe our framework can be useful for the design of cell line experiments aimed at deciphering the dynamics of phenotypic switching. The experimental design process includes deciding on what data to collect, when to collect it, and how many experimental replicates to perform. For example, preliminary experiments can first be conducted, from which initial parameter estimates and confidence intervals can be derived. Based on the confidence intervals, one can construct a set of likely values for the parameters, which can be used to evaluate the expected improvement in estimation accuracy depending on the experimental design, see e.g. [143]. Once good experimental designs have been identified, one can evaluate whether the expected improvement in estimation accuracy justifies the additional experimental resources. If this is the case, additional experiments can be performed and the process can be repeated. In a future

work, we plan to develop a tool for the optimal selection of experimental designs, to facilitate more efficient utilization of experimental resources.

References

[1] Guillaume Achaz. Frequency spectrum neutrality tests: one for all and all for one. *Genetics*, 183(1):249–258, 2009.

[2] George E Andrews, Richard Askey, and Ranjan Roy. *Special functions*. Number 71. Cambridge university press, 1999.

[3] Tibor Antal and PL Krapivsky. Exact solution of a two-type branching process: models of tumor progression. *J. Stat. Mech. Theory Exp.*, 2011(08):P08018, 2011.

[4] Peter Armitage and Richard Doll. The age distribution of cancer and a multi-stage theory of carcinogenesis. *British Journal of Cancer*, 8(1):1–12, 1954.

[5] Peter Armitage and Richard Doll. A two-stage theory of carcinogenesis in relation to the age distribution of human cancer. *British Journal of Cancer*, 11(2):161–169, 1957.

[6] Krishna B Athreya and Peter E Ney. *Branching processes*. Courier Corporation, 2004.

[7] Stefano Avanzini and Tibor Antal. Cancer recurrence times from a branching process model. *PLoS Comput. Biol.*, 15(11):e1007423, 2019.

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

[9] Sugandha Bhatia, James Monkman, Tony Blick, Cletus Pinto, Mark Waltham, Shivashankar H Nagaraj, and Erik W Thompson. Interrogation of phenotypic

plasticity between epithelial and mesenchymal states in breast cancer. *Journal of Clinical Medicine*, 8(6):893, 2019.

[10] Teena Bhatla, Jinhua Wang, Debra J Morrison, Elizabeth A Raetz, Michael J Burke, Patrick Brown, and William L Carroll. Epigenetic reprogramming reverses the relapse-specific gene expression signature and restores chemosensitivity in childhood b-lymphoblastic leukemia. *Blood*, 119(22):5201–5210, 2012.

[11] Lacramioara Bintu, John Yong, Yaron E Antebi, Kayla McCue, Yasuhiro Kazuki, Narumi Uno, Mitsuo Oshimura, and Michael B Elowitz. Dynamics of epigenetic regulation at the single-cell level. *Science*, 351(6274):720–724, 2016.

[12] Antara Biswas and Subhajyoti De. Drivers of dynamic intratumor heterogeneity and phenotypic plasticity. *American Journal of Physiology-Cell Physiology*, 320 (5):C750–C760, 2021.

[13] Ivan Borisov and Evgeny Metelkin. Confidence intervals by constrained optimization—an algorithm and software package for practical identifiability analysis in systems biology. *PLOS Computational Biology*, 16(12):e1008495, 2020.

[14] Ivana Bozic, Tibor Antal, Hisashi Ohtsuki, Hannah Carter, Dewey Kim, Sining Chen, Rachel Karchin, Kenneth W Kinzler, Bert Vogelstein, and Martin A Nowak. Accumulation of driver and passenger mutations during tumor progression. *Proceedings of the National Academy of Sciences*, 107(43):18545–18550, 2010.

[15] Ivana Bozic, Johannes G Reiter, Benjamin Allen, Tibor Antal, Krishnendu Chatterjee, Preya Shah, Yo Sup Moon, Amin Yaqubie, Nicole Kelly, Dung T Le, et al. Evolutionary dynamics of cancer in response to targeted combination therapy. *eLife*, 2:e00747, 2013.

[16] Ivana Bozic, Jeffrey M Gerold, and Martin A Nowak. Quantifying clonal and subclonal passenger mutations in cancer evolution. *PLoS Comput. Biol.*, 12(2): e1004731, 2016.

[17] Ivana Bozic, Chay Paterson, and Bartlomiej Waclaw. On measuring selection in cancer from subclonal mutation frequencies. *PLoS Comput. Biol.*, 15(9):e1007368, 2019.

[18] B. Braakhuis, M. Tabor, J. Kummer, C. Leemans, and R. Brakenhoff. A genetic explanation of Slaughter's concept of field cancerization evidence and clinical implications. *Cancer Research*, 63(8):1727–1730, 2003.

[19] Maury Bramson and David Griffeath. On the Williams-Bjerknes tumour growth model. II. *Math. Proc. Cambridge Philos. Soc.*, 88(2):339–357, 1980.

[20] Maury Bramson and David Griffeath. On the Williams-Bjerknes tumour growth model. I. *Ann. Probab.*, 9(2):173–185, 1981.

[21] Amy Brock, Hannah Chang, and Sui Huang. Non-genetic heterogeneity - a mutation-independent driving force for the somatic evolution of tumours. *Nature Reviews Genetics*, 10(5):336, 2009.

[22] Robert Brown, Edward Curry, Luca Magnani, Charlotte S Wilhelm-Benartzi, and Jane Borley. Poised epigenetic states and acquired drug resistance in cancer. *Nature Reviews Cancer*, 14(11):747, 2014.

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

[24] Thomas Buder, Andreas Deutsch, Michael Seifert, and Anja Voss-Böhme. Celltrans: an r package to quantify stochastic cell state transitions. *Bioinformatics and Biology Insights*, 11:1–14, 2017.

[25] Rebecca A Burrell, Nicholas McGranahan, Jiri Bartek, and Charles Swanton. The causes and consequences of genetic heterogeneity in cancer evolution. *Nature*, 501 (7467):338–345, 2013.

[26] C. Cannone. A short note on Poisson tail bounds. 2017. `http://www.cs.columbia.edu/~ccanonne/files/misc/2017-poissonconcentration.pdf`.

[27] Giulio Caravagna, Timon Heide, Marc J Williams, Luis Zapata, Daniel Nichol,

Ketevan Chkhaidze, William Cross, George D Cresswell, Benjamin Werner, Ahmet Acar, et al. Subclonal reconstruction of tumors by using machine learning and population genetics. *Nat. Genet.*, 52(9):898–907, 2020.

[28] George Casella and Roger L Berger. *Statistical inference*. Cengage Learning, 2021.

[29] Christine L Chaffer, Ines Brueckmann, Christina Scheel, Alicia J Kaestli, Paul A Wiggins, Leonardo O Rodrigues, Mary Brooks, Ferenc Reinhardt, Ying Su, Kornelia Polyak, et al. Normal and neoplastic nonstem cells can spontaneously convert to a stem-like state. *Proceedings of the National Academy of Sciences*, 108(19): 7950–7955, 2011.

[30] H. Chai and R. Brown. Field effect in cancer–an update. *Annals of Clinical & Laboratory Science*, 39(4):331–337, 2009.

[31] Nicolas Champagnat, Amaury Lambert, and Mathieu Richard. Birth and death processes with neutral mutations. *Int. J. Stoch. Anal.*, 2012, 2012.

[32] Parakrama T. Chandrasoma. Chapter 4 - histologic definition and diagnosis of epithelia in the esophagus and proximal stomach. In Parakrama T. Chandrasoma, editor, *GERD*, pages 73–107. Academic Press, 2018.

[33] Hannah H Chang, Martin Hemberg, Mauricio Barahona, Donald E Ingber, and Sui Huang. Transcriptome-wide noise controls lineage choice in mammalian progenitor cells. *Nature*, 453(7194):544, 2008.

[34] Adithya Chedere, Kishore Hari, Saurav Kumar, Annapoorni Rangarajan, and Mohit Kumar Jolly. Multi-stability and consequent phenotypic plasticity in ampk-akt double negative feedback loop in cancer cells. *Journal of Clinical Medicine*, 10(3):472, 2021.

[35] David Cheek and Tibor Antal. Mutation frequencies in a birth–death branching process. *Ann. Appl. Probab.*, 28(6):3922–3947, 2018.

[36] Juliann Chmielecki, Jasmine Foo, Geoffrey R Oxnard, Katherine Hutchinson, Kadoaki Ohashi, Romel Somwar, Lu Wang, Katherine R Amato, Maria Arcila, Martin L Sos, et al. Optimization of dosing for egfr-mutant non–small cell lung

cancer with evolutionary cancer modeling. *Science Translational Medicine*, 3(90): 90ra59–90ra59, 2011.

[37] Ariel A Cohen, Naama Geva-Zatorsky, Eran Eden, Milana Frenkel-Morgenstern, Iirina Issaeva, Alex Sigal, Ron Milo, Cellina Cohen-Saidon, Yuvalal Liron, Zvi Kam, et al. Dynamic proteomics of individual cancer cells in response to a drug. *Science*, 322(5907):1511–1516, 2008.

[38] Morgan Craig, Kamran Kaveh, Alec Woosley, Andrew S Brown, David Goldman, Elliot Eton, Ravindra M Mehta, Andrew Dhawan, Kazuya Arai, M Mamunur Rahman, et al. Cooperative adaptation to therapy (cat) confers resistance in heterogeneous non-small cell lung cancer. *PLoS Computational Biology*, 15(8), 2019.

[39] K Curtius, NA Wright, and TA Graham. An evolutionary perspective on field cancerization. *Nat. Rev. Cancer*, 18(1):19–32, 2018.

[40] Alexander Davis, Ruli Gao, and Nicholas Navin. Tumor evolution: Linear, branching, neutral or punctuated? *Biochim. Biophys. Acta Rev. Cancer*, 1867(2):151–161, 2017.

[41] Ugo Del Monte. Does the cell number 109 still really fit one gram of tumor tissue? *Cell Cycle*, 8(3):505–506, 2009.

[42] Vimalathithan Devaraj and Biplab Bose. Morphological state transition dynamics in egf-induced epithelial to mesenchymal transition. *Journal of Clinical Medicine*, 8(7):911, 2019.

[43] Neeraj Dhar and John D McKinney. Microbial phenotypic heterogeneity and antibiotic tolerance. *Current Opinion in Microbiology*, 10(1):30–38, 2007.

[44] Khanh N Dinh, Roman Jaksik, Marek Kimmel, Amaury Lambert, Simon Tavaré, et al. Statistical inference for the evolutionary history of cancer genomes. *Stat. Sci.*, 35(1):129–144, 2020.

[45] Anne Dirkse, Anna Golebiewska, Thomas Buder, Petr V Nazarov, Arnaud Muller,

Suresh Poovathingal, Nicolaas HC Brons, Sonia Leite, Nicolas Sauvageot, Dzjemma Sarkisjan, et al. Stem cell-associated heterogeneity in glioblastoma results from intrinsic tumor plasticity shaped by the microenvironment. *Nature Communications*, 10(1):1–16, 2019.

[46] R. Durrett, J. Foo, and K. Leder. Spatial Moran models, II: cancer initiation in spatially structured tissue. *J. Math. Biol.*, 72(5):1369–1400, 2016. ISSN 0303-6812. doi: 10.1007/s00285-015-0912-1. URL `https://doi.org/10.1007/s00285-015-0912-1`.

[47] Richard Durrett. Oriented percolation in two dimensions. *Ann. Probab.*, 12(4): 999–1040, 1984.

[48] Richard Durrett. *Lecture notes on particle systems and percolation.* The Wadsworth & Brooks/Cole Statistics/Probability Series. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1988. ISBN 0-534-09462-7.

[49] Richard Durrett. *Probability models for DNA sequence evolution.* Springer Science & Business Media, 2008.

[50] Richard Durrett. Branching process models of cancer. In *Branching Process Models of Cancer*, pages 1–63. Springer, 2015.

[51] Richard Durrett and Stephen Moseley. Spatial Moran models I. Stochastic tunneling in the neutral case. *Ann. Appl. Probab.*, 25(1):104–115, 2015.

[52] Rick Durrett. Ten lectures on particle systems. In *Lectures on probability theory (Saint-Flour, 1993)*, volume 1608 of *Lecture Notes in Math.*, pages 97–201. Springer, Berlin, 1995.

[53] Rick Durrett. Population genetics of neutral mutations in exponentially growing cancer cell populations. *Ann. Appl. Propab.*, 23(1):230, 2013.

[54] Rick Durrett and Iljana Zähle. On the width of hybrid zones. *Stochastic Process. Appl.*, 117(12):1751–1763, 2007.

[55] A. Dvoretzky and P. Erdös. Some problems on random walk in space. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*, pages 353–367. University of California Press, Berkeley and Los Angeles, 1951.

[56] Hariharan Easwaran, Hsing-Chen Tsai, and Stephen B Baylin. Cancer epigenetics: tumor heterogeneity, plasticity of stem-like states, and drug resistance. *Molecular Cell*, 54(5):716–727, 2014.

[57] Manel Esteller. Epigenetics in cancer. *New England Journal of Medicine*, 358(11): 1148–1159, 2008.

[58] Justin C Fay and Chung-I Wu. Hitchhiking under positive darwinian selection. *Genetics*, 155(3):1405–1413, 2000.

[59] Jian-Ming Feng, Ze-Hong Miao, Yi Jiang, Yi Chen, Jia-Xin Li, Lin-Jiang Tong, Jin Zhang, Yi-Ran Huang, and Jian Ding. Characterization of the conversion between cd133+ and cd133-cells in colon cancer sw620 cell line. *Cancer Biology & Therapy*, 13(14):1396–1406, 2012.

[60] Samuel M Fischer and Mark A Lewis. A robust and efficient algorithm to find profile likelihood confidence intervals. *Statistics and Computing*, 31(4):1–17, 2021.

[61] Kjersti Flatmark, Ragnhild V Nome, Sigurd Folkvord, Åse Bratland, Heidi Rasmussen, Mali Strand Ellefsen, Øystein Fodstad, and Anne Hansen Ree. Radiosensitization of colorectal carcinoma cell lines by histone deacetylase inhibition. *Radiation Oncology*, 1(1):25, 2006.

[62] William A Flavahan, Elizabeth Gaskell, and Bradley E Bernstein. Epigenetic plasticity and the hallmarks of cancer. *Science*, 357(6348):eaal2380, 2017.

[63] Jasmine Foo, Kevin Leder, and Marc D. Ryser. Multifocality and recurrence risk: a quantitative model of field cancerization. *J. Theoret. Biol.*, 355:170–184, 2014.

[64] Jasmine Foo, Kevin Leder, and Junfeng Zhu. Escape times for branching processes with random mutational fitness effects. *Stochastic Process. Appl.*, 124(11):3661–3697, 2014.

[65] Jasmine Foo, Einar Bjarki Gunnarsson, Kevin Leder, and Kathleen Storey. Spread of premalignant mutant clones and cancer initiation in multilayered tissue. *arXiv preprint arXiv:2007.03366*, 2020.

[66] Jasmine Foo, Kevin Leder, and Jason Schweinsberg. Mutation timing in a spatial model of evolution. *Stochastic Process. Appl.*, 130(10):6388–6413, 2020.

[67] Yun-Xin Fu and Wen-Hsiung Li. Statistical tests of neutrality of mutations. *Genetics*, 133(3):693–709, 1993.

[68] Bernadett Gaál, Jonathan W Pitchford, and A Jamie Wood. Exact results for the evolution of stochastic switching in variable asymmetric environments. *Genetics*, 184(4):1113–1119, 2010.

[69] K. Geboes. Squamous mucosa and reflux. `http://www.hon.ch/OESO/books/Vol_3_Eso_Mucosa/Articles/ART004.HTML`, 1994.

[70] Yue Geng, Siddarth Chandrasekaran, Sivaprakash Agastin, Jiahe Li, and Michael R King. Dynamic switch between two adhesion phenotypes in colorectal cancer cells. *Cellular and Molecular Bioengineering*, 7(1):35–44, 2014.

[71] Alan George. Nested dissection of a regular finite element mesh. *SIAM J. Numer. Anal.*, 10(2):345–363, 1973.

[72] Aaron Goldman, Biswanath Majumder, Andrew Dhawan, Sudharshan Ravi, David Goldman, Mohammad Kohandel, Pradip K Majumder, and Shiladitya Sengupta. Temporally sequenced anticancer drugs overcome adaptive resistance by targeting a vulnerable chemotherapy-induced phenotypic transition. *Nature Communications*, 6:6139, 2015.

[73] Steven D Gore, Stephen Baylin, Elizabeth Sugar, Hetty Carraway, Carole B Miller, Michael Carducci, Michael Grever, Oliver Galm, Tianna Dauses, Judith E Karp, et al. Combined DNA methyltransferase and histone deacetylase inhibition in the treatment of myeloid neoplasms. *Cancer Research*, 66(12):6361–6369, 2006.

[74] Robert C Griffiths and Anthony G Pakes. An infinite-alleles version of the simple branching process. *Adv. Appl. Propab.*, pages 489–524, 1988.

[75] Einar Bjarki Gunnarsson, Subhajyoti De, Kevin Leder, and Jasmine Foo. Understanding the role of phenotypic switching in cancer drug resistance. *Journal of Theoretical Biology*, 490:110162, 2020.

[76] Einar Bjarki Gunnarsson, Kevin Leder, and Jasmine Foo. Exact site frequency spectra of neutrally evolving tumors: A transition between power laws reveals a signature of cell viability. *Theoretical Population Biology*, 142:67–90, 2021.

[77] Piyush B Gupta, Christine M Fillmore, Guozhi Jiang, Sagi D Shapira, Kai Tao, Charlotte Kuperwasser, and Eric S Lander. Stochastic state transitions give rise to phenotypic equilibrium in populations of cancer cells. *Cell*, 146(4):633–644, 2011.

[78] Aaron N Hata, Matthew J Niederst, Hannah L Archibald, Maria Gomez-Caraballo, Faria M Siddiqui, et al. Tumor cells can follow distinct evolutionary paths to become resistant to epidermal growth factor receptor inhibition. *Nature Medicine*, 22(3):262–269, 2016.

[79] Timon Heide, Luis Zapata, Marc J Williams, Benjamin Werner, Giulio Caravagna, Chris P Barnes, Trevor A Graham, and Andrea Sottoriva. Reply to 'neutral tumor evolution?'. *Nat. Genet.*, 50(12):1633–1637, 2018.

[80] NIH National Cancer Institute. SEER training cancer classification. `https://training.seer.cancer.gov/disease/categories/classification.html`.

[81] Yoh Iwasa, Martin A Nowak, and Franziska Michor. Evolution of resistance during clonal expansion. *Genetics*, 172(4):2557–2566, 2006.

[82] N Suhas Jagannathan, Mario O Ihsan, Xiao Xuan Kin, Roy E Welsch, Marie-Véronique Clément, and Lisa Tucker-Kellogg. Transcompp: understanding phenotypic plasticity by estimating markov transition rates for cell state transitions. *Bioinformatics*, 36(9):2813–2820, 2020.

[83] Mohit Kumar Jolly, Prakash Kulkarni, Keith Weninger, John Orban, and Herbert Levine. Phenotypic plasticity, bet-hedging, and androgen independence in prostate cancer: Role of non-genetic heterogeneity. *Frontiers in Oncology*, 8:50, 2018.

[84] Peter A Jones and Stephen B Baylin. The epigenomics of cancer. *Cell*, 128(4): 683–692, 2007.

[85] Siân Jones, Wei-dong Chen, Giovanni Parmigiani, Frank Diehl, Niko Beerenwinkel, Tibor Antal, Arne Traulsen, Martin A Nowak, Christopher Siegel, Victor E Velculescu, et al. Comparative lesion sequencing provides insights into tumor evolution. *Proc. Natl. Acad. Sci. USA*, 105(11):4283–4288, 2008.

[86] Nicole Vincent Jordan, Aditya Bardia, Ben S Wittner, Cyril Benes, Matteo Ligorio, Yu Zheng, Min Yu, Tilak K Sundaresan, Joseph A Licausi, Rushil Desai, et al. Her2 expression identifies dynamic functional states within circulating breast cancer cells. *Nature*, 537(7618):102–106, 2016.

[87] Rosalyn A Juergens, John Wrangle, Frank P Vendetti, Sara C Murphy, Ming Zhao, Barbara Coleman, Rosa Sebree, Kristen Rodgers, Craig M Hooker, Noreli Franco, et al. Combination epigenetic therapy has efficacy in patients with refractory advanced non–small cell lung cancer. *Cancer Discovery*, 1(7):598–607, 2011.

[88] Peter Keller and Tibor Antal. Mutant number distribution in an exponentially growing population. *J. Stat. Mech. Theory Exp.*, 2015(1):P01011, 2015.

[89] David A Kessler and Herbert Levine. Large population solution of the stochastic Luria–Delbrück evolution model. *Proc. Natl. Acad. Sci. USA*, 110(29):11682–11687, 2013.

[90] David A Kessler and Herbert Levine. Scaling solution in the large population limit of the general asymmetric stochastic Luria–Delbrück evolution process. *J. Stat. Phys.*, 158(4):783–805, 2015.

[91] Motoo Kimura. Genetic variability maintained in a finite population due to mutational production of neutral and nearly neutral isoalleles. *Genet. Res.*, 11(3): 247–270, 1968.

[92] Motoo Kimura. The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics*, 61(4):893, 1969.

[93] John FC Kingman. On the genealogy of large populations. *J. Appl. Propab.*, 19: 27–43, 1982.

[94] John Frank Charles Kingman. The coalescent. *Stoch. Process. Their Appl.*, 13 (3):235–248, 1982.

[95] Birgit Knoechel, Justine E Roderick, Kaylyn E Williamson, Jiang Zhu, Jens G Lohr, Matthew J Cotton, Shawn M Gillespie, Daniel Fernandez, Manching Ku, Hongfang Wang, et al. An epigenetic mechanism of resistance to targeted therapy in t cell acute lymphoblastic leukemia. *Nature Genetics*, 46(4):364, 2014.

[96] A. Knudson. Two genetic hits (more or less) to cancer. *Nature Reviews Cancer*, 1(2):157–161, 2001.

[97] Alfred G Knudson. Mutation and cancer: statistical study of retinoblastoma. *Proceedings of the National Academy of Sciences*, 68(4):820–823, 1971.

[98] Natalia L. Komarova. Spatial stochastic models for cancer initiation and progression. *Bull. Math. Biol.*, 68(7):1573–1599, 2006. ISSN 0092-8240. doi: 10.1007/s11538-005-9046-8. URL `https://doi.org/10.1007/s11538-005-9046-8`.

[99] Natalia L Komarova, Lin Wu, and Pierre Baldi. The fixed-size luria–delbruck model with a nonzero death rate. *Math. Biosci.*, 210(1):253–290, 2007.

[100] Jack Kuipers, Katharina Jahn, Benjamin J Raphael, and Niko Beerenwinkel. Single-cell sequencing data reveal widespread recurrence and loss of mutational hits in the life histories of tumors. *Genome Res.*, 27(11):1885–1894, 2017.

[101] Amaury Lambert. The allelic partition for coalescent point processes. *Markov Process. Relat. Fields*, 15(3):359–386, 2009.

[102] Gregory F. Lawler and Vlada Limic. *Random walk: a modern introduction*, volume 123 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010. ISBN 978-0-521-51918-2. doi: 10.1017/CBO9780511750854. URL `https://doi.org/10.1017/CBO9780511750854`.

[103] Kevin Leder, Ken Pitter, Quincey LaPlant, Dolores Hambardzumyan, Brian D Ross, Timothy A Chan, Eric C Holland, and Franziska Michor. Mathematical

modeling of pdgf-driven glioblastoma reveals optimized radiation dosing schedules. *Cell*, 156(3):603–616, 2014.

[104] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. ISBN 978-0-8218-4739-8. doi: 10.1090/mbk/058. URL `https://doi.org/10.1090/mbk/058`. With a chapter by James G. Propp and David B. Wilson.

[105] Xin Li and Dave Thirumalai. A mathematical model for phenotypic heterogeneity in breast cancer with implications for therapeutic strategies. *bioRxiv*, 2021.

[106] Brian B Liau, Cem Sievers, Laura K Donohue, Shawn M Gillespie, William A Flavahan, Tyler E Miller, Andrew S Venteicher, Christine H Hebert, Christopher D Carey, Scott J Rodig, et al. Adaptive chromatin remodeling drives glioblastoma stem cell plasticity and drug tolerance. *Cell Stem Cell*, 20(2):233–246, 2017.

[107] Thomas M. Liggett. *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin, 2005. ISBN 3-540-22617-6. doi: 10.1007/b138374. URL `https://doi.org/10.1007/b138374`. Reprint of the 1985 original.

[108] Shaoping Ling, Zheng Hu, Zuyu Yang, Fang Yang, Yawei Li, Pei Lin, Ke Chen, Lili Dong, Lihua Cao, Yong Tao, et al. Extremely high genetic diversity in a single tumor points to prevalence of non-darwinian cell evolution. *Proc. Natl. Acad. Sci. USA*, 112(47):E6496–E6505, 2015.

[109] Andriy Marusyk, Vanessa Almendro, and Kornelia Polyak. Intra-tumour heterogeneity: a looking glass for cancer? *Nature Reviews Cancer*, 12(5):323, 2012.

[110] Thomas O McDonald, Shaon Chakrabarti, and Franziska Michor. Currently available bulk sequencing data do not necessarily support a model of neutral tumor evolution. *Nat. Genet.*, 50(12):1620–1623, 2018.

[111] Nicholas McGranahan and Charles Swanton. Clonal heterogeneity and tumor evolution: past, present, and the future. *Cell*, 168(4):613–628, 2017.

[112] Lauren MF Merlo, John W Pepper, Brian J Reid, and Carlo C Maley. Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*, 6(12):924, 2006.

[113] Franziska Michor, Martin A Nowak, and Yoh Iwasa. Evolution of resistance to cancer therapy. *Current Pharmaceutical Design*, 12(3):261–271, 2006.

[114] Richard L Momparler and J Ayoub. Potential of 5-aza-2-deoxycytidine (decitabine) a potent inhibitor of DNA methylation for therapy of advanced non-small cell lung cancer. *Lung Cancer*, 34:111–115, 2001.

[115] J Müller, BA Hense, TM Fuchs, M Utz, and Ch Pötzsche. Bet-hedging in stochastically switching environments. *Journal of Theoretical Biology*, 336:144–157, 2013.

[116] Michael C Neale and Michael B Miller. The use of likelihood-based confidence intervals in genetic models. *Behavior Genetics*, 27(2):113–120, 1997.

[117] Cyril Neftel, Julie Laffy, Mariella G Filbin, Toshiro Hara, Marni E Shore, Gilbert J Rahme, Alyssa R Richman, Dana Silverbush, McKenzie L Shaw, Christine M Hebert, et al. An integrative model of cellular states, plasticity, and genetics for glioblastoma. *Cell*, 178(4):835–849, 2019.

[118] Daniel Nichol, Mark Robertson-Tessi, Peter Jeavons, and Alexander RA Anderson. Stochasticity in the genotype-phenotype map: implications for the robustness and persistence of bet-hedging. *Genetics*, pages genetics–116, 2016.

[119] Mario Niepel, Sabrina L Spencer, and Peter K Sorger. Non-genetic cell-to-cell variability and the consequences for pharmacology. *Current Opinion in Chemical Biology*, 13(5-6):556–561, 2009.

[120] Thomas M Norman, Nathan D Lord, Johan Paulsson, and Richard Losick. Stochastic switching of cell fate in microbes. *Annual Review of Microbiology*, 69:381–403, 2015.

[121] M. A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, Cambridge, MA, 2006.

[122] Peter C Nowell. The clonal evolution of tumor cell populations. *Science*, 194 (4260):23–28, 1976.

[123] Neil O'Connell. Yule process approximation for the skeleton of a branching process. *J. Appl. Propab.*, 30(3):725–729, 1993.

[124] Hisashi Ohtsuki and Hideki Innan. Forward and backward evolutionary processes and allele frequency spectrum in a cancer cell population. *Theor. Popul. Biol.*, 117:43–50, 2017.

[125] Anthony G Pakes. An infinite alleles version of the markov branching process. *J. Aust. Math. Soc.*, 46(1):146–169, 1989.

[126] Angela Oliveira Pisco, Amy Brock, Joseph Zhou, Andreas Moor, Mitra Mojtahedi, Dean Jackson, and Sui Huang. Non-Darwinian dynamics in therapy-induced cancer drug resistance. *Nature Communications*, 4:2467, 2013.

[127] Elsa Quintana, Mark Shackleton, Hannah R Foster, Douglas R Fullen, Michael S Sabel, Timothy M Johnson, and Sean J Morrison. Phenotypic heterogeneity among tumorigenic melanoma cells from patients that is reversible and not hierarchically organized. *Cancer Cell*, 18(5):510–523, 2010.

[128] Michael Ramirez, Satwik Rajaram, Robert J Steininger, Daria Osipchuk, Maike A Roth, Leanna S Morinishi, Louise Evans, Weiyue Ji, Chien-Hsiang Hsu, Kevin Thurley, et al. Diverse drug-resistance mechanisms can emerge from drug-tolerant cancer persister cells. *Nature Communications*, 7:10690, 2016.

[129] Andreas Raue, Clemens Kreutz, Thomas Maiwald, Julie Bachmann, Marcel Schilling, Ursula Klingmüller, and Jens Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009.

[130] DA Rew and GD Wilson. Cell production rates in human tissues and tumours and their significance. part ii: clinical data. *Eur. J. Surg. Oncol.*, 26(4):405–417, 2000.

[131] José Reyes and Galit Lahav. Leveraging and coping with uncertainty in the response of individual cells to therapy. *Current Opinion in Biotechnology*, 51: 109–115, 2018.

[132] Daniel Richardson. Random growth in a tessellation. *Proc. Cambridge Philos. Soc.*, 74:515–528, 1973. ISSN 0008-1981. doi: 10.1017/s0305004100077288. URL `https://doi.org/10.1017/s0305004100077288`.

[133] Alexander Roesch, Mizuho Fukunaga-Kalabis, Elizabeth C Schmidt, Susan E Zabierowski, Patricia A Brafford, Adina Vultur, Devraj Basu, Phyllis Gimotty, Thomas Vogt, and Meenhard Herlyn. A temporarily distinct subpopulation of slow-cycling melanoma cells is required for continuous tumor growth. *Cell*, 141 (4):583–594, 2010.

[134] Alexander Roesch, Adina Vultur, Ivan Bogeski, Huan Wang, Katharina M Zimmermann, David Speicher, Christina Körbel, Matthias W Laschke, Phyllis A Gimotty, Stephan E Philipp, et al. Overcoming intrinsic multidrug resistance in melanoma by blocking the mitochondrial respiratory chain of slow-cycling jarid1bhigh cells. *Cancer Cell*, 23(6):811–825, 2013.

[135] Thomas J Rothenberg. Identification in parametric models. *Econometrica: Journal of the Econometric Society*, pages 577–591, 1971.

[136] Ravi Salgia and Prakash Kulkarni. The genetic/non-genetic duality of drug 'resistance'in cancer. *Trends in Cancer*, 4(2):110–118, 2018.

[137] Amélie Clémentine Seghers, Sofie Wilgenhof, Céleste Lebbé, and Bart Neyns. Successful rechallenge in two patients with BRAF-V600-mutant melanoma who experienced previous progression during treatment with a selective BRAF inhibitor. *Melanoma Research*, 22(6):466–472, 2012.

[138] Sydney M Shaffer, Margaret C Dunagin, Stefan R Torborg, Eduardo A Torre, Benjamin Emert, Clemens Krepler, Marilda Beqiri, Katrin Sproesser, Patricia A Brafford, Min Xiao, et al. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658):431, 2017.

[139] Sreenath V Sharma, Diana Y Lee, Bihua Li, Margaret P Quinlan, Fumiyuki Takahashi, Shyamala Maheswaran, Ultan McDermott, Nancy Azizian, Lee Zou, Michael A Fischbach, et al. A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1):69–80, 2010.

[140] Alex Sigal, Ron Milo, Ariel Cohen, Naama Geva-Zatorsky, Yael Klein, Yuvalal Liron, Nitzan Rosenfeld, Tamar Danon, Natalie Perzov, and Uri Alon. Variability and memory of protein levels in human cells. *Nature*, 444(7119):643, 2006.

[141] Mikhail V Simkin and Vwani P Roychowdhury. Re-inventing willis. *Phys. Rep.*, 502(1):1–35, 2011.

[142] Andrea Sottoriva, Haeyoun Kang, Zhicheng Ma, Trevor A Graham, Matthew P Salomon, Junsong Zhao, Paul Marjoram, Kimberly Siegmund, Michael F Press, Darryl Shibata, et al. A big bang model of human colorectal tumor growth. *Nat. Genet.*, 47(3):209–216, 2015.

[143] Bernhard Steiert, Andreas Raue, Jens Timmer, and Clemens Kreutz. Experimental design for parameter estimation of gene regulatory networks. *PloS One*, 7(7): e40052, 2012.

[144] Yapeng Su, Wei Wei, Lidia Robert, Min Xue, Jennifer Tsoi, Angel Garcia-Diaz, Blanca Homet Moreno, Jungwoo Kim, Rachel H Ng, Jihoon W Lee, et al. Single-cell analysis resolves the cell state transition and signaling dynamics associated with melanoma drug-induced resistance. *Proceedings of the National Academy of Sciences*, 114(52):13679–13684, 2017.

[145] Chong Sun, Liqin Wang, Sidong Huang, Guus JJE Heynen, Anirudh Prahallad, Caroline Robert, John Haanen, Christian Blank, Jelle Wesseling, Stefan M Willems, et al. Reversible and adaptive resistance to BRAF (V600E) inhibition in melanoma. *Nature*, 508(7494):118, 2014.

[146] Hyuna Sung, Jacques Ferlay, Rebecca L Siegel, Mathieu Laversanne, Isabelle Soerjomataram, Ahmedin Jemal, and Freddie Bray. Global cancer statistics 2020: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer Journal for Clinicians*, 71(3):209–249, 2021.

[147] Fumio Tajima. Statistical method for testing the neutral mutation hypothesis by dna polymorphism. *Genetics*, 123(3):585–595, 1989.

[148] Maxime Tarabichi, Iñigo Martincorena, Moritz Gerstung, Armand M Leroi, Florian Markowetz, Paul T Spellman, Quaid D Morris, Ole Christian Lingjærde, David C Wedge, and Peter Van Loo. Neutral tumor evolution? *Nat. Genet.*, 50 (12):1630–1633, 2018.

[149] Cristian Tomasetti, Bert Vogelstein, and Giovanni Parmigiani. Half or more of the somatic mutations in cancers of self-renewing tissues originate prior to tumor initiation. *Proc. Natl. Acad. Sci. USA*, 110(6):1999–2004, 2013.

[150] Samra Turajlic, Andrea Sottoriva, Trevor Graham, and Charles Swanton. Resolving genetic heterogeneity in cancer. *Nat. Rev. Genet.*, 20(7):404–416, 2019.

[151] Subramanian Venkatesan and Charles Swanton. Tumor evolutionary principles: how intratumor heterogeneity influences cancer treatment and outcome. *Am. Soc. Clin. Oncol. Educ. Book*, 36:e141–e149, 2016.

[152] DJ Venzon and SH Moolgavkar. A method for computing profile-likelihood-based confidence intervals. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 37(1):87–94, 1988.

[153] Kavya Vipparthi, Kishore Hari, Priyanka Chakraborty, Subhashis Ghosh, Ankit Patel, Arnab Ghosh, Nidhan K Biswas, Rajeev Sharan, Pattatheyil Arun, Mohit Kumar Jolly, et al. Spontaneous emergence of non-convertible cell states with CD24-high phenotype results in phenotypic heterogeneity that associates with poor prognosis in oral cancer. *bioRxiv*, 2021.

[154] Bert Vogelstein, Nickolas Papadopoulos, Victor E Velculescu, Shibin Zhou, Luis A Diaz, and Kenneth W Kinzler. Cancer genome landscapes. *Science*, 339(6127): 1546–1558, 2013.

[155] Eric-Jan Wagenmakers and Simon Farrell. Aic model selection using akaike weights. *Psychonomic bulletin & review*, 11(1):192–196, 2004.

[156] Weikang Wang, Yi Quan, Qibin Fu, Yu Liu, Ying Liang, Jingwen Wu, Gen Yang, Chunxiong Luo, Qi Ouyang, and Yugang Wang. Dynamics between cancer cell subpopulations reveals a model coordinating with both hierarchical and stochastic concepts. *PloS One*, 9(1):e84654, 2014.

[157] GA Watterson. On the number of segregating sites in genetical models without recombination. *Theor. Popul. Biol.*, 7(2):256–276, 1975.

[158] Benjamin Werner, Marc J Williams, Chris P Barnes, Trevor A Graham, and Andrea Sottoriva. Reply to 'currently available bulk sequencing data do not necessarily support a model of neutral tumor evolution'. *Nat. Genet.*, 50(12):1624–1626, 2018.

[159] Benjamin Werner, Jack Case, Marc J Williams, Ketevan Chkhaidze, Daniel Temko, Javier Fernández-Mateos, George D Cresswell, Daniel Nichol, William Cross, Inmaculada Spiteri, et al. Measuring single cell divisions in human tissues from multi-region sequencing data. *Nat. Commun.*, 11(1):1–9, 2020.

[160] Marc J Williams, Benjamin Werner, Chris P Barnes, Trevor A Graham, and Andrea Sottoriva. Identification of neutral tumor evolution across cancer types. *Nat. Genet.*, 48(3):238, 2016.

[161] Marc J Williams, Benjamin Werner, Timon Heide, Christina Curtis, Chris P Barnes, Andrea Sottoriva, and Trevor A Graham. Quantification of subclonal selection in cancer from bulk sequencing data. *Nat. Genet.*, 50(6):895–903, 2018.

[162] T. Williams and R. Bjerknes. Stochastic model for abnormal clone spread through epithelial basal layer. *Nature*, 236:19–21, 1972.

[163] Roel H Wilting and Jan-Hermen Dannenberg. Epigenetic mechanisms in tumorigenesis, tumor cell heterogeneity and drug resistance. *Drug Resistance Updates*, 15(1-2):21–38, 2012.

[164] Dominik Wodarz and Natalia L. Komarova. *Dynamics of cancer: Mathematical foundations of oncology.* World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2014.

[165] Christopher L Woodcock. Chromatin architecture. *Current Opinion in Structural Biology*, 16(2):213–220, 2006.

[166] Xiaowei Wu and Marek Kimmel. Modeling neutral evolution using an infinite-allele markov branching process. *Int. J. Stoch. Anal.*, 2013:1–10, 2013.

[167] Andrei Y Yakovlev and Nikolay M Yanev. Relative frequencies in multitype branching processes. *The Annals of Applied Probability*, 19(1):1–14, 2009.

[168] G Yang, Y Quan, W Wang, Q Fu, J Wu, T Mei, J Li, Y Tang, C Luo, Q Ouyang, et al. Dynamic equilibrium between cancer stem cells and non-stem cancer cells in human SW620 and MCF-7 cancer cell populations. *British Journal of Cancer*, 106(9):1512–1519, 2012.

[169] George Udny Yule. II. —-A mathematical theory of evolution, based on the conclusions of Dr. J.C. Willis, F.R. S. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.*, 213(402-410):21–87, 1925.

[170] Kai Zeng, Yun-Xin Fu, Suhua Shi, and Chung-I Wu. Statistical tests for detecting positive selection by utilizing high-frequency variants. *Genetics*, 174(3):1431–1439, 2006.

[171] Joseph Xu Zhou, Angela Oliveira Pisco, Hong Qian, and Sui Huang. Nonequilibrium population dynamics of phenotype conversion of cancer cells. *PloS One*, 9 (12):e110714, 2014.

Appendix A

Appendix to Chapter 2

A.1 Proofs of lemmas

A.1.1 Proof of Lemma 2.1

Lemma A.1 (LCLT on $\mathbb{Z}^d \times \mathbb{Z}_w$)**.** *Let $(S_n)_{n\geq 0}$ be the discrete-time SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ with $S_0 = 0$. Set $b_1 := 2$, $b_w := 2$ if $w > 1$ is even and $b_w := 1$ if $w > 1$ is odd, and define*

$$p_{w,d} := \begin{cases} 1, & w = 1, \\ 2d/(2d+1), & w = 2, \\ d/(d+1), & w > 2, \end{cases}$$

as the probability that $(S_n)_{n\geq 0}$ takes a step in the $\mathbb{Z}^d$-direction, with $p_{w,2} = p_w$ as defined in (2.6). Then, for any $n \geq 1$ and $x \in \mathbb{Z}^d \times \mathbb{Z}_w$ so that $\mathbb{P}(S_n = x) > 0$,

$$n^{d/2}\mathbb{P}(S_n = x) = (b_w/w)\big(d/(2\pi p_{w,d})\big)^{d/2} + o(1).$$

Proof. For $w = 1$, the SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ is equivalent to the SSRW on $\mathbb{Z}^d$, and it suffices to refer to the LCLT on $\mathbb{Z}^d$, given by (A.2) below. We can therefore assume that $w > 1$.

Let $\hat{N}(n)$ denote the number of steps taken along the first d dimensions of $\mathbb{Z}^d \times \mathbb{Z}_w$ by time n. Conditional on $\hat{N}(n) = k$, we can decompose S_n as $S_n = (\hat{S}_k, S^w_{n-k})$, where $(\hat{S}_j)_{j\geq 0}$ and $(S^w_j)_{j\geq 0}$ are independent SSRWs on $\mathbb{Z}^d$ and $\mathbb{Z}_w$, respectively. In other

words, for any $x \in \mathbb{Z}^d \times \mathbb{Z}_w$ with $x = (\hat{x}, x_{d+1})$,

$$\mathbb{P}(S_n = x | \hat{N}(n) = k) = \mathbb{P}(\hat{S}_k = \hat{x})\mathbb{P}(S^w_{n-k} = x_{d+1}).$$

Thus, by conditioning on the value of $\hat{N}(n)$, we can analyze the large-n asymptotics of S_n by analyzing the two random walks $(\hat{S}_j)_{j\geq 0}$ and $(S^w_j)_{j\geq 0}$ separately.

We start with the latter walk $(S^w_j)_{j\geq 0}$ on $\mathbb{Z}_w$. Note first that $(S^w_j)_{j\geq 0}$ is aperiodic if $w > 1$ is odd, while it is periodic with period 2 if $w > 1$ is even. For the w odd case, the transition matrix P_w for $(S^w_j)_{j\geq 0}$ is irreducible and aperiodic, and the uniform distribution on $\mathbb{Z}_w$ is stationary for P_w, so by Theorem 4.9 of [104], there exist constants $\gamma_1 = \gamma_1(w) \in (0,1)$ and $C_1 = C_1(w) > 0$ so that

$$|\mathbb{P}(S^w_j = x_{d+1}) - 1/w| \leq C_1 \gamma_1^j, \quad x_{d+1} \in \mathbb{Z}_w.$$

For the w even case, by rearranging the state space into odds and evens, we can write P_w^2 as a block diagonal matrix consisting of two identical irreducible, aperiodic blocks Q_w of dimension $(w/2) \times (w/2)$. If $S^w_0 = 0$ and x_{d+1} is even, we can apply the same theorem as above to get constants $\gamma_2 = \gamma_2(w) \in (0,1)$ and $C_2 = C_2(w) > 0$ so that

$$|\mathbb{P}(S^w_{2j} = x_{d+1}) - 2/w| \leq C_2 \gamma_2^j = C_2 \big(\gamma_2^{1/2}\big)^{2j}.$$

If x_{d+1} is odd, we condition on the first step and obtain as before

$$|\mathbb{P}(S^w_{2j+1} = x_{d+1}) - 2/w| \leq C_2 \gamma_2^j = \big(C_2 \gamma_2^{-1/2}\big)\big(\gamma_2^{1/2}\big)^{2j+1}.$$

Recall that $b_w = 1$ for $w > 1$ odd and $b_w = 2$ for $w > 1$ even. Combining the above observations, we see that for the SSRW $(S^w_j)_{j\geq 0}$ on $\mathbb{Z}_w$, there exist $\gamma = \gamma(w) \in (0,1)$ and $C = C(w) > 0$ so that for each $x_{d+1} \in \mathbb{Z}_w$ with $\mathbb{P}(S^w_j = x_{d+1}) > 0$,

$$|\mathbb{P}(S^w_j = x_{d+1}) - (b_w/w)| \leq C\gamma^j. \tag{A.1}$$

We next turn to the random walk $(\hat{S}_j)_{j\geq 0}$ on $\mathbb{Z}^d$. For $\hat{x} \in \mathbb{Z}^d$ and $j \geq 1$, set

$$p_j(\hat{x}) := \big(d/(2\pi j)\big)^{d/2} \exp\big(-d\|\hat{x}\|^2/(2j)\big),$$

where $\|\cdot\|$ is the Euclidean norm on $\mathbb{Z}^d$. By Theorem 2.1.3 of [102], there exists $c > 0$ so that for all $\hat{x} \in \mathbb{Z}^d$,

$$|\mathbb{P}(\hat{S}_j = \hat{x}) + \mathbb{P}(\hat{S}_{j+1} = \hat{x}) - 2p_j(\hat{x})| \leq c/j^{(d+2)/2}.$$

Since $(\hat{S}_j)_{j\geq 0}$ is periodic with period 2, we obtain for j and $\hat{x}$ such that $\mathbb{P}(\hat{S}_j = \hat{x}) > 0$,

$$j^{d/2}\mathbb{P}(\hat{S}_j = \hat{x}) = 2\big(d/(2\pi)\big)^{d/2} + O(1/j). \tag{A.2}$$

We finally establish bounds on $\hat{N}(n)$, the number of steps the SSRW on $\mathbb{Z}^d \times \mathbb{Z}_w$ takes along the first d dimensions by time n. Since $\hat{N}(n)$ is binomially distributed with success probability $p_{w,d}$, we obtain by Hoeffding's inequality for any $\nu > 0$,

$$\mathbb{P}(|\hat{N}(n) - np_{w,d}| > n^{1/2+\nu}) \leq 2\exp(-2n^{2\nu}). \tag{A.3}$$

For each n and each $\nu \in (0, 1/2)$, define the neighborhoods

$$A_n(\nu) := \{1 \leq k \leq n : |k - np_{w,d}| \leq n^{1/2+\nu}\},$$

and note that by (A.3), $\hat{N}(n)$ takes values in $A_n(\nu)$ with high probability for n large.

We are now ready to carry out the main calculations. Fix $\nu \in (0, 1/2)$ and note first that

$$\begin{aligned} n^{d/2}\mathbb{P}(S_n = x) &= n^{d/2}\mathbb{P}(S_n = x|\hat{N}(n) \in A_n(\nu))\mathbb{P}(\hat{N}(n) \in A_n(\nu)) \\ &\quad + n^{d/2}\mathbb{P}(S_n = x|\hat{N}(n) \notin A_n(\nu))\mathbb{P}(\hat{N}(n) \notin A_n(\nu)). \end{aligned}$$

Write $x = (\hat{x}, x_{d+1})$. Since $\mathbb{P}(\hat{N}(n) \notin A_n(\nu)) \leq 2\exp(-2n^{2\nu})$ by (A.3), and $n^{d/2} = o(\exp(2n^{2\nu}))$, it suffices to study the large-n asymptotics of

$$\begin{aligned} &n^{d/2}\mathbb{P}(S_n = x|\hat{N}(n) \in A_n(\nu)) \\ &= n^{d/2}\textstyle\sum_{k\in A_n(\nu)} \mathbb{P}(\hat{S}_k = \hat{x})\mathbb{P}(S^w_{n-k} = x_{d+1})\mathbb{P}(\hat{N}(n) = k|\hat{N}(n) \in A_n(\nu)). \end{aligned}$$

For n and x so that $\mathbb{P}(S_n = x) > 0$, we can assume without loss of generality that $\mathbb{P}(\hat{S}_k = \hat{x}) > 0$ and $\mathbb{P}(S^w_{n-k} = x_{d+1}) > 0$ for $k \in A_n(\nu)$ even and n sufficiently large.

Define

$$A_n^2(\nu) := \{k \in A_n(\nu) : k \text{ even}\}.$$

For $k \in A_n^2(\nu)$, we can write $k^{d/2}\mathbb{P}(\hat{S}_k = \hat{x}) = 2(d/(2\pi))^{d/2} + O(1/k)$ by (A.2). By the definition of $A_n(\nu)$, we furthermore have

$$(1/p_{w,d}) \cdot \big(1/(1 + p_{w,d}^{-1} n^{\nu-1/2})\big) \le n/k \le (1/p_{w,d}) \cdot \big(1/(1 - p_{w,d}^{-1} n^{\nu-1/2})\big),$$

which implies that $n/k = 1/p_{w,d} + O(n^{\nu-1/2})$. Therefore,

$$\begin{aligned}
&n^{d/2}\mathbb{P}(S_n = x | \hat{N}(n) \in A_n(\nu)) \\
&= \textstyle\sum_{k \in A_n^2(\nu)} (n/k)^{d/2} \cdot k^{d/2}\mathbb{P}(\hat{S}_k = \hat{x}) \cdot \mathbb{P}(S_{n-k}^w = x_{d+1})\mathbb{P}(\hat{N}(n) = k | \hat{N}(n) \in A_n(\nu)) \\
&= 2(d/(2\pi p_{w,d}))^{d/2} \textstyle\sum_{k \in A_n^2(\nu)} \mathbb{P}(S_{n-k}^w = x_{d+1})\mathbb{P}(\hat{N}(n) = k | \hat{N}(n) \in A_n(\nu)) + o(1).
\end{aligned}$$

By (A.1), there exist constants $\gamma = \gamma(w) \in (0,1)$ and $C = C(w) > 0$ so that for $k \in A_n^2(\nu)$,

$$|\mathbb{P}(S_{n-k}^w = x_{d+1}) - (b_w/w)| \le C\gamma^{n(1-p_{w,d}) - n^{1/2+\nu}},$$

which implies that $\mathbb{P}(S_{n-k}^w = x_{d+1}) = b_w/w + O(\gamma^{n(1-p_{w,d}) - n^{1/2+\nu}})$. We thus obtain

$$\begin{aligned}
&n^{d/2}\mathbb{P}(S_n = x | \hat{N}(n) \in A_n(\nu)) \\
&= 2(b_w/w)(d/(2\pi p_{w,d}))^{d/2} \textstyle\sum_{k \in A_n^2(\nu)} \mathbb{P}(\hat{N}(n) = k | \hat{N}(n) \in A_n(\nu)) + o(1).
\end{aligned}$$

The remaining sum is the probability that $\hat{N}(n)$ is even given that $\hat{N}(n) \in A_n(\nu)$. Since for $X \sim \mathrm{Bin}(n,p)$, $\mathbb{P}(X \text{ is even}) = 1/2 + (1/2)(1-2p)^n$, and the probability that $\hat{N}(n) \in A_n(\nu)$ converges to 1 as $n \to \infty$, the sum converges to 1/2 as $n \to \infty$. The result follows. □

A.1.2 Proof of Lemma 2.2

Lemma A.2 (Asymptotic tail of T_0)**.** *Let $(Z_t)_{t \ge 0}$ be the SSRW on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate $\alpha > 0$, started at a nearest neighbor of the origin. Set $T_0 := \inf\{t \ge 0 : Z_t = 0\}$*

and define

$$\mu_w := p_w \pi w = \begin{cases} \pi, & w = 1, \\ (4/5)\pi w, & w = 2, \\ (2/3)\pi w, & w > 2, \end{cases}$$

where p_w is the probability given by (2.6)*. Then*

$$\mathbb{P}(T_0 > t) \sim \mu_w / \log t, \quad t \to \infty.$$

Proof. We begin by proving the result for the embedded discrete-time SSRW $(S_n)_{n\geq 0}$ on $\mathbb{Z}^2 \times \mathbb{Z}_w$, defined by $S_n := Z_{\sigma_n}$, where σ_n is the time of the n-th jump of $(Z_t)_{t\geq 0}$. Set $\tau_0 := \min\{n > 0 : S_n = 0\}$. We want to show that

$$\mathbb{P}(\tau_0 > n) \sim \mu_w / \log n, \quad n \to \infty.$$

The $w = 1$ case has already been established by Dvoretzky and Erdös in [55]. We can easily extend their argument to the $w > 1$ case by substituting the LCLT on $\mathbb{Z}^2$ by our LCLT on $\mathbb{Z}^2 \times \mathbb{Z}_w$. We sketch the argument briefly below.

Note first that instead of assuming that $(S_n)_{n\geq 0}$ is started at a nearest neighbor of the origin, we can assume that $(S_n)_{n\geq 0}$ is started at the origin itself, since for $n \geq 1$,

$$\mathbb{P}(\tau_0 > n | S_0 = 0) = \mathbb{P}(\tau_0 > n - 1 | S_0 = e_1).$$

Define $\gamma_2(n) := \mathbb{P}(\tau_0 > n - 1 | S_0 = 0)$ and $u_2(n) := \mathbb{P}(S_n = 0 | S_0 = 0)$ using the notation of [55]. Clearly, $\gamma_2(n)$ in decreasing in n with $\gamma_2(1) = 1$ and $\gamma_2(n) > 0$ for all $n \geq 1$. Assume that w is even, in which case $(S_n)_{n\geq 0}$ is periodic with period 2. Then $u_2(2n - 1) = 0$ for $n \geq 1$, and by the discrete-time LCLT of Lemma 2.1,

$$u_2(2n) = \mathbb{P}(S_{2n} = 0 | S_0 = 0) = (1/\mu_w)(1/n) + o(1/n), \quad n \geq 1. \tag{A.4}$$

If $S_0 = 0$, then for any $n \geq 1$, the walk visits the origin at least once by time $n - 1$ with probability 1. Decomposing this event in terms of the last return to the origin by time

$n-1$, and setting $m := n/2-1$ for even n and $m := (n-1)/2$ for odd n, we can write

$$u_2(0)\gamma_2(n) + u_2(2)\gamma_2(n-2) + \cdots + u_2(2m)\gamma_2(n-2m) = 1.$$

It follows from (A.4) that $\sum_{i=0}^{k} u_2(2i) = (1/\mu_w)(\log k)(1+o(1))$. Using the monotonicity of γ_2, one can now show that $\gamma_2(n) = (\mu_w + o(1))/\log n$, see page 356 of [55]. The argument for odd w, in which case the walk is aperiodic, follows along similar lines.

It remains to translate the above discrete-time result into continuous time. Fix $\varepsilon > 0$. Let $T_0 = \inf\{t \geq 0 : Z_t = 0\}$ be the time of first visit of $(Z_t)_{t\geq 0}$ to the origin, and let $N(t)$ denote the number of jumps Z_t makes by time t. Then

$$\mathbb{P}(T_0 > t) = \mathbb{P}(\tau_0 > N(t)) \leq \mathbb{P}\left(\tau_0 > \alpha t(1-\varepsilon)\right) + \mathbb{P}\left(N(t) < \alpha t(1-\varepsilon)\right),$$

and

$$\mathbb{P}(T_0 > t) = \mathbb{P}(\tau_0 > N(t)) \geq \mathbb{P}\left(\tau_0 > \alpha t(1+\varepsilon)\right) - \mathbb{P}\left(N(t) > \alpha t(1+\varepsilon)\right).$$

Since $(N(t))_{t\geq 0}$ is a Poisson process with rate α, we have $N(t)/(\alpha t) \to 1$ almost surely as $t \to \infty$. The result follows. □

A.1.3 Proof of Lemma 2.3

Lemma A.3. *Let $(Z_t^1)_{t\geq 0}$ and $(Z_t^2)_{t\geq 0}$ be independent SSRWs on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate 1, started at 0 and a nearest neighbor of 0. Set $T_0 := \inf\{t \geq 0 : Z_t^1 = Z_t^2\}$, and let S be an exponential random variable with mean $1/\beta$, independent of $(Z_t^1)_{t\geq 0}$ and $(Z_t^2)_{t\geq 0}$. Then*

(1) $\alpha_0(\beta) := \mathbb{P}(T_0 \leq \min\{S, \tau(\beta)\}) \to 1$ *as* $\beta \to 0$.

(2) $\alpha_1(\beta) := \mathbb{P}(S \leq \min\{T_0, \tau(\beta)\}) = \Theta\big(1/(\log(1/\beta))^{3/2}\big)$ *as* $\beta \to 0$,

(3) $\alpha_2(\beta) := \mathbb{P}(\tau(\beta) \leq \min\{S, T_0\}) \sim \mu_w/\log(1/\beta)$ *as* $\beta \to 0$.

Furthermore, there exists $C > 0$ so that for sufficiently small β,

(4) $\mathbb{E}\big[\sup_{t\leq\tau(\beta)} ||Z_t^1||^j \,\big|\, T_0 \leq \min\{S, \tau(\beta)\}\big] \leq Cj!\tau(\beta)^{j/2}, \quad j \geq 1,$

(5) $\mathbb{E}\big[\sup_{t\le\tau(\beta)}||Z_t^1||^j \,\big|\, S\le\min\{T_0,\tau(\beta)\}\big]\le Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2},\quad j\ge 1,$

(6) $\mathbb{E}\big[\sup_{t\le\tau(\beta)}||Z_t^1||^j \,\big|\, \tau(\beta)\le\min\{S,T_0\}\big]\le Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2},\quad j\ge 1.$

Each of (4)-(6) continues to hold if Z_t^1 is replaced by Z_t^2 or $\bar{Z}_t = Z_t^1 - Z_t^2$.

Proof. Recall that $\tau(\beta)=(1/\beta)\big(1/\sqrt{\log(1/\beta)}\big)$, and define

$$a(\beta) := (1/\beta)\big(1/(\log(1/\beta))^{3/2}\big) = o(\tau(\beta)).$$

(1) Follows from (2) and (3).

(2) Since S is independent of T_0, we can write

$$\begin{aligned}\mathbb{P}(S\le\min\{T_0,\tau(\beta)\}) &= \textstyle\int_0^{\tau(\beta)}\mathbb{P}(s\le T_0)\cdot\beta e^{-\beta s}ds\\ &= \textstyle\int_0^{a(\beta)}\mathbb{P}(s\le T_0)\cdot\beta e^{-\beta s}ds + \int_{a(\beta)}^{\tau(\beta)}\mathbb{P}(s\le T_0)\cdot\beta e^{-\beta s}ds.\end{aligned}$$

For the former integral, it is easy to see that

$$\begin{aligned}\textstyle\int_0^{a(\beta)}\mathbb{P}(s\le T_0)\cdot\beta e^{-\beta s}ds &\le \textstyle\int_0^{a(\beta)}\beta e^{-\beta s}ds\\ &= 1-e^{-\beta a(\beta)}\sim 1/(\log(1/\beta))^{3/2},\quad \beta\to 0.\end{aligned}$$

For the latter integral, we can write

$$\textstyle\int_{a(\beta)}^{\tau(\beta)}\mathbb{P}(s\le T_0)\cdot\beta e^{-\beta s}ds \le (\tau(\beta)-a(\beta))\cdot\beta e^{-\beta a(\beta)}\cdot\mathbb{P}(a(\beta)\le T_0). \tag{A.5}$$

Now, $(\tau(\beta)-a(\beta))\sim\tau(\beta)=(1/\beta)\big(1/\sqrt{\log(1/\beta)}\big)$ as $\beta\to 0$, $e^{-\beta a(\beta)}\to 1$ as $\beta\to 0$, and by Lemma 2.2, $\mathbb{P}(a(\beta)\le T_0)\sim\mu_w/\log(1/\beta)$ as $\beta\to 0$. The right-hand side of (A.5) is therefore of order $1/(\log(1/\beta))^{3/2}$ as $\beta\to 0$. Thus,

$$\mathbb{P}\big(S\le\min\{T_0,\tau(\beta)\}\big) = O\big(1/(\log(1/\beta))^{3/2}\big).$$

On the other hand, by independence,

$$\begin{aligned}\mathbb{P}(S \le \min\{T_0, \tau(\beta)\}) &\ge \mathbb{P}(S \le \tau(\beta) \le T_0)\\ &= \mathbb{P}(S \le \tau(\beta)) \cdot \mathbb{P}(\tau(\beta) \le T_0)\\ &\sim \mu_w/(\log(1/\beta))^{3/2}, \quad \beta \to 0,\end{aligned}$$

since $\mathbb{P}(S \le \tau(\beta)) = 1 - e^{-\beta\tau(\beta)} \sim 1/\sqrt{\log(1/\beta)}$ as $\beta \to 0$, and $\mathbb{P}(\tau(\beta) \le T_0) \sim \mu_w/\log(1/\beta)$ as $\beta \to 0$ by Lemma 2.2. Therefore,

$$\mathbb{P}(S \le \min\{T_0, \tau(\beta)\}) = \Omega\big(1/(\log(1/\beta))^{3/2}\big).$$

Since $\mathbb{P}(S \le \min\{T_0, \tau(\beta)\})$ is both $O\big(1/(\log(1/\beta))^{3/2}\big)$ and $\Omega\big(1/(\log(1/\beta))^{3/2}\big)$ as $\beta \to 0$, we have established part (2).

(3) By independence,

$$\begin{aligned}\mathbb{P}(\tau(\beta) \le \min\{S, T_0\}) &= \mathbb{P}(\tau(\beta) \le S) \cdot \mathbb{P}(\tau(\beta) \le T_0)\\ &= e^{-\beta\tau(\beta)} \cdot \mathbb{P}(\tau(\beta) \le T_0)\\ &\sim \mu_w/\log(1/\beta), \quad \beta \to 0,\end{aligned}$$

which follows from the fact that $\exp(-\beta\tau(\beta)) \to 1$ as $\beta \to 0$, and that $\mathbb{P}(\tau(\beta) \le T_0) \sim \mu_w/\log(1/\beta)$ as $\beta \to 0$ by Lemma 2.2.

(4) Note that we can write $Z_t^1 = (Z_{1,t}^1, Z_{2,t}^1, Z_{3,t}^1)$, where $Z_{1,t}^1$ and $Z_{2,t}^1$ are SSRWs on $\mathbb{Z}$ with jump rate $p_w/2$ each, and $Z_{3,t}^1$ is the SSRW on $\mathbb{Z}_w$ with jump rate $1 - p_w$, where p_w is defined as in (2.6). All walks are started at 0. By part (1) above, we have $\mathbb{P}(T_0 \le \min\{S, \tau(\beta)\}) \ge 1/2$ for sufficiently small β, which implies that for any $x > 0$,

$$\begin{aligned}&\mathbb{P}\big(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} ||Z_t^1|| > x \,\big|\, T_0 \le \min\{S, \tau(\beta)\}\big)\\ &\le 2\,\mathbb{P}\big(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} ||Z_t^1|| > x\big).\end{aligned} \tag{A.6}$$

Then note that for any $x > 0$,

$$\mathbb{P}\big(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} ||Z_t^1|| > x\big) \le 4\,\mathbb{P}\big(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} Z_{1,t}^1 > x/2\big), \tag{A.7}$$

since $Z_{1,t}^1$ and $Z_{2,t}^1$ have the same distribution. Now, $Z_{1,t}^1$ takes steps ± 1 with equal probability at rate $p_w/2$. The steps have moment generating function

$$\phi(\theta) = (e^{\theta} + e^{-\theta})/2,$$

so $Z_{1,t}^1$ has moment generating function

$$\begin{aligned} \psi_t(\theta) &= \mathbb{E}\big[\exp(\theta Z_{1,t}^1)\big] \\ &= \textstyle\sum_{n=0}^{\infty} e^{-(p_w/2)t}\big(((p_w/2)t)^n/n!\big) \cdot \phi(\theta)^n \\ &= \exp\big((p_w/2)t \cdot (\phi(\theta) - 1)\big). \end{aligned} \tag{A.8}$$

For any $x > 0$ and $\theta > 0$, we thus obtain by Doob's inequality,

$$\begin{aligned} &\mathbb{P}(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} Z_{1,t}^1 > x/2) \\ &= \mathbb{P}\big(\sup_{t \le \tau(\beta)} \exp(\theta \tau(\beta)^{-1/2} Z_{1,t}^1) > \exp(\theta x/2)\big) \\ &\le \exp\big((p_w/2)\tau(\beta) \cdot (\phi(\theta\tau(\beta)^{-1/2}) - 1)\big) \cdot \exp(-\theta x/2). \end{aligned}$$

Since $\phi(\theta\tau(\beta)^{-1/2}) = 1 + (1/2)\theta^2\tau(\beta)^{-1} + \cdots$, we can find $C_1 = C_1(\theta) > 0$ so that for sufficiently small β,

$$(p_w/2)\tau(\beta) \cdot (\phi(\theta\tau(\beta)^{-1/2}) - 1) \le \log C_1.$$

If we take $\theta = 2$, we then obtain for sufficiently small β,

$$\mathbb{P}(\tau(\beta)^{-1/2} \sup_{t \le \tau(\beta)} Z_{1,t}^1 > x/2) \le C_1 \exp(-x). \tag{A.9}$$

Now fix $j \geq 1$ and write, by (A.6), (A.7) and (A.9),

$$\begin{aligned}
&\tau(\beta)^{-j/2}\mathbb{E}\big[\sup_{t\leq\tau(\beta)}||Z_t^1||^j \,\big|\, T_0 \leq \min\{S,\tau(\beta)\}\big] \\
&= \int_0^\infty \mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x^{1/j} \,\big|\, T_0 \leq \min\{S,\tau(\beta)\}\big)dx \\
&\leq C_2\int_0^\infty \exp\big(-x^{1/j}\big)dx,
\end{aligned}$$

where $C_2 := 8C_1 > 0$. Using the substitution $u = x^{1/j}$, $x = u^j$, $dx = ju^{j-1}du$, we obtain

$$\tau(\beta)^{-j/2}\mathbb{E}\big[\sup_{t\leq\tau(\beta)}||Z_t^1||^j \,\big|\, T_0 \leq \min\{S,\tau(\beta)\}\big] \leq C_2 j\int_0^\infty e^{-u}u^{j-1}du = C_2 j!,$$

since $\int_0^\infty e^{-u}u^{j-1}du = (j-1)!$. The result follows.

(5) We begin by noting that for any $x > 0$, by independence and Cauchy-Schwarz,

$$\begin{aligned}
&\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x, S \leq \min\{T_0,\tau(\beta)\}\big) \\
&= \int_0^{\tau(\beta)} \mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x, s \leq T_0\big)\cdot\beta e^{-\beta s}ds \\
&\leq \sqrt{\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x\big)}\cdot\int_0^{\tau(\beta)}\sqrt{\mathbb{P}(s \leq T_0)}\cdot\beta e^{-\beta s}ds.
\end{aligned}$$

By the same analysis as in part (2) above, the integral is $O\big(1/\log(1/\beta)\big)$. Since

$$\mathbb{P}(S \leq \min\{T_0,\tau(\beta)\}) = \Theta\big(1/(\log(1/\beta))^{3/2}\big)$$

by part (2), we obtain for some $C_3 > 0$ and sufficiently small β,

$$\begin{aligned}
&\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x \,\big|\, S \leq \min\{T_0,\tau(\beta)\}\big) \\
&\leq C_3(\log(1/\beta))^{1/2}\sqrt{\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\leq\tau(\beta)}||Z_t^1|| > x\big)}.
\end{aligned}$$

The same argument as in part (4) above now yields the desired result. The only modification is that we take $\theta = 4$ instead of $\theta = 2$ in the calculations due to the square root.

(6) We begin by writing for any $x > 0$, by Cauchy-Schwarz,

$$\begin{aligned}&\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\le\tau(\beta)}||Z_t^1|| > x, \tau(\beta) \le \min\{S, T_0\}\big)\\ &\le \sqrt{\mathbb{P}\big(\tau(\beta)^{-1/2}\sup_{t\le\tau(\beta)}||Z_t^1|| > x\big)}\cdot\sqrt{\mathbb{P}(\tau(\beta)\le\min\{S,T_0\})},\end{aligned}$$

and the same argument as in part (5) yields the desired result, the only difference being that we appeal to the result of part (3) instead of part (2). □

A.1.4 Proof of Lemma 2.4

Lemma A.4. *Fix $a > 0$ and $0 < \rho < 1$, and define $A_r := [r, \infty) \times \mathbb{R} \times \mathbb{Z}_w$. For each $\delta > 0$, there exist $M > 0$ and $\beta_0 > 0$ so that*

$$\mathbb{P}\big(\tilde{\zeta}_t^{\beta,0} \cap A_{b+at} \neq \varnothing\big) \le 4\,\mathbb{P}\big(\mathring{\phi}_t^{\beta,0} \cap A_{b+\rho at} \neq \varnothing\big) + Me^{-\delta t}, \quad \beta \le \beta_0,\ t > 0,\ b \in \mathbb{R}.$$

Proof. To avoid notational overload, we suppress the initial conditions of the processes used in the proof, and assume that each is started with a single particle at the origin.

As in Section 2.3.4, let $\mathring{\phi}_t^{(k)}$ for $k = 0, 1, 2$ be the subprocess of the unaltered BRW $\bar{\phi}_t$ obtained by gathering offsprings of particles in the pruned BRW $\mathring{\phi}_t$ that have just been born through a type-k branching event and whose decision period has not yet passed. Recall that by (2.18),

$$\tilde{\zeta}_t \subseteq \mathring{\phi}_t \cup \mathring{\phi}_t^{(0)} \cup \mathring{\phi}_t^{(1)} \cup \mathring{\phi}_t^{(2)}.$$

Since $0 < \rho < 1$, it follows that

$$\mathbb{P}\big(\tilde{\zeta}_t^{\beta} \cap A_{b+at} \neq \varnothing\big) \le \mathbb{P}\big(\mathring{\phi}_t^{\beta} \cap A_{b+\rho at} \neq \varnothing\big) + \textstyle\sum_{k=0}^{2}\mathbb{P}\big(\mathring{\phi}_t^{(k),\beta} \cap A_{b+at} \neq \varnothing\big). \tag{A.10}$$

To analyze the terms in the sum, consider $\mathbb{P}\big(\mathring{\phi}_t^{(k),\beta} \cap A_{b+at} \neq \varnothing\big)$ for $k = 1$. We begin by writing

$$\begin{aligned}&\mathbb{P}\big(\mathring{\phi}_t^{(1),\beta} \cap A_{b+at} \neq \varnothing\big)\\ &\le \mathbb{P}\big(\mathring{\phi}_t^{\beta} \cap A_{b+\rho at} \neq \varnothing\big) + \mathbb{P}\big(\mathring{\phi}_t^{(1),\beta} \cap A_{b+at} \neq \varnothing, \mathring{\phi}_t^{\beta} \cap A_{b+\rho at} = \varnothing\big).\end{aligned} \tag{A.11}$$

Enumerate the particles in $\mathring{\phi}_t^{(1),\beta}$ as $X_t^{\beta,1}, X_t^{\beta,2}, \ldots, X_t^{\beta,N_t^{(1),\beta}}$, where $N_t^{(1),\beta}$ is the number of particles in $\mathring{\phi}_t^{(1),\beta}$. For each $X_t^{\beta,i}$, let $Y_t^{\beta,i}$ denote the position of its parent particle in $\mathring{\phi}_t^{\beta}$. Then

$$\begin{aligned}
&\mathbb{P}\big(\mathring{\phi}_t^{(1),\beta} \cap A_{b+at} \neq \varnothing, \mathring{\phi}_t^{\beta} \cap A_{b+\rho at} = \varnothing\big) \\
&\leq \mathbb{P}\big(\textstyle\bigcup_{i=1}^{N_t^{(1),\beta}} \{||X_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at\}\big) \\
&\leq \textstyle\sum_{i=1}^{\infty} \mathbb{P}\big(||X_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at \,\big|\, i \leq N_t^{(1),\beta}\big)\mathbb{P}\big(i \leq N_t^{(1),\beta}\big). \qquad \text{(A.12)}
\end{aligned}$$

Let $(Z_s^1)_{s\geq 0}$ and $(Z_s^2)_{s\geq 0}$ be independent SSRWs on $\mathbb{Z}^2 \times \mathbb{Z}_w$, started at 0 and a randomly chosen neighbor of 0, let T_0 be the time at which they first meet, and let S be an exponential random variable with mean $1/\beta$, independent of $(Z_s^1)_{s\geq 0}$ and $(Z_s^2)_{s\geq 0}$. Since each particle in $\mathring{\phi}_t^{(1)}$ has existed for at most $\tau(\beta)$ time units, and each particle is conditioned on producing its own offspring quickly, we can write for any $i \geq 1$,

$$\begin{aligned}
&\mathbb{P}\big(||X_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at \,\big|\, i \leq N_t^{(1),\beta}\big) \\
&\leq \mathbb{P}\big(h(\beta)^{-1/2} \sup_{s\leq\tau(\beta)} \|Z_s^1 - Z_s^2\| > (1-\rho)at \,\big|\, S \leq \min\{T_0, \tau(\beta)\}\big). \qquad \text{(A.13)}
\end{aligned}$$

For ease of notation, let $(\tilde{Z}_s)_{s\geq 0}$ denote a walk on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with the distribution of $(Z_s^1 - Z_s^2)_{s\geq 0}$ conditional on $\{S \leq \min\{T_0, \tau(\beta)\}\}$. Write $\tilde{Z}_s = (\tilde{Z}_{1,s}, \tilde{Z}_{2,s}, \tilde{Z}_{3,s})$, and note that

$$\begin{aligned}
&\mathbb{P}\big(h(\beta)^{-1/2} \sup_{s\leq\tau(\beta)} \|\tilde{Z}_s\| > (1-\rho)at\big) \\
&\leq 4\,\mathbb{P}\big(h(\beta)^{-1/2} \sup_{s\leq\tau(\beta)} \tilde{Z}_{1,s} > (1/2)(1-\rho)at\big). \qquad \text{(A.14)}
\end{aligned}$$

We now consider the moment generating function

$$\mathbb{E}\big[\exp\big(\nu h(\beta)^{-1/2} \tilde{Z}_{1,\tau(\beta)}\big)\big].$$

By part (5) in Lemma 2.3, there exists $C > 0$ so that for sufficiently small β,

$$\mathbb{E}|\tilde{Z}_{1,\tau(\beta)}|^j \leq Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2}, \quad j \geq 1.$$

We therefore obtain, with $\tau(\beta) = (1/\beta)\big(1/\sqrt{\log(1/\beta)}\big)$ and $h(\beta) = (1/\beta)\cdot\log(1/\beta)$,

$$h(\beta)^{-j}\mathbb{E}|\tilde{Z}_{1,\tau(\beta)}|^{2j} \le C(2j)!(\log(1/\beta))^{(1-3j)/2} \le C(2j)!(\log(1/\beta))^{-j}, \quad j \ge 1,$$

where we use that $(1-3j)/2 \le -j$ whenever $j \ge 1$. Using that $\tilde{Z}_{1,\tau(\beta)} \overset{d}{=} -\tilde{Z}_{1,\tau(\beta)}$ by symmetry, we obtain for sufficiently small β,

$$\begin{aligned}\mathbb{E}\big[\exp\big(\nu h(\beta)^{-1/2}\tilde{Z}_{1,\tau(\beta)}\big)\big] &= \textstyle\sum_{j=0}^{\infty}\big(1/(2j)!\big)\cdot\nu^{2j}h(\beta)^{-j}\mathbb{E}|\tilde{Z}_{1,\tau(\beta)}|^{2j}\\ &\le 1 + C\textstyle\sum_{j=1}^{\infty}\big(1/(2j)!\big)\cdot\nu^{2j}(2j)!(\log(1/\beta))^{-j}\\ &= 1 + O_{C,\nu}\big((\log(1/\beta))^{-1}\big), \end{aligned} \tag{A.15}$$

where $O_{C,\nu}$ signifies that the error term depends on C and ν. By Doob's inequality,

$$\begin{aligned}&\mathbb{P}\big(h(\beta)^{-1/2}\sup_{s\le\tau(\beta)}\tilde{Z}_{1,s} > (1/2)(1-\rho)at\big)\\ &\le \exp\big(-(1/2)(1-\rho)\nu at\big)\cdot\mathbb{E}\big[\exp\big(\nu h(\beta)^{-1/2}\tilde{Z}_{1,\tau(\beta)}\big)\big], \quad t > 0.\end{aligned}$$

By first choosing ν sufficiently large and then β sufficiently small, we can for any $\delta > 0$ find a constant $M_1 > 0$ so that for sufficiently small β,

$$\mathbb{P}\big(h(\beta)^{-1/2}\sup_{s\le\tau(\beta)}\tilde{Z}_{1,s} > (1/2)(1-\rho)at\big) \le M_1 e^{-(\delta+2\mu_w)t}, \quad t > 0.$$

It then follows from (A.12), (A.13) and (A.14) that for $M_2 := 4M_1$ and sufficiently small β,

$$\begin{aligned}\mathbb{P}\big(\mathring{\phi}_t^{(1),\beta}\cap A_{b+at}\ne\varnothing, \mathring{\phi}_t^{\beta}\cap A_{b+\rho at} = \varnothing\big) &\le M_2 e^{-(\delta+2\mu_w)t}\cdot\textstyle\sum_{i=1}^{\infty}\mathbb{P}\big(i \le N_t^{(1),\beta}\big)\\ &= M_2 e^{-(\delta+2\mu_w)t}\cdot\mathbb{E}\Big[N_t^{(1),\beta}\Big], \quad t > 0.\end{aligned}$$

Now, $\mathbb{E}\Big[N_t^{(1),\beta}\Big] = \mathbb{E}\big|\mathring{\phi}_t^{(1),\beta}\big| \le \mathbb{E}\big|\mathring{\phi}_t^{\beta}\big| \le \mathbb{E}\big|\phi_t^{\beta}\big|$, where ϕ_t^{β} denotes the simple BRW defined in Section 2.3.4. By Lemma 2.3, ϕ_t^{β} branches at exponential rate $\gamma(\beta) = \mu_w + o(1)$, with mean number of new particles $\ell(\beta) = 1 + o(1)$ introduced per branching event. Since $\gamma(\beta)\cdot\ell(\beta) \le 2\mu_w$ for sufficiently small β, we can write

$$\mathbb{E}\Big[N_t^{(1),\beta}\Big] \le \mathbb{E}\big|\phi_t^{\beta}\big| \le \exp(2\mu_w t)$$

for sufficiently small β and any $t > 0$. We thus obtain for sufficiently small β,

$$\mathbb{P}\big(\mathring{\phi}_t^{(1),\beta} \cap A_{b+at} \neq \varnothing, \mathring{\phi}_t^{\beta} \cap A_{b+\rho at} = \varnothing\big) \leq M_2 e^{-\delta t}, \quad t > 0.$$

The same bound holds for $k = 0$ and $k = 2$ by parts (4) and (6) of Lemma 2.3. By setting $M := 3M_2$, the desired result then follows from (A.10) and (A.11). □

A.1.5 Proof of Lemma 2.5

Lemma A.5. *For a particle chosen uniformly at random from $\phi_t^{\beta,0}$, let $(Y_s^{\beta})_{s\leq t}$ be the path followed by this particle and its ancestors, and let $(\mathring{Y}_s^{\beta})_{s\leq t}$ be the corresponding path in $\mathring{\phi}_t^{\beta,0}$. Then, for any $r > 0$ and $\delta > 0$, there exist $M > 0$ and $\beta_0 > 0$ so that*

$$\mathbb{P}\big(\sup_{s\leq t} ||Y_s^{\beta} - \mathring{Y}_s^{\beta}|| > rt\big) \leq M e^{-\delta t}, \quad \beta \leq \beta_0,\ t > 0.$$

Proof. As in the proof of Lemma 2.4, we suppress the initial conditions of the processes from the notation, and assume that each is started with a single particle at the origin.

Consider the pruned BRW $(\mathring{\phi}_s)_{s\geq 0}$ (in the unscaled spacetime). Type-1 and type-2 branching events occur at total rate $\sim \mu_w/h(\beta)$ as $\beta \to 0$ by Lemma 2.3, where we recall that $h(\beta) = (1/\beta) \cdot \log(1/\beta)$. In the simple BRW $(\phi_s)_{s\geq 0}$, we modify the path of parent and daughter at each type-1 and type-2 event as outlined in Section 2.3.4. Type-0 branching events, where the daughter is not introduced to $(\mathring{\phi}_s)_{s\geq 0}$, occur at rate $\sim \beta$ as $\beta \to 0$ by Lemma 2.3. In the simple BRW $(\phi_s)_{s\geq 0}$, we modify the parent's path at a type-0 event with probability $1 - \alpha_0(\beta) \sim \mu_w/\log(1/\beta)$ as $\beta \to 0$, so the modification rate for type-0 events is $\sim \mu_w/h(\beta)$. The total rate of modifications for type-0, type-1 and type-2 events is therefore $\sim 2\mu_w/h(\beta)$ as $\beta \to 0$, which translates into $\sim 2\mu_w$ in the scaled spacetime.

Consider now a type-1 branching event in $(\mathring{\phi}_s)_{s\geq 0}$. Assume for simplicity that the parent is at the origin at the time of branching. Let $(Z_t^1)_{t\geq 0}$ and $(Z_t^2)_{t\geq 0}$ be the SSRWs followed by parent and daughter from the time of branching, let T_0 be the time at which they first meet, and let S be the time at which the daughter first branches. Note that the two paths are conditioned on $\{S \leq \min\{T_0, \tau(\beta)\}\}$. Next, let $(Z_t^3)_{t\geq 0}$ and $(Z_t^4)_{t\geq 0}$ be independent SSRWs with jump rate 1 that are independent of $(Z_t^1)_{t\geq 0}$, $(Z_t^2)_{t\geq 0}$ and

S, both started at the origin. Define

$$\hat{Z}_t^1 := \begin{cases} Z_t^3, & t \le S, \\ Z_S^3 + (Z_t^1 - Z_S^1), & t \ge S, \end{cases}$$

and

$$\hat{Z}_t^2 := \begin{cases} Z_t^4, & t \le S, \\ Z_S^4 + (Z_t^2 - Z_S^2), & t \ge S. \end{cases}$$

In the simple BRW $(\phi_s)_{s\ge 0}$, we replace $(Z_t^1)_{t\ge 0}$ and $(Z_t^2)_{t\ge 0}$ by $(\hat{Z}_t^1)_{t\ge 0}$ and $(\hat{Z}_t^2)_{t\ge 0}$, respectively. Note that in the definition of $(\hat{Z}_t^1)_{t\ge 0}$, we connect the paths $(Z_t^3)_{t\le S}$ and $(Z_t^1)_{t\ge S}$ at time S (see Figure 2.5 in the main text), which requires perturbing $(Z_t^1)_{t\ge S}$ by $Z_S^3 - Z_S^1$ conditional on $\{S \le \min\{T_0, \tau(\beta)\}\}$. Call this perturbation X. By the same argument as used to establish part (5) in Lemma 2.3, there exists $C > 0$ so that for sufficiently small β,

$$\mathbb{E}||X||^j = \mathbb{E}\big[||Z_S^3 - Z_S^1||^j \,\big|\, S \le \min\{T_0, \tau(\beta)\}\big] \le Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2}, \quad j \ge 1.$$

The perturbation of $(Z_t^2)_{t\ge S}$ in the definition of $(\hat{Z}_t^2)_{t\ge 0}$ has the same upper bound by Lemma 2.3. Furthermore, the same argument, using part (6) of Lemma 2.3 instead of part (5), will show that this upper bound also applies to perturbations on type-2 branching events.

On type-0 branching events, $(Z_t^1)_{t\ge 0}$ and $(Z_t^2)_{t\ge 0}$ will be conditioned on $\{T_0 \le \min\{S, \tau(\beta)\}\}$. Here, $(Z_t^3)_{t\ge 0}$ and $(Z_t^4)_{t\ge 0}$ will be conditioned on $\{S' \le \min\{T_0', \tau(\beta)\}\}$ or $\{\tau(\beta) \le \min\{S', T_0'\}\}$, with S' and T_0' defined in terms of $(Z_t^3)_{t\ge 0}$ and $(Z_t^4)_{t\ge 0}$. Let X be the perturbation required to connect $(Z_t^3)_{t\le\tau(\beta)}$ and $(Z_t^1)_{t\ge\tau(\beta)}$ at time $\tau(\beta)$ (recall that we only modify the parent's path on a type-0 branching event). By the same argument as used to establish parts (4)-(6) in Lemma 2.3, we can show that

$$\mathbb{E}||X||^j \le Cj!(\log(1/\beta))^{1/2}\tau(\beta)^{j/2}, \quad j \ge 1, \tag{A.16}$$

the same as for the type-1 and type-2 branching events.

Consider next the scaled processes $(\mathring{\phi}_s^\beta)_{s\ge 0}$ and $(\phi_s^\beta)_{s\ge 0}$. We have already observed

that each particle path in $(\mathring{\phi}_s^\beta)_{s\geq 0}$ is perturbed at total rate $\sim 2\mu_w$ as $\beta \to 0$ to produce the corresponding path in $(\phi_s^\beta)_{s\geq 0}$. Some of the perturbations occur on type-0 branching events, in which case no new particle is introduced, while the remaining perturbations occur on type-1 and type-2 branching events, in which case one new particle is introduced. To keep track of the perturbations, we define a branching process embedded in $(\mathring{\phi}_s^\beta)_{s\geq 0}$ as follows:

- Each individual in the branching process is associated with a particle in $(\mathring{\phi}_s^\beta)_{s\geq 0}$.
- At each type-0 perturbation in $(\mathring{\phi}_s^\beta)_{s\geq 0}$, the corresponding individual in the branching process is killed and replaced by another individual.
- At each type-1 or type-2 perturbation in $(\mathring{\phi}_s^\beta)_{s\geq 0}$, the corresponding individual in the branching process is killed and replaced by two individuals.

For a given particle path in $(\mathring{\phi}_s^\beta)_{s\geq 0}$, the number of perturbations up until time t is the generation number at time t of the corresponding individual in the branching process. The branching process has branching rate $\sim 2\mu_w$ and mean number of offspring $\sim 3/2$ per branching event as $\beta \to 0$. To obtain an upper bound on the number of perturbations by time t, it therefore suffices to establish the following lemma.

Lemma A.5.1. *Let $(\varphi_s)_{s\geq 0}$ be a continuous-time branching process with exponential branching rate $\alpha > 0$ and offspring distribution $(p_j)_{j=0}^\infty$ with $p_0 = 0$, started with a single individual. Let G_t be the generation number of an individual selected uniformly at random from φ_t. Set $m := \sum_j jp_j$ and assume $1 < m \leq m_0$. Then, for any $c > m_0$,*

$$\mathbb{P}(G_t > c\alpha t) \leq \exp\big(\alpha t \cdot (2c(m_0 - 1) - (c - m_0)^2)/(2c)\big), \quad t > 0.$$

Proof. Let $\varphi_{j,t}$ be the number of individuals at time t that are in generation j. Then

$$\mathbb{P}(G_t > c\alpha t) = \mathbb{E}\big[(1/\varphi_t) \cdot \textstyle\sum_{j>c\alpha t} \varphi_{j,t}\big] \leq \textstyle\sum_{j>c\alpha t} \mathbb{E}\big[\varphi_{j,t}\big],$$

where we use that $\varphi_t \geq 1$ since $p_0 = 0$. Note that $\mathbb{E}\big[\varphi_{j,t}\big] = m^j \mathbb{P}(R_t = j)$, where R_t is Poisson distributed with mean αt. Let $\hat{R}_t$ be Poisson distributed with mean $m_0 \alpha t$.

Then

$$\mathbb{P}(G_t > c\alpha t) \leq \sum_{j>c\alpha t} m_0^j \, \mathbb{P}(R_t = j) = \exp\big(\alpha t(m_0 - 1)\big) \cdot \mathbb{P}(\hat{R}_t > c\alpha t), \quad t > 0. \tag{A.17}$$

Note next that by Theorem 1 of [26], for $c > m_0$,

$$\mathbb{P}(\hat{R}_t > c\alpha t) \leq \exp\big(-\alpha t \cdot (c - m_0)^2/(2c)\big), \quad t > 0.$$

The result then follows from (A.17). □

We are now ready to begin the main calculations. First, select a particle uniformly at random from ϕ_t^β. Let $(Y_s^\beta)_{s\leq t}$ be the path followed by this particle and its ancestors, and let $(\mathring{Y}_s^\beta)_{s\leq t}$ be the corresponding path in $\mathring{\phi}_t^\beta$. The number of perturbations G_t^β between the two paths by time t is the generation number of a particle selected uniformly at random from the embedded branching process. By the previous lemma, for a given $\delta > 0$, we can select $c > 0$ sufficiently large so that for sufficiently small β,

$$\mathbb{P}(G_t^\beta > ct) \leq \exp(-\delta t), \quad t > 0.$$

This implies

$$\begin{aligned} &\mathbb{P}\big(\sup_{s\leq t} ||Y_s^\beta - \mathring{Y}_s^\beta|| > rt\big) \\ &\leq \mathbb{P}\big(\sup_{s\leq t} ||Y_s^\beta - \mathring{Y}_s^\beta|| > rt, G_t^\beta \leq ct\big) + e^{-\delta t}, \quad t > 0. \end{aligned} \tag{A.18}$$

To analyze the probability in (A.18), note that by the above observations, we can write

$$Y_t^\beta = \mathring{Y}_t^\beta + \sum_{k=1}^{G_t^\beta} X_k^\beta,$$

where $X_1^\beta, X_2^\beta, \ldots$ is the (independent) sequence of scaled perturbations, and we assume for simplicity that the time point t does not occur during a decision period. It follows that

$$\mathbb{P}\big(\sup_{s\leq t} ||Y_s^\beta - \mathring{Y}_s^\beta|| > rt, G_t^\beta \leq ct\big) \leq \mathbb{P}\big(\sup_{m\leq ct} \big\| \sum_{k=1}^m X_k^\beta \big\| > rt\big). \tag{A.19}$$

To estimate this probability, we begin by noting that

$$\begin{aligned}
&\mathbb{P}\big(\sup_{m\le ct}\big\|\textstyle\sum_{k=1}^m X_k^\beta\big\| > rt\big)\\
&= \mathbb{P}\big(\sup_{m\le ct} h(\beta)^{-1/2}\big\|\textstyle\sum_{k=1}^m X_k\big\| > rt\big)\\
&\le 4\,\mathbb{P}\big(\sup_{m\le ct} h(\beta)^{-1/2}\textstyle\sum_{k=1}^m X_{1,k} > (1/2)rt\big),
\end{aligned} \tag{A.20}$$

where we write $X_k = (X_{1,k}, X_{2,k}, X_{3,k})$. We next consider the moment generating function

$$\mathbb{E}\big[\exp\big(\nu h(\beta)^{-1/2}\textstyle\sum_{k=1}^{\lfloor ct\rfloor} X_{1,k}\big)\big] = \prod_{k=1}^{\lfloor ct\rfloor}\mathbb{E}\big[\exp(\nu h(\beta)^{-1/2}X_{1,k})\big].$$

Using the moment bound (A.16), we can show that

$$\mathbb{E}\big[\exp(\nu h(\beta)^{-1/2}X_{1,k})\big] = 1 + O_{C,\nu}\big((\log(1/\beta))^{-1}\big), \quad k\ge 1,$$

using the same argument we used to establish (A.15) in the proof of Lemma 2.4. For a given $\nu > 0$, take $\beta_0 = \beta_0(\nu) > 0$ so that for $\beta \le \beta_0$,

$$\mathbb{E}\big[\exp(\nu h(\beta)^{-1/2}X_{1,k})\big] \le \exp\big((1/4)(r\nu/c)\big), \quad k\ge 1,$$

which implies that for $\beta \le \beta_0$,

$$\textstyle\prod_{k=1}^{\lfloor ct\rfloor}\mathbb{E}\big[\exp(\nu h(\beta)^{-1/2}X_{1,k})\big] \le \exp\big((1/4)r\nu t\big), \quad t > 0.$$

We then obtain by Doob's inequality for $\beta \le \beta_0$,

$$\mathbb{P}\big(\sup_{m\le ct} h(\beta)^{-1/2}\textstyle\sum_{k=1}^m X_{1,k} > (1/2)rt\big) \le \exp\big(-(1/4)r\nu t\big), \quad t > 0.$$

Choosing ν appropriately, we obtain by (A.20) for sufficiently small β,

$$\mathbb{P}\big(\sup_{m\le ct}\big\|\textstyle\sum_{k=1}^m X_k^\beta\big\| > rt\big) \le 4\,e^{-\delta t}, \quad t > 0.$$

Finally, combining with (A.18) and (A.19), we obtain the desired result. □

A.1.6 Proof of Lemma 2.6

Lemma A.6. *Fix $a > 0$ and $0 < \rho < 1$, and define $A_r := [r, \infty) \times \mathbb{R} \times \mathbb{Z}_w$. For each $\delta > 0$, there exist $M > 0$ and $\beta_0 > 0$ so that*

$$\mathbb{P}\big(\mathring{\phi}_t^{\beta,0} \cap A_{b+at} \neq \varnothing\big) \leq \mathbb{P}\big(\phi_t^{\beta,0} \cap A_{b+\rho at} \neq \varnothing\big) + Me^{-\delta t}, \quad \beta \leq \beta_0,\ t > 0,\ b \in \mathbb{R}.$$

Proof. As in the proofs of Lemmas 2.4 and 2.5, we suppress the initial conditions of the processes from the notation, and assume that each is started with a single particle at the origin. Since $0 < \rho < 1$, we can begin by writing

$$\begin{aligned}&\mathbb{P}\big(\mathring{\phi}_t^{\beta} \cap A_{b+at} \neq \varnothing\big)\\ &\leq \mathbb{P}\big(\phi_t^{\beta} \cap A_{b+\rho at} \neq \varnothing\big) + \mathbb{P}\big(\mathring{\phi}_t^{\beta} \cap A_{b+at} \neq \varnothing, \phi_t^{\beta} \cap A_{b+\rho at} = \varnothing\big).\end{aligned}$$

Enumerate the particles in ϕ_t^β as $Y_t^{\beta,1}, Y_t^{\beta,2}, \ldots, Y_t^{\beta,N_t^\beta}$, where N_t^β is the number of particles in ϕ_t^β. For each $Y_t^{\beta,i}$, let $\mathring{Y}_t^{\beta,i}$ denote the position of the corresponding particle in $\mathring{\phi}_t^\beta$. Then

$$\mathbb{P}\big(\mathring{\phi}_t^{\beta} \cap A_{b+at} \neq \varnothing, \phi_t^{\beta} \cap A_{b+\rho at} = \varnothing\big) \leq \mathbb{P}\big(\textstyle\bigcup_{i=1}^{N_t^\beta}\{||\mathring{Y}_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at\}\big),$$

and by Markov's inequality,

$$\mathbb{P}\big(\textstyle\bigcup_{i=1}^{N_t^\beta}\{||\mathring{Y}_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at\}\big) \leq \mathbb{E}\Big[\sum_{i=1}^{N_t^\beta} 1_{\{||\mathring{Y}_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at\}}\Big].$$

Let I be the index of a particle chosen uniformly at random from $Y_t^{\beta,1}, Y_t^{\beta,2}, \ldots, Y_t^{\beta,N_t^\beta}$, i.e. $\mathbb{P}\big(I = j \big| N_t^\beta\big) = 1/N_t^\beta$ for $j = 1, \ldots, N_t^\beta$. Then

$$\mathbb{E}\Big[\sum_{i=1}^{N_t^\beta} 1_{\{||\mathring{Y}_t^{\beta,i} - Y_t^{\beta,i}|| > (1-\rho)at\}}\Big] = \mathbb{E}\Big[N_t^\beta \cdot 1_{\{||\mathring{Y}_t^{\beta,I} - Y_t^{\beta,I}|| > (1-\rho)at\}}\Big],$$

and by Cauchy-Schwarz,

$$\mathbb{E}\Big[N_t^\beta \cdot 1_{\{||\mathring{Y}_t^{\beta,I} - Y_t^{\beta,I}|| > (1-\rho)at\}}\Big] \leq \sqrt{\mathbb{E}\big[(N_t^\beta)^2\big]} \cdot \sqrt{\mathbb{P}\big(||\mathring{Y}_t^{\beta,I} - Y_t^{\beta,I}|| > (1-\rho)at\big)}.$$

Now, ϕ_t^β branches at exponential rate $\gamma(\beta) = \mu_w + o(1)$, with mean number of new

particles $\ell(\beta) = 1 + o(1)$ introduced per branching event. Since $\gamma(\beta) \cdot \ell(\beta) \leq 2\mu_w$ for sufficiently small β, by Lemma 5 in [64], there exists $M_1 > 0$ so that for sufficiently small β,

$$\mathbb{E}\big[(N_t^\beta)^2\big] \leq M_1(e^{2\mu_w t})^2, \quad t > 0.$$

Furthermore, for any $\delta > 0$, by Lemma 2.5 above, there exists $M_2 > 0$ so that for sufficiently small β,

$$\mathbb{P}\big(||\mathring{Y}_t^{\beta,I} - Y_t^{\beta,I}|| > (1-\rho)at\big) \leq M_2\, e^{-2(\delta+2\mu_w)t}, \quad t > 0.$$

Combining the above, there exists $M > 0$ so that for sufficiently small β,

$$\mathbb{P}\big(\mathring{\phi}_t^\beta \cap A_{b+at} \neq \varnothing, \phi_t^\beta \cap A_{b+\rho at} = \varnothing\big) \leq Me^{-\delta t}, \quad t > 0,$$

and the result follows. □

A.1.7 Proof of Lemma 2.7

Lemma A.7. *Define $\tau_\varnothing^A := \min\{t \geq 0 : \xi_t^A = \varnothing\}$ for $A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$. For each $\kappa > 1$, there exists a family of random variables $(S_\beta)_{\beta>0}$, with $\mathbb{P}(S_\beta < \infty | \tau_\varnothing^0 = \infty) = 1$ for each $\beta > 0$, so that*

$$\lim_{\beta\to 0} \liminf_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1 \notin \xi^0_{S_\beta + h(\beta)t} \,\big|\, \tau_\varnothing^0 = \infty\big) = 1,$$

where $a_w := p_w\sqrt{\pi w}$.

Proof. Take $\varepsilon > 0$. We segment the proof into three main steps.

Step 1: Remove conditioning on nonextinction. Let B_r denote a box in $\mathbb{Z}^2 \times \mathbb{Z}_w$ centered at 0 with radius r, i.e.

$$B_r := \{(x_1, x_2, x_3) \in \mathbb{Z}^2 \times \mathbb{Z}_w : \max\{|x_1|, |x_2|\} \leq r\}.$$

If we set $r = r(\beta) = h(\beta)^{1/2}$, then by the gambler's ruin formula,

$$\mathbb{P}(\tau_{\varnothing}^{B_r} < \infty) = (1+\beta)^{-(2\lfloor r \rfloor+1)^2 w} = \exp\big(-\Theta(h(\beta)) \cdot \log(1+\beta)\big) = o(1), \quad \beta \to 0,$$

since $h(\beta) = (1/\beta) \cdot \log(1/\beta)$ and $\log(1+\beta) = \beta + o(\beta)$. Define $\sigma_{r,R} = \sigma_{r,R}(\beta)$ as

$$\sigma_{r,R} := \inf\{t \geq 0 : B_r \subseteq \xi_t^0 \subseteq B_{Rt}\}.$$

By (3) and (7) in Bramson & Griffeath [19] (i.e. the corresponding results for $\mathbb{Z}^2 \times \mathbb{Z}_w$), there exists a constant $R > 0$ so that

$$\mathbb{P}(\sigma_{r,R} < \infty \,|\, \tau_{\varnothing}^0 = \infty) = 1,$$

which implies $\mathbb{P}(\tau_{\varnothing}^0 = \infty) \leq \mathbb{P}(\sigma_{r,R} < \infty)$. We then get by the strong Markov property and the monotonicity property (2.10) of ξ_t, for any $t > 0$, $m \geq 1$ and sufficiently small β,

$$\begin{aligned}
&\mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^0_{\sigma_{r,R}+h(\beta)t}, \tau_{\varnothing}^0 = \infty\big) \\
&\geq \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^0_{\sigma_{r,R}+h(\beta)t}, \tau_{\varnothing}^0 = \infty, \sigma_{r,R} \leq m\big) \\
&= \textstyle\sum_{B_r \subseteq \Lambda \subseteq B_{mR}} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^{\Lambda}_{h(\beta)t}, \tau_{\varnothing}^{\Lambda} = \infty\big) \mathbb{P}(\sigma_{r,R} \leq m, \xi^0_{\sigma_{r,R}} = \Lambda) \\
&\geq \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^{B_{mR}}_{h(\beta)t}, \tau_{\varnothing}^{B_r} = \infty\big) \mathbb{P}(\sigma_{r,R} \leq m) \\
&\geq \big(\mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^{B_{mR}}_{h(\beta)t}\big) - \mathbb{P}\big(\tau_{\varnothing}^{B_r} < \infty\big)\big) \mathbb{P}(\sigma_{r,R} \leq m) \\
&\geq \big(\mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^{B_{mR}}_{h(\beta)t}\big) - \varepsilon\big) \mathbb{P}(\sigma_{r,R} \leq m).
\end{aligned}$$

If we can establish that for sufficiently small β,

$$\lim_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^{B_{mR}}_{h(\beta)t}\big) = 1, \quad m \geq 1,$$

then sending $m \to \infty$ will yield for sufficiently small β,

$$\liminf_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t \rceil e_1 \notin \xi^0_{\sigma_{r,R}+h(\beta)t}, \tau_{\varnothing}^0 = \infty\big) \geq (1-\varepsilon)\mathbb{P}(\sigma_{r,R} < \infty),$$

and the desired result will follow from $\mathbb{P}(\sigma_{r,R} < \infty) \geq \mathbb{P}(\tau_{\varnothing}^0 = \infty)$. Equivalently, it is

sufficient to show that for sufficiently small β,

$$\lim_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1 \in \xi_{h(\beta)t}^{B_{mR}}\big) = 0, \quad m \geq 1. \tag{A.21}$$

Step 2: Introduce duality and apply approximation scheme. Note first that by the duality relation (2.8) between ξ_t and $\tilde{\zeta}_t$, the translation invariance (2.11) and symmetry property (2.12) of the dual process $\tilde{\zeta}_t$, and the definition (2.17) of the scaled dual process $\tilde{\zeta}_t^\beta$,

$$\begin{aligned}\mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1 \in \xi_{h(\beta)t}^{B_{mR}}\big) &= \mathbb{P}\big(\tilde{\zeta}_{h(\beta)t}^{\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1} \cap B_{mR} \neq \varnothing\big) \\ &= \mathbb{P}\big(\tilde{\zeta}_{h(\beta)t}^{0} \cap (B_{mR} + \lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1) \neq \varnothing\big) \\ &= \mathbb{P}\big(\tilde{\zeta}_{t}^{\beta,0} \cap h(\beta)^{-1/2}(B_{mR} + \lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1) \neq \varnothing\big).\end{aligned} \tag{A.22}$$

Set $A_r := [r,\infty) \times \mathbb{R} \times \mathbb{Z}_w$ and take κ_1 and κ_2 so that $1 < \kappa_2 < \kappa_1 < \kappa$. For any $x = (x_1, x_2, x_3) \in B_{mR}$, we have $|x_1| \leq mR$ where $R > 0$ is a constant. We can therefore choose β sufficiently small that $|x_1^\beta| = h(\beta)^{-1/2}|x_1| \leq m$ for all $x \in B_{mR}$. Using the approximation Lemmas 2.4 and 2.6, we then obtain for sufficiently small β,

$$\begin{aligned}&\limsup_{t\to\infty} \mathbb{P}\big(\tilde{\zeta}_t^{\beta,0} \cap h(\beta)^{-1/2}(B_{mR} + \lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1) \neq \varnothing\big) \\ &\leq \limsup_{t\to\infty} \mathbb{P}(\tilde{\zeta}_t^{\beta,0} \cap A_{-m+\kappa a_w t} \neq \varnothing) \\ &\leq 4 \limsup_{t\to\infty} \mathbb{P}\big(\mathring{\phi}_t^{\beta,0} \cap A_{-m+\kappa_1 a_w t} \neq \varnothing\big) \\ &\leq 4 \limsup_{t\to\infty} \mathbb{P}\big(\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t} \neq \varnothing\big).\end{aligned} \tag{A.23}$$

We have now reduced the problem to analyzing the tail of $\phi_t^{\beta,0}$, which is straightforward.

Step 3: Analyze simple BRW. We begin by using Markov's inequality to write

$$\mathbb{P}\big(\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t} \neq \varnothing\big) \leq \mathbb{E}\big|\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t}\big|. \tag{A.24}$$

Recall that $\phi_t^{\beta,0}$ has branching rate $\mu_w + o(1)$, and on average, $1 + o(1)$ new particles

are added per branching event. Therefore,

$$\mathbb{E}\big|\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t}\big| = \exp\big((\mu_w + o(1))t\big) \cdot \mathbb{P}(h(\beta)^{-1/2} S_{h(\beta)t} \geq -m + \kappa_2 a_w t), \quad \text{(A.25)}$$

where $(S_s)_{s\geq 0}$ is the path of the SSRW on $\mathbb{Z}$ started at 0 with jump rate $p_w/2$, where p_w is defined as in (2.6). By (A.8), its moment generating function is

$$\psi_s(\theta) = \exp\big((p_w/2)s \cdot (\phi(\theta) - 1)\big),$$

where $\phi(\theta) = (e^\theta + e^{-\theta})/2$. Set $\theta_0 := 2\sqrt{\pi w}\kappa_2$, and note that

$$\psi_{h(\beta)t}(\theta_0 h(\beta)^{-1/2}) = \exp\big((p_w/2)h(\beta)t \cdot (\phi(\theta_0 h(\beta)^{-1/2}) - 1)\big).$$

Since $\phi(\theta_0 h(\beta)^{-1/2}) = 1 + (1/2)\theta_0^2 h(\beta)^{-1} + o(h(\beta)^{-1})$, and $\mu_w = p_w \pi w$, we get

$$\psi_{h(\beta)t}(\theta_0 h(\beta)^{-1/2}) = \exp\big((p_w \pi w \kappa_2^2 + o(1))t\big) = \exp\big((\mu_w \kappa_2^2 + o(1))t\big).$$

Now, since $a_w = p_w\sqrt{\pi w}$, $\theta_0 = 2\sqrt{\pi w}\kappa_2$ and $\mu_w = p_w \pi w$, we have $\kappa_2 a_w \theta_0 = 2\mu_w \kappa_2^2$. We therefore obtain by Markov's inequality:

$$\begin{aligned}
&\mathbb{P}(h(\beta)^{-1/2} S_{h(\beta)t} \geq -m + \kappa_2 a_w t) \\
&= \mathbb{P}\big(\exp(\theta_0 h(\beta)^{-1/2} S_{h(\beta)t}) \geq \exp(-m\theta_0 + \kappa_2 a_w \theta_0 t)\big) \\
&\leq \exp\big((\mu_w \kappa_2^2 + o(1))t\big) \cdot \exp\big(-2\mu_w \kappa_2^2 t\big) \cdot e^{\theta_0 m} \\
&= \exp\big((-\kappa_2^2 \mu_w + o(1))t\big) \cdot e^{\theta_0 m}.
\end{aligned}$$

Combining with (A.25), we obtain

$$\mathbb{E}\big|\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t}\big| \leq \exp\big(((1-\kappa_2^2)\mu_w + o(1))t\big) \cdot e^{\theta_0 m}.$$

Take κ_3 such that $1 < \kappa_3 < \kappa_2^2$. Then for sufficiently small β,

$$\mathbb{E}\big|\phi_t^{\beta,0} \cap A_{-m+\kappa_2 a_w t}\big| \leq \exp\big((1-\kappa_3)\mu_w t\big) \cdot e^{\theta_0 m}.$$

Combining this with (A.22), (A.23) and (A.24), we obtain for sufficiently small β and

any $m \geq 1$,

$$\limsup_{t\to\infty} \mathbb{P}\big(\lceil \kappa a_w h(\beta)^{1/2} t\rceil e_1 \in \xi^{B_{mR}}_{h(\beta)t}\big) = 0,$$

which yields (A.21), as desired. □

A.1.8 Proof of Lemma 2.8

Lemma A.8. *Set $d(\beta) := \beta^{-1/2}(\log(1/\beta))^{-1}$. Let $\mathcal{A} = \mathcal{A}(\beta)$ denote the collection of finite subsets of $\mathbb{Z}^2 \times \mathbb{Z}_w$ in which points are pairwise separated by at least $d(\beta)$. Set $A^\beta := h(\beta)^{-1/2}A$ for $A \in \mathcal{A}$ and $\mathcal{A}^\beta := h(\beta)^{-1/2}\mathcal{A}$. Then, for any $K > 0$ and $T > 0$,*

$$\sup_{A\in\mathcal{A},\, |A|\leq K} \mathbb{P}\big(\{(\hat{\zeta}^{\beta,A^\beta}_t)_{t\leq T} \neq (\mathring{\psi}^{\beta,A^\beta}_t)_{t\leq T}\} \cup \{\mathring{\psi}^{\beta,A^\beta}_T \notin \mathcal{A}^\beta\}\big) \to 0, \quad \beta \to 0.$$

Proof. Recall that the pruned dual process $\hat{\zeta}_t$ includes any particle from the dual process $\tilde{\zeta}_t$ that has not coalesced with *any other particle in the process* by time $\tau(\beta)$. To show that

$$(\hat{\zeta}^{\beta,A^\beta}_t)_{t\leq T} = (\mathring{\psi}^{\beta,A^\beta}_t)_{t\leq T}$$

with high probability for sufficiently small β, we need to show that any particle in the dual process $\tilde{\zeta}_t$ that does not coalesce with its parent by time $\tau(\beta)$ will, with high probability, (i) not coalesce with any other particle in the process before time $\tau(\beta)$, and (ii) neither coalesce with its parent nor another particle in the process after time $\tau(\beta)$.

Assume that the starting set A has at most K particles, i.e. $|A| \leq K$. On pages 1758-1759 of [54], Durrett and Zähle establish the following for the $w = 1$ case, i.e. for $\mathbb{Z}^2$:

- For any $\varepsilon > 0$, there exists $M = M(\varepsilon, K, T)$ and $\beta_0 > 0$ such that

$$\mathbb{P}(|\hat{\zeta}^{\beta,A^\beta}_T| > M) \leq \varepsilon, \quad \beta \leq \beta_0, \tag{A.26}$$

  i.e. with high probability, the total number of particles in the scaled, pruned dual process at time T is finite.

- If Z^1_t and Z^2_t are independent SSRWs on $\mathbb{Z}^2$ with jump rate 1, $\bar{Z}_t := Z^1_t - Z^2_t$, and

$x = x(\beta) \in \mathbb{Z}^2$ with $||x|| > d(\beta)$, then for any $R > 0$,

$$\mathbb{P}(||\bar{Z}_t|| \leq R \text{ for some } t \leq h(\beta)T \,|\, \bar{Z}_0 = x) \to 0, \quad \beta \to 0. \tag{A.27}$$

- If Z_t^1 and Z_t^2 are started at nearest neighbors, and T_0 denotes the time at which they first meet, then by the local central limit theorem on $\mathbb{Z}^2$,

$$\mathbb{P}(||\bar{Z}_{\tau(\beta)}|| \leq d(\beta) \,|\, T_0 > \tau(\beta)) = O\big((\log(1/\beta))^{-1/2}\big) = o(1), \quad \beta \to 0. \tag{A.28}$$

All statements continue to be true on $\mathbb{Z}^2 \times \mathbb{Z}_w$ for $w > 1$. In addition, following the same argument as used to establish (A.28), we note that if Z_t^1 and Z_t^2 are independent SSRWs on $\mathbb{Z}^2 \times \mathbb{Z}_w$ with jump rate 1, $\bar{Z}_t = Z_t^1 - Z_t^2$, and $g(\beta) = \omega\big(d(\beta)^2\big)$, i.e. $d(\beta)^2/g(\beta) \to 0$ as $\beta \to 0$, then by the local central limit theorem (2.15) on $\mathbb{Z}^2$,

$$\begin{aligned}\mathbb{P}\big(||\bar{Z}_{g(\beta)}|| \leq d(\beta) \,|\, \bar{Z}_0 = x\big) &= \textstyle\sum_{y:||y|| \leq d(\beta)} \mathbb{P}(\bar{Z}_{g(\beta)} = y \,|\, \bar{Z}_0 = x) \\ &= w(2d(\beta)+1)^2 \cdot O\big(g(\beta)^{-1}\big) \\ &= o(1)\end{aligned} \tag{A.29}$$

for all $x \in \mathbb{Z}^2 \times \mathbb{Z}_w$. The fact that $(\hat{\zeta}_t^{\beta,A^\beta})_{t \leq T} = (\mathring{\psi}_t^{\beta,A^\beta})_{t \leq T}$ with high probability (w.h.p.) will now follow from an induction argument similar to the one presented on pages 1758-1759 of [54]. To see that $\mathring{\psi}_T^{\beta,A^\beta} \in \mathcal{A}^\beta$ w.h.p., we first take M so that there are at most M particles w.h.p., by the same argument as in (A.26). In the unscaled process $(\mathring{\psi}_t^A)_{t \leq h(\beta)T}$, the probability that at least one particle gives birth in the last $2\tau(\beta)$ time units is at most $M\big(1 - \exp(-2\beta\tau(\beta))\big) = o(1)$, since $\tau(\beta) = (1/\beta)\big(1/\sqrt{\log(1/\beta)}\big)$. If at time $h(\beta)T - 2\tau(\beta)$, there exists a parent-daughter pair whose decision period has not yet passed, the daughter will be introduced to the pruned process by time $h(\beta)T - \tau(\beta)$, given that it does not coalesce with its parent. Then, applying (A.29) to each pair of particles alive at time $h(\beta)T - \tau(\beta)$, and using the fact that $\tau(\beta) = \omega\big(d(\beta)^2\big)$, we see that the particles in $(\mathring{\psi}_t^A)_{t \leq h(\beta)T}$ will be pairwise separated by at least $d(\beta)$ at time $h(\beta)T$ w.h.p. It follows that $\mathring{\psi}_T^{\beta,A^\beta} \in \mathcal{A}^\beta$ w.h.p. $\square$

A.1.9 Proof of Lemma 2.9

Lemma A.9. *For $0 < \theta < 1$ and $L > 0$, define (Fig. 2.6b)*

$$I_0^\theta := [-(1/2)\theta a_w L, (1/2)\theta a_w L]^2 \times \mathbb{Z}_w,$$
$$I_k^\theta := I_0^\theta + k \cdot \theta a_w L e_1,\ k \in \mathbb{Z},$$

and

$$\mathcal{A}^{\beta,\theta,K,k} := \{A^\beta \in \mathcal{A}^\beta : |A^\beta \cap I_k^\theta| \geq K\},$$

with $\mathcal{A}^\beta$ defined as in Lemma 2.8. Let $(\hat{\zeta}_t^{\beta,A^\beta,\theta})_{t\geq 0}$ denote a pruning of $(\hat{\zeta}_t^{\beta,A^\beta})_{t\geq 0}$ with particles killed as soon as they exit the box I_Δ^θ with $I_\Delta^\theta := [-2\theta a_w L, 2\theta a_w L]^2 \times \mathbb{Z}_w$. Then, for any $2/3 < \theta < 1$ and $\varepsilon > 0$, there exist $L = L(\theta) > 0$, $K = K(\theta, \varepsilon) > 0$ and $\beta_0 = \beta_0(\theta, \varepsilon) > 0$ such that for any $A^\beta \in \mathcal{A}^{\beta,\theta,K,0}$ with $|A^\beta| = K$, and any $\beta \leq \beta_0$,

$$\mathbb{P}(\hat{\zeta}_L^{\beta,A^\beta,\theta} \in \mathcal{A}^{\beta,\theta,K,k}) \geq 1 - \varepsilon, \quad k \in \{-1, 1\}.$$

Proof. Let $L > 0$ and $K > 0$ be constants to be selected later. Furthermore, let $(\mathring{\psi}_t^{\beta,A^\beta,\theta})_{t\geq 0}$ denote a pruning of $(\mathring{\psi}_t^{\beta,A^\beta})_{t\geq 0}$ with particles killed as soon as they exit I_Δ^θ. By Lemma 2.8, we have for all $A^\beta \in \mathcal{A}^{\beta,\theta,K,0}$ with $|A^\beta| = K$ and sufficiently small β,

$$\begin{aligned}
&\mathbb{P}(\hat{\zeta}_L^{\beta,A^\beta,\theta} \in \mathcal{A}^{\beta,\theta,K,k}) \\
&\geq \mathbb{P}\big(|\mathring{\psi}_L^{\beta,A^\beta,\theta} \cap I_k^\theta| \geq K, (\hat{\zeta}_t^{\beta,A^\beta})_{t\leq L} = (\mathring{\psi}_t^{\beta,A^\beta})_{t\leq L}, \mathring{\psi}_T^{\beta,A^\beta} \in \mathcal{A}^\beta\big) \\
&\geq \mathbb{P}\big(|\mathring{\psi}_L^{\beta,A^\beta,\theta} \cap I_k^\theta| \geq K\big) - \mathbb{P}\big(\{(\hat{\zeta}_t^{\beta,A^\beta})_{t\leq L} \neq (\mathring{\psi}_t^{\beta,A^\beta})_{t\leq L}\} \cup \{\mathring{\psi}_T^{\beta,A^\beta} \notin \mathcal{A}^\beta\}\big) \\
&\geq \mathbb{P}\big(|\mathring{\psi}_L^{\beta,A^\beta,\theta} \cap I_k^\theta| \geq K\big) - \varepsilon/4, \quad k \in \{-1, 1\}.
\end{aligned} \tag{A.30}$$

Define a simple BRW $(\psi_t^\beta)_{t\geq 0}$ in terms of $(\mathring{\psi}_t^\beta)_{t\geq 0}$ analogously to how $(\phi_t^\beta)_{t\geq 0}$ is defined in terms of $(\mathring{\phi}_t^\beta)_{t\geq 0}$ in Section 2.3.4. Take θ_1 so that $2/3 < \theta_1 < \theta$, and let $(\psi_t^{\beta,A^\beta,\theta_1})_{t\geq 0}$ denote a pruning of $(\psi_t^{\beta,A^\beta})_{t\geq 0}$ with particles killed as soon as they exit $I_\Delta^{\theta_1}$. Now, consider the event

$$\{|\psi_L^{\beta,A^\beta,\theta_1} \cap I_k^{\theta_1}| \geq K\}, \quad k \in \{-1, 1\}.$$

On the above event, pick K particles from $\psi_L^{\beta,A^\beta,\theta_1} \cap I_k^{\theta_1}$, and consider one such particle Y_L^β. Lemma 2.5 implies that the distance between the path of this particle (and its

ancestors) and the corresponding particle $\mathring{Y}_L^\beta$ in $\mathring{\psi}_L^\beta$ (and its ancestors) up until time L is upper bounded by $(1/2)(\theta-\theta_1)a_w L$ with high probability given sufficiently small β. This implies that w.h.p., $\mathring{Y}_L^\beta$ and its ancestors stay within I_Δ^θ during $[0,L]$, and $\mathring{Y}_L^\beta$ ends up in I_k^θ. Thus, for sufficiently small β,

$$\mathbb{P}\big(|\psi_L^{\beta,A^\beta,\theta_1}\cap I_k^{\theta_1}|\geq K\big)\leq \mathbb{P}(|\mathring{\psi}_L^{\beta,A^\beta,\theta}\cap I_k^{\theta}|\geq K)+\varepsilon/4,\quad k\in\{-1,1\},$$

which yields by (A.30) for sufficiently small β,

$$\mathbb{P}(\hat{\zeta}_L^{\beta,A^\beta,\theta}\in\mathcal{A}^{\beta,\theta,K,k})\geq \mathbb{P}\big(|\psi_L^{\beta,A^\beta,\theta_1}\cap I_k^{\theta_1}|\geq K\big)-\varepsilon/2,\quad k\in\{-1,1\}. \tag{A.31}$$

We now wish to estimate the probability on the right-hand side of (A.31). Recall that the branching rate of ψ_t^β is $\mu_w+o(1)$ by Lemma 2.3. Then, for any $z^\beta\in I_0^{\theta_1}\cap h(\beta)^{-1/2}(\mathbb{Z}^2\times\mathbb{Z}_w)$,

$$\mathbb{E}\big|\psi_L^{\beta,z^\beta,\theta_1}\cap I_k^{\theta_1}\big|=\exp\big((\mu_w+o(1))L\big)\cdot\mathbb{P}\big(Z_L^{\beta,z^\beta,\theta_1}\in I_k^{\theta_1}\big),\quad k\in\{-1,1\}, \tag{A.32}$$

where $(Z_t^{\beta,z^\beta})_{t\geq 0}$ is the scaled version of the SSRW $(Z_t^z)_{t\geq 0}$ on $\mathbb{Z}^2\times\mathbb{Z}_w$ with jump rate 1, and $(Z_t^{\beta,z^\beta,\theta_1})_{t\geq 0}$ is a pruned version where the walk is killed if it exits $I_\Delta^{\theta_1}$ (see e.g. (7.5) of [52] for why (A.32) is true). For the probability on the right-hand side, note that

$$\mathbb{P}\big(Z_L^{\beta,z^\beta,\theta_1}\in I_k^{\theta_1}\big)=\mathbb{P}\big(Z_L^{\beta,z^\beta}\in I_k^{\theta_1},Z_t^{\beta,z^\beta}\in I_\Delta^{\theta_1}\text{ for all }t\leq L\big),\quad k\in\{-1,1\}.$$

Without loss of generality, set $k:=1$. Take $\delta>0$ so that for any $z^\beta=(z_1^\beta,z_2^\beta,z_3^\beta)\in I_0^{\theta_1}\cap h(\beta)^{-1/2}(\mathbb{Z}^2\times\mathbb{Z}_w)$ with $z_1^\beta\geq 0$ and $z_2^\beta\geq 0$,

$$(z^\beta+([-\delta,0]^2\times\mathbb{Z}_w))\subseteq I_0^{\theta_1}.$$

For such z^β and $\delta > 0$, and $J_\Delta^{\theta_1} := [-(3/2)\theta_1 a_w L, (3/2)\theta_1 a_w L]^2 \times \mathbb{Z}_w$,

$$\begin{aligned}
&\mathbb{P}\big(Z_L^{\beta,z^\beta,\theta_1} \in I_1^{\theta_1}\big) \\
&= \mathbb{P}\big(Z_L^{\beta,z^\beta} \in I_1^{\theta_1}, Z_t^{\beta,z^\beta} \in I_\Delta^{\theta_1} \text{ for all } t \le L\big) \\
&\ge \mathbb{P}\big(Z_L^{\beta,0} \in ([-\delta,0]^2 \times \mathbb{Z}_w) + \theta_1 a_w L e_1, Z_t^{\beta,0} \in J_\Delta^{\theta_1} \text{ for all } t \le L\big) \\
&\ge \mathbb{P}\big(Z_L^{\beta,0} \in ([-\delta,0]^2 \times \mathbb{Z}_w) + \theta_1 a_w L e_1\big) - \mathbb{P}(Z_t^{\beta,0} \notin J_\Delta^{\theta_1} \text{ for some } t \le L), \qquad \text{(A.33)}
\end{aligned}$$

where we use the translation invariance of $(Z_t^\beta)_{t\ge 0}$. To analyze the latter term in (A.33), write $Z_t = (Z_{1,t}, Z_{2,t}, Z_{3,t})$, where $Z_{1,t}$ and $Z_{2,t}$ are i.i.d. copies of the SSRW on $\mathbb{Z}$ with jump rate $p_w/2$ each, and $Z_{3,t}$ is the SSRW on $\mathbb{Z}_w$ with jump rate $1 - p_w$, where p_w is defined as in (2.6). Then note that

$$\begin{aligned}
&\mathbb{P}(Z_t^{\beta,0} \notin J_\Delta^{\theta_1} \text{ for some } t \le L) \\
&\le 2\,\mathbb{P}(Z_{1,t}^{\beta,0} \notin [-(3/2)\theta_1 a_w L, (3/2)\theta_1 a_w L] \text{ for some } t \le L) \\
&\le 4\,\mathbb{P}\big(h(\beta)^{-1/2} \sup_{t \le h(\beta)L} Z_{1,t}^0 > (3/2)\theta_1 a_w L\big).
\end{aligned}$$

By assumption, $(3/2)\theta_1 > 1$. The same argument we used to analyze (A.25) in the proof of Lemma 2.7 (take $\kappa_2 := (3/2)\theta_1 > 1$ and $m := 0$, and use Doob's inequality to handle the supremum) will show that we can take $\gamma_2 > 0$ so that for sufficiently small β,

$$\exp\big((\mu_w + o(1))L\big) \cdot \mathbb{P}\big(h(\beta)^{-1/2} \sup_{t \le h(\beta)L} Z_{1,t}^0 > (3/2)\theta_1 a_w L\big) \le \exp(-\gamma_2 \mu_w L). \qquad \text{(A.34)}$$

To analyze the former term in (A.33), note that the local central limit theorem (2.15) on $\mathbb{Z}^2$ implies that for all $y \in \mathbb{Z}^2$,

$$\mathbb{P}\big((Z_{1,h(\beta)L}^0, Z_{2,h(\beta)L}^0) = y\big) = \big(p_w \pi h(\beta) L\big)^{-1} \exp\big(-||y||^2/(p_w h(\beta) L)\big) + o_L\big(h(\beta)^{-1}\big),$$

where o_L signifies that the error term depends on L. Since

$$\begin{aligned}
&\mathbb{P}\big(Z_L^{\beta,0} \in ([-\delta,0]^2 \times \mathbb{Z}_w) + \theta_1 a_w L e_1\big) \\
&= \mathbb{P}\big((Z^0_{1,h(\beta)L}, Z^0_{2,h(\beta)L}) \in ([-\delta h(\beta)^{1/2},0]^2 + \theta_1 a_w h(\beta)^{1/2} L e_1)\big) \\
&= \textstyle\sum_{y \in ([-\delta h(\beta)^{1/2},0]^2 + \theta_1 a_w h(\beta)^{1/2} L e_1)} \mathbb{P}\big((Z^0_{1,h(\beta)L}, Z^0_{2,h(\beta)L}) = y\big),
\end{aligned}$$

and the number of terms in the sum is of order $h(\beta)$, we obtain for some $C > 0$,

$$\mathbb{P}\big(Z_L^{\beta,0} \in ([-\delta,0]^2 \times \mathbb{Z}_w) + \theta_1 a_w L e_1\big) \geq (C/L) \cdot \exp(-\theta_1^2 \mu_w L) + o_L(1),$$

where we use that $a_w^2/p_w = \mu_w$ since $a_w = p_w\sqrt{\pi w}$ and $\mu_w = p_w \pi w$. Since $\theta_1 < 1$, we can then find $\gamma_1 > 0$ so that for sufficiently small β,

$$\begin{aligned}
&\exp\big((\mu_w + o(1))L\big) \cdot \mathbb{P}\big(Z_L^{\beta,0} \in ([-\delta,0]^2 \times \mathbb{Z}_w) + \theta_1 a_w L e_1\big) \\
&\geq (C/L) \cdot \exp(\gamma_1 \mu_w L). \qquad \text{(A.35)}
\end{aligned}$$

Combining (A.34) and (A.35) with (A.32) and (A.33), we obtain for sufficiently small β,

$$\mathbb{E}|\psi_L^{\beta,z^\beta,\theta_1} \cap I_1^{\theta_1}| \geq (C/L) \cdot \exp\big(\gamma_1 \mu_w L\big) - 4\exp\big(-\gamma_2 \mu_w L\big),$$

where $\gamma_1, \gamma_2 > 0$. Since the former term tends to ∞ as $L \to \infty$ and the latter term tends to 0 as $L \to \infty$, we can select $L = L(\theta_1)$ large enough so that

$$\mathbb{E}|\psi_L^{\beta,z^\beta,\theta_1} \cap I_1^{\theta_1}| \geq 2, \quad z^\beta \in I_0^{\theta_1} \cap h(\beta)^{-1/2}(\mathbb{Z}^2 \times \mathbb{Z}_w),\ z_1^\beta \geq 0,\ z_2^\beta \geq 0.$$

The same is true for $k = -1$, as well as z^β in the second, third and fourth quadrant of $I_0^{\theta_1} \cap h(\beta)^{-1/2}(\mathbb{Z}^2 \times \mathbb{Z}_w)$. Now, using the same argument as in Durrett and Zähle ([54], p. 1760), for the given L and any $\varepsilon > 0$, we can select $K = K(\theta_1, \varepsilon) > 0$ large enough so that for all $A^\beta \in \mathcal{A}^{\beta,\theta,K,0}$ with $|A^\beta| = K$ and sufficiently small β,

$$\mathbb{P}(|\psi_L^{\beta,A^\beta,\theta_1} \cap I_k^{\theta_1}| < K) \leq \varepsilon/2, \quad k \in \{-1, 1\}.$$

Combining with (A.31), which holds for sufficiently small β given fixed K and L, we

obtain the desired result. □

A.1.10 Proof of Lemma 2.10

Lemma A.10. *Define* $\tau_\varnothing^A := \min\{t \geq 0 : \xi_t^A = \varnothing\}$ *for* $A \subseteq \mathbb{Z}^2 \times \mathbb{Z}_w$. *For each* $2/3 < \rho < 1$, *there exists a constant* $L > 0$ *and a family of random variables* $(S_\beta)_{\beta>0}$, *with* $\mathbb{P}(S_\beta < \infty | \tau_\varnothing^0 = \infty) = 1$ *for each* $\beta > 0$, *so that*

$$\lim_{\beta\to 0} \liminf_{n\to\infty} \mathbb{P}\big(\xi^0_{S_\beta+2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing \,\big|\, \tau_\varnothing^0 = \infty\big) = 1,$$

where $a_w := p_w\sqrt{\pi w}$.

Proof. Take $\varepsilon > 0$. We segment the proof into three main steps.

Step 1: Remove conditioning on nonextinction. Let C_r denote a box in $\mathbb{Z}^2 \times \mathbb{Z}_w$, centered at 0 with side lengths $10r$ and $2r$, i.e.

$$C_r := \{(x_1, x_2, x_3) \in \mathbb{Z}^2 \times \mathbb{Z}_w : |x_1| \leq 5r, |x_2| \leq r\},$$

and set $\sigma_r := \inf\{t \geq 0 : C_r \subseteq \xi_t^0\}$. As in the proof of Lemma 2.7, we note that if $r = r(\beta) = M \cdot h(\beta)^{1/2}$ for some $M > 0$, then $\mathbb{P}(\tau_\varnothing^{C_r} < \infty) = o(1)$ as $\beta \to 0$, and $\mathbb{P}(\sigma_r < \infty \,|\, \tau_\varnothing^0 = \infty) = 1$, so $\mathbb{P}(\sigma_r < \infty) \geq \mathbb{P}(\tau_\varnothing^0 = \infty)$. Let $L > 0$ and $\rho < \rho_1 < 1$, and set

$$\gamma := (1/2)\rho_1 a_w L \quad \text{and} \quad r := \gamma h(\beta)^{1/2}.$$

We then get by the strong Markov property and the monotonicity property (2.10) of ξ_t, for any $n \geq 1$ and sufficiently small β,

$$\begin{aligned}
&\mathbb{P}\big(\xi^0_{\sigma_r+2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing, \tau_\varnothing^0 = \infty\big) \\
&= \textstyle\sum_{C_r \subseteq \Lambda} \int_s \mathbb{P}\big(\xi^\Lambda_{2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing, \tau_\varnothing^\Lambda = \infty\big)\mathbb{P}\big(\sigma_r \in ds, \xi^0_{\sigma_r} = \Lambda\big) \\
&\geq \mathbb{P}\big(\xi^{C_r}_{2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing, \tau_\varnothing^{C_r} = \infty\big)\mathbb{P}(\sigma_r < \infty) \\
&\geq \big(\mathbb{P}\big(\xi^{C_r}_{2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing\big) - \mathbb{P}\big(\tau_\varnothing^{C_r} < \infty\big)\big)\mathbb{P}(\sigma_r < \infty) \\
&\geq \big(\mathbb{P}\big(\xi^{C_r}_{2nLh(\beta))} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing\big) - \varepsilon\big)\mathbb{P}(\tau_\varnothing^0 = \infty),
\end{aligned}$$

from which it follows that

$$\begin{aligned}&\mathbb{P}\big(\xi^0_{\sigma_r+2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2},\infty\big)e_1 \neq \varnothing \,\big|\, \tau^0_\varnothing = \infty\big)\\ &\geq \mathbb{P}\big(\xi^{C_r}_{2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2},\infty\big)e_1 \neq \varnothing\big) - \varepsilon. \end{aligned} \tag{A.36}$$

Step 2: Introduce duality. Let $K > 0$ be a constant to be selected later, and set $d(\beta) := \beta^{-1/2}(\log(1/\beta))^{-1}$ as in Lemma 2.8. Define

$$\begin{aligned}A_0 = A_0(\beta) &:= \lceil d(\beta)\rceil \cdot [\![-K+1,0]\!] \cdot e_1,\\ A_0^\beta &:= h(\beta)^{-1/2}A_0,\end{aligned}$$

with $[\![m,n]\!] = \{m, m+1, \ldots, n\}$ for integers $m < n$ (with possibly $n = \infty$). Note that A_0 has K points which are pairwise separated by at least $d(\beta)$. Since $d(\beta) = o\big(\sqrt{h(\beta)}\big)$, it follows that $A_0^\beta \in \mathcal{A}^{\beta,\rho_1,K,0}$ for sufficiently small β, where $\mathcal{A}^{\beta,\rho_1,K,0}$ is defined as in Lemma 2.9. Next define

$$g(n) = g(n,\beta) := (2\rho_1 a_w Lh(\beta)^{1/2})^{-1}\lceil 2n\rho a_w Lh(\beta)^{1/2}\rceil,$$

and note that for fixed β, $g(n)/n \to \rho/\rho_1 < 1$ as $n \to \infty$. Continuing on from (A.36), we obtain using the duality relation (2.8) between ξ_t and $\tilde{\zeta}_t$, the monotonicity property (2.10) of $\tilde{\zeta}_t$, the translation invariance (2.11) and symmetry property (2.12) of $\tilde{\zeta}_t$, and the definition (2.17) of the scaled dual process $\tilde{\zeta}^\beta_t$,

$$\begin{aligned}&\mathbb{P}\big(\xi^{C_r}_{2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2},\infty\big)e_1 \neq \varnothing\big)\\ &= \mathbb{P}\big(\tilde{\zeta}^{[\![\lceil 2n\rho a_w Lh(\beta)^{1/2}\rceil,\infty]\!]e_1}_{2nLh(\beta)} \cap C_r \neq \varnothing\big)\\ &\geq \mathbb{P}\big(\tilde{\zeta}^{-A_0+\lceil 2n\rho a_w Lh(\beta)^{1/2}\rceil e_1}_{2nLh(\beta)} \cap C_r \neq \varnothing\big)\\ &= \mathbb{P}\big(\tilde{\zeta}^{A_0}_{2nLh(\beta)} \cap (C_r + \lceil 2n\rho a_w Lh(\beta)^{1/2}\rceil e_1) \neq \varnothing\big)\\ &= \mathbb{P}\big(\tilde{\zeta}^{\beta,A_0^\beta}_{2nL} \cap \big(([-5\gamma,5\gamma]\times[-\gamma,\gamma]\times\mathbb{Z}_w) + 2g(n)\cdot\rho_1 a_w Le_1\big) \neq \varnothing\big).\end{aligned} \tag{A.37}$$

Recall that $\gamma = (1/2)\rho_1 a_w L$, and define $I_0^{\rho_1} := [-\gamma,\gamma]^2 \times \mathbb{Z}_w$ and $I_k^{\rho_1} := I_0^{\rho_1} + k\cdot 2\gamma e_1$

for $k \in \mathbb{Z}$ as in Lemma 2.9. Then

$$\begin{aligned}
&\mathbb{P}\big(\tilde{\zeta}_{2nL}^{\beta,A_0^\beta} \cap \big(([-5\gamma,5\gamma]\times[-\gamma,\gamma]\times\mathbb{Z}_w) + 2g(n)\cdot\rho_1 a_w L e_1\big) \neq \varnothing\big)\\
&\geq \mathbb{P}\big(\tilde{\zeta}_{2nL}^{\beta,A_0^\beta} \cap \big(I_0^{\rho_1} + 2\lceil g(n)\rceil\cdot\rho_1 a_w L e_1\big) \neq \varnothing\big)\\
&\geq \mathbb{P}\big(\hat{\zeta}_{2nL}^{\beta,A_0^\beta} \cap I_{2\lceil g(n)\rceil}^{\rho_1} \neq \varnothing\big),
\end{aligned} \tag{A.38}$$

where in the last step, we use the lower-bounding property (2.19) of the pruned dual process $\hat{\zeta}_t$. We are now ready to apply the percolation construction of Lemma 2.9.

Step 3: Compare with oriented percolation. Let $(\hat{\zeta}_t^{\beta,A_0^\beta,\rho_1})_{t\geq 0}$ be a pruning of $(\hat{\zeta}_t^{\beta,A^\beta})_{t\geq 0}$ with particles killed as soon as they exit the box $I_\Delta^{\rho_1}$ with $I_\Delta^{\rho_1} := [-2\rho_1 a_w L, 2\rho_1 a_w L]^2 \times \mathbb{Z}_w$. Note first that for any $K > 0$,

$$\mathbb{P}\big(\hat{\zeta}_{2nL}^{\beta,A_0^\beta} \cap I_{2\lceil g(n)\rceil}^{\rho_1} \neq \varnothing\big) \geq \mathbb{P}\big(\hat{\zeta}_{2nL}^{\beta,A_0^\beta,\rho_1} \in \mathcal{A}^{\beta,\rho_1,K,2\lceil g(n)\rceil}\big). \tag{A.39}$$

By assumption, $2/3 < \rho_1 < 1$, so by Lemma 2.9, we can choose K and L so that for any $A^\beta \in \mathcal{A}^{\beta,\rho_1,K,0}$ with $|A^\beta| = K$ and sufficiently small β,

$$\mathbb{P}(\hat{\zeta}_L^{\beta,A^\beta,\rho_1} \in \mathcal{A}^{\beta,\rho_1,K,k}) \geq 1 - \varepsilon/2, \quad k \in \{-1,1\}.$$

Now set

$$X_n := \{k \in \mathbb{Z} : k+n \text{ even and } \hat{\zeta}_{nL}^{\beta,A_0^\beta,\rho_1} \in \mathcal{A}^{\beta,\rho_1,K,k}\}, \quad n \geq 0.$$

By Theorem 4.3 of [52], X_n dominates a one-dependent oriented percolation process $\{\omega_n^0\}_{n\geq 0}$ with density $\geq 1-\varepsilon$ and $\omega_0^0 = \{0\}$, i.e. $\omega_n^0 \subseteq X_n$ for all n. Let Ω_∞^0 denote the event $\{|\bigcup_n \omega_n^0| = \infty\}$, i.e. the event that percolation occurs, and let $l_n^0 = \min \omega_n^0$ (resp. $r_n^0 = \max \omega_n^0$) denote the left (resp. right) edge of the process. Take ρ_2 so that $\rho/\rho_1 < \rho_2 < 1$. By Theorem 4.1 of [52], we have for sufficiently small ε,

$$\mathbb{P}(\Omega_\infty^0) \geq 1 - 55\varepsilon^{1/9}, \tag{A.40}$$

and by Theorem 3.21 on page 300 of [107], we have for sufficiently small ε,

$$\mathbb{P}(r_n^0/n \geq \rho_2 \,|\, \Omega_\infty^0) \geq 1 - 3^{-n+1}. \tag{A.41}$$

Since for fixed β, $g(n)/n \to \rho/\rho_1$ as $n \to \infty$, we further have for sufficiently large n,

$$\lceil g(n) \rceil / n \leq \rho_2. \tag{A.42}$$

Continuing on from (A.39), we obtain for sufficiently small β and sufficiently large n,

$$\begin{aligned}\mathbb{P}\big(\hat{\zeta}_{2nL}^{\beta, A_0^\beta, \rho_1} \in \mathcal{A}^{\beta,\rho_1,K,2\lceil g(n)\rceil}\big) &\geq \mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^0) \\ &\geq \mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^0, r_{2n}^0 \geq 2\rho_2 n, \Omega_\infty^0).\end{aligned}$$

On Ω_∞^0, we have $\omega_n^0 = \omega_n^{2\mathbb{Z}} \cap [\![l_n^0, r_n^0]\!]$ (see Section 8 of [47]). By (A.42), we can therefore write for sufficiently large n,

$$\begin{aligned}\mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^0, r_{2n}^0 \geq 2\rho_2 n, \Omega_\infty^0) &= \mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^{2\mathbb{Z}}, r_{2n}^0 \geq 2\rho_2 n, \Omega_\infty^0) \\ &\geq \mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^{2\mathbb{Z}}) + \mathbb{P}(r_{2n}^0 \geq 2\rho_2 n, \Omega_\infty^0) - 1 \\ &= \mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^{2\mathbb{Z}}) + \mathbb{P}(r_{2n}^0 \geq 2\rho_2 n \,|\, \Omega_\infty^0)\mathbb{P}(\Omega_\infty^0) - 1.\end{aligned}$$

Furthermore, ω_n^0 is self-dual (see Section 8 of [47]), so

$$\mathbb{P}(2\lceil g(n)\rceil \in \omega_{2n}^{2\mathbb{Z}}) = \mathbb{P}(\omega_{2n}^{2\lceil g(n)\rceil} \neq \varnothing) = \mathbb{P}(\omega_{2n}^0 \neq \varnothing) \geq \mathbb{P}(\Omega_\infty^0).$$

By (A.40) and (A.41), we therefore obtain for sufficiently small β and sufficiently large n,

$$\mathbb{P}\big(\hat{\zeta}_{2nL}^{\beta, A_0^\beta, \rho_1} \in \mathcal{A}^{\beta,\rho_1,K,2\lceil g(n)\rceil}\big) \geq (2 - 3^{-2n+1})(1 - 55\varepsilon^{1/9}) - 1.$$

Combining this with (A.36), (A.37), (A.38) and (A.39), we finally obtain for sufficiently

small β and sufficiently large n,

$$\begin{aligned}&\mathbb{P}\big(\xi^0_{\sigma_r+2nLh(\beta)} \cap \big[2n\rho a_w Lh(\beta)^{1/2}, \infty\big)e_1 \neq \varnothing \,\big|\, \tau^0_\varnothing = \infty\big)\\ &\geq \big((2-3^{-2n+1})(1-55\varepsilon^{1/9})-1\big)-\varepsilon,\end{aligned}$$

and the result follows. □

A.1.11 Extension of Bramson-Griffeath shape theorem

In this section, we discuss how the Bramson-Griffeath shape theorem on $\mathbb{Z}^2$ can be extended to $\mathbb{Z}^2 \times \mathbb{Z}_w$. Bramson and Griffeath's proof is split across two papers. In [20], they show that with probability 1, the biased voter model on $\mathbb{Z}^d$, conditioned on nonextinction, eventually contains a ball which expands linearly in time. Then, in [19], they use results from [20] to show that the biased voter model on $\mathbb{Z}^d$ satisfies five conditions formulated by Richardson [132], which together guarantee the existence of a linearly-expanding asymptotic shape.

The key to extending the BG shape theorem to $\mathbb{Z}^2 \times \mathbb{Z}_w$ is to interpret Richardson's conditions with the appropriate notion of spatial scaling and distance. Richardson's conditions deal with the "first infection times" of distant points of the form x/Δ for small $\Delta > 0$, i.e. the times at which the process first reaches these points. For long-run growth to be linear, the infection time of x/Δ should be of order $1/\Delta$ as $\Delta \to 0$, i.e. the infection time of x/Δ scaled by Δ should be finite. In [132], Richardson proposes five conditions on these scaled infection times, which together imply the existence of a linearly-expanding asymptotic shape. On $\mathbb{Z}^2 \times \mathbb{Z}_w$, the mutant clone will expand without bound along the first two coordinates, while the third coordinate remains bounded throughout. As a result, we have defined the scalar multiplication (2.2) and the distance from the origin (2.3) in the main text as only applying to the first two coordinates, and we view the scaled points x/Δ as living on $\mathbb{R}^2 \times \mathbb{Z}_w$.

Once Richardson's conditions have been interpreted appropriately, Bramson and Griffeath's arguments on $\mathbb{Z}^2$ extend naturally to $\mathbb{Z}^2 \times \mathbb{Z}_w$ for the most part, since they largely rely on fundamental properties of the biased voter model, such as the ones presented in (2.9)-(2.12) of Section 2.3.1 and the strong Markov property. Where modifications are necessary, one can generally focus on movement along the first two coordinates

and use the same arguments as on $\mathbb{Z}^2$, taking into account that the SSRW on $\mathbb{Z}^2 \times \mathbb{Z}_w$ takes a step in the $\mathbb{Z}^2$-direction with probability p_w given by (2.6). To avoid repeating the arguments, we will focus our discussion on how to interpret Richardson's conditions on $\mathbb{Z}^2 \times \mathbb{Z}_w$, and on how Richardson's conditions give rise to an asymptotic shape of the form $D = \bigcup_{i \in \mathbb{Z}_w} (X \times \{i\})$ given in the main text.

Richardson's conditions for the existence of an asymptotic shape

We begin by stating a slightly stronger version of Richardson's conditions on $\mathbb{Z}^d$, as presented by Bramson and Griffeath in [19]. First, let $(\bar{\xi}_t^0)_{t \geq 0}$ be the biased voter model on $\mathbb{Z}^d$ started at the origin and conditioned on nonextinction, i.e. the process $(\xi_t^0 | \tau_\varnothing^0 = \infty)_{t \geq 0}$. For $x \in \mathbb{Z}^d$, define the first infection time of x as

$$t(x) := \inf\{t : x \in \bar{\xi}_t^0\},$$

with $\inf \varnothing = \infty$. Then introduce the scaled infection time

$$t_\Delta(x) := \Delta \cdot t(x/\Delta),$$

where $\Delta > 0$. Note that x/Δ is in general not in $\mathbb{Z}^d$. We therefore extend the definition of $t(y)$ so that for $y \in \mathbb{R}^d$, it takes the same value as the nearest lattice location in $\mathbb{Z}^d$.

When stating Richardson's conditions, we use the same symbol $\|\cdot\|$ for the Euclidean norm on $\mathbb{R}^d$ as for the distance from the origin (2.3) on $\mathbb{R}^2 \times \mathbb{Z}_w$, since they coincide for $d = 2$ and $w = 1$. Let $V(x, y)$ denote a random function of $x, y \in \mathbb{R}^d$ which satisfies

$$\mathbb{E}[V^2(x, y)] = O(||x + y||),$$

and define

$$V_\Delta(x, y) := \Delta \cdot V(x/\Delta, y/\Delta).$$

Let $\varepsilon_i(\Delta)$ denote an $O(\Delta)$ function, and $\varepsilon_i(\Delta, x)$ denote an $O(\Delta)$ function that depends on x. Let $x, y \in \mathbb{R}^d$. Richardson's conditions are, in Bramson and Griffeath's formulation:

(A1) (Scaled infection times are essentially subadditive.) For all $\alpha \in \mathbb{R}$:

$$\mathbb{P}(t_\Delta(x+y) \leq \alpha) \geq \mathbb{P}(t_\Delta(x) + s_\Delta(y) + V_\Delta(x,y) \leq \alpha),$$

where $s_\Delta(y)$ is a copy of $t_\Delta(y)$ which is independent of $t_\Delta(x)$.

(A2) (Long-run growth is at most linear.) For some $L > 0$,

$$\mathbb{E}[t_\Delta(x)] \geq ||x||/L + \varepsilon_1(\Delta, x).$$

(A3) (Nearby sites tend to be infected at similar times.) For some $r' > 0$,

$$\mathbb{P}(\{y : ||x-y|| \leq r'\delta\} \subseteq \{y : |t_\Delta(x) - t_\Delta(y)| \leq \delta\}) \geq 1 - \varepsilon_2(\Delta, \delta).$$

(A4) (Moment bound on scaled infection times.) $\mathbb{E}[t_\Delta^2(x)]$ exists, and for some $r > 0$,

$$\mathbb{E}[t_\Delta^2(x)] \leq \|x/r\|^2 + \varepsilon_3(\Delta).$$

(A5) (Symmetry condition.)

$$\mathbb{E}[t_\Delta(x)] \leq \mathbb{E}[t_\Delta(-x)] + \varepsilon_4(\Delta, x).$$

Under slightly weaker conditions, Richardson shows in [132] that if (A1)-(A5) hold for a given growth process on $\mathbb{R}^d$, then $\lim_{\Delta\to 0} \mathbb{E}[t_\Delta(x)]$ exists. He defines

$$N(x) := \lim_{\Delta\to 0} \mathbb{E}[t_\Delta(x)], \tag{A.43}$$

and shows that N is a norm on $\mathbb{R}^d$. He also shows that $\mathrm{Var}[t_\Delta(x)] = \varepsilon_5(\Delta, x)$, which implies that $\lim_{\Delta\to 0} t_\Delta(x) = N(x)$ in probability. Let $D'_R := \{x \in \mathbb{R}^d : N(x) \leq R\}$ be the ball of radius R under N. Richardson's main result is that if (A1)-(A5) hold for a given growth process, then for any $\varepsilon > 0$, there exists $t_* < \infty$ so that

$$\mathbb{P}\big(D'_{(1-\varepsilon)t} \cap \mathbb{Z}^d \subseteq \{x : t(x) \leq t\} \subseteq D'_{(1+\varepsilon)t}\big) \geq 1 - \varepsilon, \quad t \geq t_*. \tag{A.44}$$

On $\mathbb{R}^2 \times \mathbb{Z}_w$, we reinterpret Richardson's conditions (A1)-(A5) in terms of the scalar multiplication in (2.2) and the distance function in (2.3). Set $e_3 = (0, 0, 1)$. By definition, $\mathbb{E}[t_\Delta(e_3)] = \Delta\mathbb{E}[t(e_3)] = O(\Delta)$, and by condition (A4), $\mathbb{E}[t_\Delta^2(e_3)] = O(\Delta)$. Thus, the effect of jumping between layers is insignificant under the scaling by Δ. In particular, $N(e_3) = 0$, so N does not separate points on $\mathbb{R}^2 \times \mathbb{Z}_w$, which is a consequence of the fact that $\|\cdot\|$ on $\mathbb{R}^2 \times \mathbb{Z}_w$ does not separate points. However, N satisfies the triangle inequality and $N(tx) = |t|N(x)$ for $t \in \mathbb{R}$ and $x \in \mathbb{R}^2 \times \mathbb{Z}_w$, i.e. N has the properties of a seminorm. (We do not explicitly refer to N as a seminorm since $\mathbb{R}^2 \times \mathbb{Z}_w$ is not a vector space.) From the triangle inequality, it follows that $N(x + e_3) = N(x)$ for all $x \in \mathbb{R}^2 \times \mathbb{Z}_w$.

Apart from the above, Richardson's proof that (A.44) follows from (A1)-(A5) extends to $\mathbb{R}^2 \times \mathbb{Z}_w$ for the most part. A slight modification is needed in his Lemma 5, where he makes use of the topology on $\mathbb{R}^d$. The lemma states that the convergence $\lim_{\Delta\to 0} \mathbb{E}[t_\Delta(x)] = N(x)$ is uniform in x on any bounded set around the origin. On $\mathbb{R}^2 \times \mathbb{Z}_w$, we can use Richardson's argument to establish uniform convergence layer by layer. Then, since there are only finitely many layers, we still get the desired conclusion that for any $R > 0$, $|N(x) - \mathbb{E}[t_\Delta(x)]| = O(\Delta)$ for all x with $\|x\| \le R$. In his Lemma 7, Richardson shows that

$$\mathrm{Var}[t_\Delta(x)] \le \mathrm{Var}\big[t_\Delta(\tfrac{1}{2}x) + s_\Delta(\tfrac{1}{2}x)\big] + \varepsilon_6(\Delta, x),$$

where $s_\Delta(\frac{1}{2}x)$ is an independent copy of $t_\Delta(\frac{1}{2}x)$. The proof uses that $x = \frac{1}{2}x + \frac{1}{2}x$, which is not true in general on $\mathbb{R}^2 \times \mathbb{Z}_w$. It is true, however, up to addition by a vector parallel to e_3, whose scaled infection time has negligible variance by the above. In his Lemma 10, Richardson uses the fact that if a site x/Δ is first infected at some time ≤ 1 in the scaled spacetime, there must be a path from the origin to x/Δ on which first infection times are ≤ 1, and along which $N(y)$ varies continuously. Since $N(y + e_3) = N(y)$, adding jumps across layers will not disrupt this continuity property.

Note that (A.44) is a weaker statement than the Bramson-Griffeath shape theorem, both because it does not hold almost surely, and because $\{x : t(x) \le t\}$ describes the set of points that have been infected at some point before time t, whereas the shape theorem (2.4) makes a claim on the state of the process at time t. This is the reason

that Bramson and Griffeath use slightly stronger conditions than Richardson's original conditions to prove their result.

Final result on $\mathbb{Z}^2 \times \mathbb{Z}_w$

As previously mentioned, Bramson and Griffeath's proof that (A1)-(A5) hold for the biased voter model on $\mathbb{Z}^2$, conditioned on nonextinction, will carry over to $\mathbb{Z}^2 \times \mathbb{Z}_w$ with minor modifications. To state the final result on $\mathbb{Z}^2 \times \mathbb{Z}_w$, let N be defined by (A.43) and set $D'_R := \{x \in \mathbb{R}^2 \times \mathbb{Z}_w : N(x) \leq R\}$. Then, for all $\varepsilon > 0$,

$$\mathbb{P}\big(\exists\, t_* < \infty : D'_{(1-\varepsilon)t} \cap (\mathbb{Z}^2 \times \mathbb{Z}_w) \subseteq \xi_t^0 \subseteq D'_{(1+\varepsilon)t}\ \forall t \geq t_* \mid \tau_\varnothing^0 = \infty\big) = 1,$$

which is the stronger version of (A.44). In the statement (2.4) of the shape theorem in the main text, the set D is the unit ball on $\mathbb{R}^2 \times \mathbb{Z}_w$ under N, i.e. $D := D'_1$. Define

$$X := \{(x_1, x_2) \in \mathbb{R}^2 : (x_1, x_2, 0) \in D\} = \{(x_1, x_2) \in \mathbb{R}^2 : N((x_1, x_2, 0)) \leq 1\}.$$

Since $N(x + e_3) = N(x)$ for all $x \in \mathbb{R}^2 \times \mathbb{Z}_w$, we must have $D = \cup_{i \in \mathbb{Z}_w}(X \times \{i\})$. Since N satisfies the triangle inequality and $N(tx) = |t|N(x)$ for $t \in \mathbb{R}$ and $x \in \mathbb{R}^2 \times \mathbb{Z}_w$, the set X is convex on $\mathbb{R}^2$. Finally, X inherits all symmetries of the biased voter model.

A.2 Boundary condition comparison

Here, we use simulation to compare the propagation speed of the biased voter model on w layers of two-dimensional integer lattices for two different boundary conditions along the third dimension. On the one hand, we consider $\mathbb{Z}^2 \times \mathbb{Z}_w$ with a periodic boundary condition, and on the other hand, we consider $\mathbb{Z}^2 \times [\![0, w-1]\!]$ with a reflecting boundary condition, i.e. cells on the top (resp. bottom) layer can only replace cells on the same layer and the layer immediately below (resp. above). In Figure A.1, we show results of simulations of these two processes given tissue thickness $w = 2, 3, 4, 5$ and fitness advantage $\beta = 0.01, 0.05, 0.1$. We ran at least 30 simulations for each set of parameters and stopped each simulation when the process reached $(100, 0, 0)$ or $(-100, 0, 0)$. We then used this data to determine an average speed and a 95% confidence interval for each set of parameters.

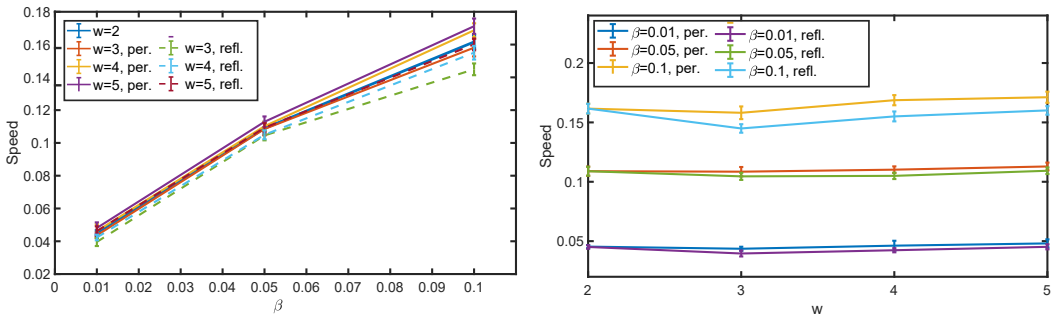

Figure A.1: Simulation comparison of the propagation speed on $\mathbb{Z}^2 \times \mathbb{Z}_w$, with the third dimension equipped with a periodic boundary condition, and on $\mathbb{Z}^2 \times [\![0, w-1]\!]$, with the third dimension equipped with a reflecting boundary condition. In **(a)**, we show the propagation speed as a function of β for $w = 2$ to $w = 5$. In **(b)**, we show the propagation speed as a function of w for $\beta = 0.01, 0.05, 0.1$. Error bars indicate 95% confidence intervals.

Note first that the two boundary conditions are equivalent for $w = 2$ two layers. When $w > 2$, equipping the model with a reflecting boundary condition along the third dimension will result in a smaller propagation speed than for the periodic case, due to the decreased ability of type-1 cells on the top and bottom layers to spread out. However, the difference is small, especially for smaller values of the fitness advantage β, which indicates that our modeling decision to equip the third layer with a periodic boundary condition is a reasonable approximation for small β (recall that in [14], $\beta = 0.004$ is estimated to be a typical value).

Appendix B

Appendix to Chapter 3

B.1 Proof of Proposition 3.1

In this section, we prove Proposition 3.1 on the expected SFS of the skeleton. We first present a brief outline of the proof.

In part (1), the SFS is observed at the fixed time $\tilde{t}_N$. To compute the expected number of mutations that end up in $j \geq 1$ skeleton cells at time $\tilde{t}_N$, we decompose the time interval $[0, \tilde{t}_N]$ into infinitesimal intervals $[t, t + \Delta t]$. We can compute how many mutations occur during each small interval using the expected mutation rate (3.6) in the main text. Then, to only count the mutations that end up in j skeleton cells at time $\tilde{t}_N$, we multiply this number by the probability that a single-cell derived skeleton clone has size j at time $\tilde{t}_N$, which has a known expression. We finally integrate over time to add up the contributions of the infinitesimal intervals.

In part (2), the SFS is observed at the stochastic time $\tilde{\tau}_N$, at which the skeleton reaches size N. In this case, we decompose into population size levels instead of into small time intervals. We know how many mutations accumulate on population size level k by (3.9) in the main text. Then, to get the expected number of mutations that accumulate on level k and end up in $j \geq 1$ skeleton cells at time $\tilde{\tau}_N$, we need to compute the probability that starting from one skeleton cell carrying a particular mutation and $k - 1$ cells without it, j cells carry the mutation when the skeleton reaches size N. To that end, we define a Markov chain that keeps track of how many skeleton cells carry the mutation as the skeleton increases in size, and we compute its hitting probabilities.

We finally sum over k to add up the contributions of each population size level.

Proof of Proposition 3.1. (1) By (3.6) in the main text, the expected mutation rate per skeleton cell per unit time is wr_0. To stratify mutations based on their frequencies at time $\tilde{t}_N$, we define

$$\tilde{p}_j(s) := \mathbb{P}(\tilde{Z}_0(s) = j | \tilde{Z}_0(0) = 1), \quad j \geq 1, s \geq 0,$$

as the size-distribution at time s of a single-cell derived skeleton clone. Since $(\tilde{Z}_0(t))_{t\geq 0}$ is a Yule process with birth rate λ_0, this distribution has an explicit expression,

$$\tilde{p}_j(s) = (1/e^{\lambda_0 s})(1 - 1/e^{\lambda_0 s})^{j-1}, \quad j \geq 1,$$

which is the geometric distribution with support $\{1, 2, \ldots\}$ and success probability $1/e^{\lambda_0 s}$, see e.g. Section 3 of Durrett [50] (the support does not include 0 since skeleton clones do not go extinct). For $0 \leq t \leq \tilde{t}_N$, let $\tilde{S}_{j,\tilde{t}_N}(t)$ denote the number of mutations that accumulate in the time interval $[0, t]$ and are found in $j \geq 1$ skeleton cells at time $\tilde{t}_N$. We write $\tilde{S}_j(\tilde{t}_N) := \tilde{S}_{j,\tilde{t}_N}(\tilde{t}_N)$ for the site frequency spectrum of the skeleton at time $\tilde{t}_N$. If a mutation occurs during an infinitesimal time interval $[t, t + \Delta t]$, the clone started by the cell carrying the mutation has size j at time $\tilde{t}_N$ with probability $\tilde{p}_j(\tilde{t}_N - t) + O(\Delta t)$, where $f(x) = O(x)$ means that there exists $C > 0$ so that $|f(x)| \leq Cx$ for sufficiently small $x > 0$. The expected number of mutations that accumulate in $[t, t + \Delta t]$ and are present in $j \geq 1$ skeleton cells at time $\tilde{t}_N$ is therefore

$$\mathbb{E}[\tilde{S}_{j,\tilde{t}_N}(t + \Delta t)] - \mathbb{E}[\tilde{S}_{j,\tilde{t}_N}(t)] = wr_0 \Delta t \cdot e^{\lambda_0 t} \cdot \tilde{p}_j(\tilde{t}_N - t) + o(\Delta t),$$

where we use that $\mathbb{E}[\tilde{Z}_0(t)] = e^{\lambda_0 t}$ is the mean skeleton size at time t, and $f(x) = o(x)$ means that $f(x)/x \to 0$ as $x \to 0$. This calculation is somewhat heuristic in that we have simply multiplied an expected mutation rate by an expected population size, which is in turn multiplied by the probability that a particular mutation ends up in j cells at time $\tilde{t}_N$. In our proof of part (1) of Proposition 3.2 for the total population, we present a more detailed argument which can be used to obtain this expression more rigorously.

Integrating over time, and using that $q_0 = \lambda_0/r_0$ by expression (E.3) and $N = e^{\lambda_0 \tilde{t}_N}$ by expression (3.10), we obtain

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] &= \int_0^{\tilde{t}_N} w r_0 \tilde{p}_j(\tilde{t}_N - t) e^{\lambda_0 t} dt \\ &= (w/q_0)N \cdot \int_0^{\tilde{t}_N} (e^{\lambda_0 t}/N)(1 - e^{\lambda_0 t}/N)^{j-1} \cdot \lambda_0 (e^{\lambda_0 t}/N) dt.\end{aligned}$$

Substituting $y := 1 - e^{\lambda_0 t}/N$, $dy = -\lambda_0(e^{\lambda_0 t}/N)dt$, this implies

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] &= (w/q_0)N \cdot \int_0^{1-1/N} (1-y) y^{j-1} dy \\ &= (w/q_0)N \cdot \big(1 - \tfrac{1}{N}\big)^j \big(\tfrac{1}{j(j+1)} + \tfrac{1}{N}\tfrac{1}{j+1}\big),\end{aligned}$$

the desired result. Clearly, for fixed $j \geq 1$, then as $N \to \infty$,

$$\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] \sim (w/q_0)N \cdot 1/(j(j+1)). \tag{B.1}$$

The asymptotic expression (B.1) can also be derived more heuristically as follows, which gives another way of interpreting the expression. If the skeleton is observed at a large time t, the age s of an arbitrary mutation has approximate density $\lambda_0 e^{-\lambda_0 s}$. A mutation with age s at time t is found in $j \geq 1$ skeleton cells at time t with probability $\tilde{p}_j(s)$. The probability that an arbitrary mutation is found in $j \geq 1$ skeleton cells at time t is therefore, as $t \to \infty$,

$$\int_0^\infty \tilde{p}_j(s) \cdot \lambda_0 e^{-\lambda_0 s} ds = \int_0^1 (1-y) y^{j-1} dy = 1/(j(j+1)),$$

using the substitution $y := 1 - e^{-\lambda_0 s}$, $dy = \lambda_0 e^{-\lambda_0 s} ds$. Next, we can compute the expected total number of mutations up until time $\tilde{t}_N$ via

$$\int_0^{\tilde{t}_N} w r_0 e^{\lambda_0 s} ds = (w/q_0)(N-1) \sim (w/q_0)N,$$

which is given as (3.16) in the main text. Finally, we can obtain (B.1) as the expected total number of mutations up until time $\tilde{t}_N$ multiplied by the probability $1/(j(j+1))$ of finding an arbitrary mutation in j skeleton cells at time $\tilde{t}_N$ as $N \to \infty$. The distribution $j \mapsto 1/(j(j+1))$ is a special case of the *Yule-Simon distribution*, which was originally computed by Yule [169] as the distribution of the number of

species within a genus, where a species mutates to a new species within the same genus at some rate s, and a genus mutates to a new genus at some rate g. In the previous paragraph, we have adapted Yule's basic argument to our setting with $g = s = \lambda_0$. We refer to Simkin and Roychowdhury [141] for a comprehensive discussion of how the Yule-Simon distribution and variants thereof have appeared in a wide variety of scientific contexts since its original conception.

(2) In (3.8) of Section 3.3.1, we showed that on average, wp_0/q_0 mutations accumulate on type-1 divisions in between two type-2 divisions, while type-2 divisions add w mutations and change the skeleton population size level. Since the type-1 mutations on level $k = 1$ are the clonal mutations, the expected number of clonal mutations is wp_0/q_0, which is the $j = N$ case of the desired result. For $k = 2, \ldots, N-1$, the expected number of mutations on level k is $w + wp_0/q_0 = w/q_0$, which includes the type-2 division that starts the level.

For $1 \leq j \leq N-1$, let $\tilde{h}^j_{(1,k-1)}$ be the probability that starting with one skeleton cell carrying a particular mutation and $k-1$ cells without it, j cells carry the mutation when the skeleton reaches size N. Since for levels $k = 2, \ldots, N-1$, there are w/q_0 mutations on average per level, and each mutation on level k contributes $\tilde{h}^j_{(1,k-1)}$ to the expected number of mutations found in j skeleton cells at level N, we obtain for $1 \leq j \leq N-1$,

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] &= (w/q_0) \textstyle\sum_{k=2}^{N-1} \tilde{h}^j_{(1,k-1)} + w\delta_{1,j} \\ &= (w/q_0) \textstyle\sum_{k=1}^{N-2} \tilde{h}^j_{(1,k)} + w\delta_{1,j},\end{aligned}$$

where the extra $w\delta_{1,j}$ term is due to mutations that occur on the final type-2 division that changes levels from $N-1$ to N, each of which is found in one skeleton cell. As was the case for the proof of part (1), the fact that for each level k, we can simply multiply the expected number of mutations with the probability that each particular mutation ends up in j cells when the skeleton reaches size N can be justified more rigorously using an argument similar to the one we present in part (2) of Proposition 3.2 in Appendix B.2.

It remains to compute the probabilities $\tilde{h}^j_{(1,k)}$. To this end, define a two-dimensional

discrete-time Markov chain on the state space $\{(\ell, m) : \ell, m \geq 1, \ell + m \leq N\}$, where ℓ is the number of skeleton cells carrying a particular mutation and m is the number of cells without it. Since each skeleton cell, with or without the mutation, divides into two cells at rate λ_0, the transition probabilities for this chain are given by

$$\begin{aligned}(\ell, m) &\to (\ell + 1, m) \quad \text{w.p.} \quad \ell/(\ell + m),\\ (\ell, m) &\to (\ell, m + 1) \quad \text{w.p.} \quad m/(\ell + m),\end{aligned}$$

for $\ell, m \geq 1$ and $\ell + m < N$. The states $(\ell, N - \ell)$ for $1 \leq \ell \leq N - 1$ are absorbing. A diagram for this Markov chain is shown in Figure B.1.

Let $\tilde{h}^r_{(\ell,m)}$ denote the probability that the above chain is absorbed in state $(r, N-r)$ when started from state (ℓ, m). It is immediate that $\tilde{h}^r_{(\ell,m)} = 0$ if $\ell > r$ or $m > N-r$. For (ℓ, m) with $\ell \leq r$, $m \leq N - r$ and $\ell + m < N$, by conditioning on whether the first transition out of state (ℓ, m) is to $(\ell + 1, m)$ or $(\ell, m + 1)$, we obtain the following recursion for $\tilde{h}^r_{(\ell,m)}$:

$$(\ell + m)\tilde{h}^r_{(\ell,m)} = \ell\tilde{h}^r_{(\ell+1,m)} + m\tilde{h}^r_{(\ell,m+1)}. \tag{B.2}$$

The boundary conditions are $\tilde{h}^r_{(\ell,N-\ell)} = \delta_{\ell,r}$ for $1 \leq \ell \leq N - 1$. It is actually possible to compute $\tilde{h}^r_{(\ell,m)}$ directly as the sum of probabilities of all possible paths from (ℓ, m) to $(r, N - r)$ without using the above recursion. By noting that there are $\binom{N-(\ell+m)}{r-\ell}$ possible paths, and that each path has the same probability, we can obtain

$$\tilde{h}^r_{(\ell,m)} = \binom{N-(\ell+m)}{r-\ell} \cdot \prod_{n=0}^{r-\ell-1} \frac{\ell+n}{\ell+m+n} \cdot \prod_{n=0}^{N-m-r-1} \frac{m+n}{r+m+n}, \tag{B.3}$$

with $\prod_{\varnothing} := 1$. As verification, it is straightforward to check that (B.3) solves (B.2). To obtain $\tilde{h}^j_{(1,k)}$ for $1 \leq k \leq N - 2$ and $1 \leq j \leq N - 1$, note first that $\tilde{h}^j_{(1,k)} = 0$ for

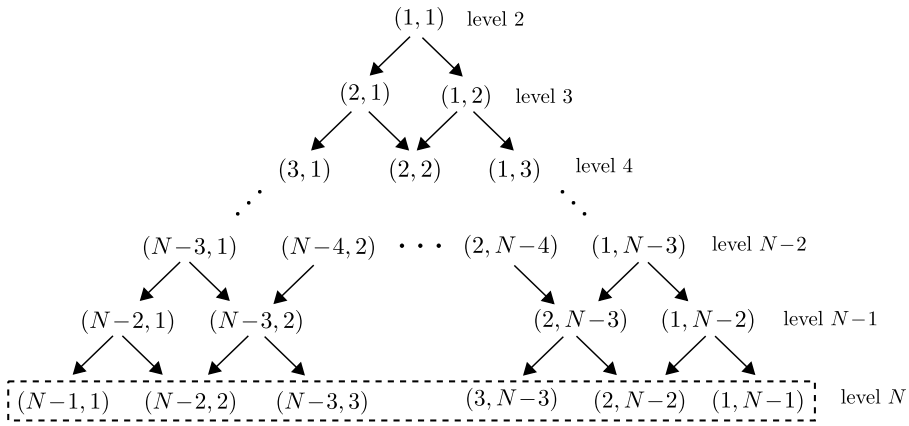

Figure B.1: A diagram of the discrete-time Markov chain on $\{(\ell, m) : \ell, m \geq 1, \ell + m \leq N\}$, where ℓ is the number of skeleton cells carrying a particular mutation and m is the number of cells without it. Each type-2 division increases the skeleton population size level by one, and since skeleton cells do not die, the chain never returns to the lower levels. The states $(\ell, N-\ell)$ with $1 \leq \ell \leq N-1$ are absorbing (dashed box), since we are only interested in the evolution up until level N.

$k > N - j$. For $1 \leq k \leq \min(N-j, N-2)$, we can simplify (B.3) as follows:

$$\begin{aligned}
\tilde{h}^j_{(1,k)} &= \tbinom{N-(k+1)}{j-1} \cdot \textstyle\prod_{n=0}^{j-2} \frac{1+n}{k+1+n} \cdot \prod_{n=0}^{N-k-j-1} \frac{k+n}{k+j+n} \\
&= \tfrac{(N-k-1)!}{(j-1)!(N-k-j)!} \cdot \tfrac{(j-1)!}{(k+1)\cdots(k+j-1)} \cdot \tfrac{k\cdots(N-j-1)}{(k+j)\cdots(N-1)} \\
&= k \cdot \tfrac{(N-k-1)!}{(N-k-j)!} \cdot \tfrac{(N-j-1)!}{(N-1)!} \\
&= \tbinom{N-j-1}{k-1}\tbinom{N-1}{k}^{-1}.
\end{aligned}$$

Thus, for $1 \leq j \leq N-1$,

$$\begin{aligned}
\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] &= (w/q_0) \textstyle\sum_{k=1}^{N-2} \tilde{h}^j_{(1,k)} + w\delta_{1,j} \\
&= (w/q_0) \textstyle\sum_{k=1}^{\min(N-j,N-2)} \tbinom{N-j-1}{k-1}\tbinom{N-1}{k}^{-1} + w\delta_{1,j}. \qquad \text{(B.4)}
\end{aligned}$$

For $j = 1$, it is easy to compute

$$\begin{aligned}\mathbb{E}[\tilde{S}_1(\tilde{\tau}_N)] &= (w/q_0)(1/(N-1))\big(\textstyle\sum_{k=1}^{N-2} k\big) + w \\ &= (1/2)(w/q_0)(N-2) + w \\ &= (1/2)(w/q_0)N - wp_0/q_0.\end{aligned}$$

To simplify (B.4) for $2 \le j \le N-1$, we first note that for any positive integers a and b,

$$\textstyle\int_0^1 t^{a-1}(1-t)^{b-1}dt = \frac{(a-1)!(b-1)!}{(a+b-1)!},$$

see e.g. Theorem 1.1.4 of Andrews et al. [2]. By observing that

$$N\textstyle\int_0^1 t^k(1-t)^{N-1-k}dt = \binom{N-1}{k}^{-1},$$

we can rewrite (B.4) as follows:

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] &= (w/q_0)N \cdot \textstyle\int_0^1 \big(\sum_{k=1}^{N-j} \binom{N-j-1}{k-1} t^k(1-t)^{N-1-k}\big)dt \\ &= (w/q_0)N \cdot \textstyle\int_0^1 t(1-t)^{j-1}\big(\sum_{k=1}^{N-j} \binom{N-j-1}{k-1} t^{k-1}(1-t)^{(N-j-1)-(k-1)}\big)dt.\end{aligned}$$

The sum inside the integral is the total probability mass ($= 1$) of the binomial distribution with number of trials $N-j-1$ and success probability t. Therefore, for $2 \le j \le N-1$,

$$\begin{aligned}\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] &= (w/q_0)N \cdot \textstyle\int_0^1 t(1-t)^{j-1}dt \\ &= (w/q_0)N \cdot 1/(j(j+1)).\end{aligned}$$

This concludes the proof. □

B.2 Proof of Proposition 3.2

In this section, we prove Proposition 3.2 on the expected SFS of the total population. The proof strategy is the same as in the proof of Proposition 3.1 for the skeleton. There

are some added complications, however, since the tumor as a whole follows a birth-death process, whereas the skeleton subpopulation follows a pure-birth process.

In part (1), we wish to compute the expected number of mutations that accumulate in an infinitesimal time interval and end up in $j \geq 1$ tumor cells at the fixed time t_N. In this case, we need to condition on survival up until time t_N, which was not necessary in the same computation for the skeleton. In the proof below, we first give an informal argument for how to handle the conditioning on survival, and then present detailed calculations at the end of the proof. Another complication is that the size distribution of a single-cell derived clone now has a more complex form than for the skeleton, which translates into more work simplifying the integral that results from adding up the infinitesimal interval contributions.

In part (2), we decompose into population size levels. To compute how many mutations accumulate on population size level k, we first compute how many times the population hits this level. This computation was not necessary for the skeleton, since the skeleton only increases in size. We then need to compute the probability that starting from one cell carrying a particular mutation and $k-1$ cells without it, j cells carry the mutation when the population reaches size N. This probability can be computed as a hitting probability of a Markov chain, as in the proof of Proposition 3.1, but the Markov chain now has a more complicated structure. This time, we are not able to compute the hitting probabilities explicitly, and we instead provide a linear system which determines them.

Proof of Proposition 3.2. (1) We begin by defining

$$p_j(s) := \mathbb{P}(Z_0(s) = j | Z_0(0) = 1), \quad j \geq 0, s \geq 0,$$

as the size-distribution at time s of a single-cell derived clone. This distribution has an explicit expression: Setting

$$g(t) := \tfrac{p_0(e^{\lambda_0 t}-1)}{e^{\lambda_0 t}-p_0} \quad \text{and} \quad h(t) := \tfrac{e^{\lambda_0 t}-1}{e^{\lambda_0 t}-p_0},$$

we can write

$$p_0(t) = \mathbb{P}(Z_0(t) = 0 | Z_0(0) = 1) = g(t),$$

and for $j \geq 1$,

$$\begin{aligned} p_j(t) &= \mathbb{P}(Z_0(t) = j | Z_0(0) = 1) \\ &= (1 - g(t))(1 - h(t))(h(t))^{j-1}, \end{aligned}$$

see e.g. (8) of Durrett [50]. Simplifying, we obtain

$$\begin{aligned} p_0(t) &= \tfrac{p_0(e^{\lambda_0 t}-1)}{e^{\lambda_0 t}-p_0}, \\ p_j(t) &= \tfrac{q_0^2 e^{\lambda_0 t}}{(e^{\lambda_0 t}-p_0)^2} \cdot \left(\tfrac{e^{\lambda_0 t}-1}{e^{\lambda_0 t}-p_0}\right)^{j-1}, \quad j \geq 1. \end{aligned} \tag{B.5}$$

The probability that a single-cell derived clone is still alive at time t is then given by

$$\mathbb{P}(Z_0(t) > 0 | Z_0(0) = 1) = 1 - p_0(t) = q_0 e^{\lambda_0 t}/(e^{\lambda_0 t} - p_0). \tag{B.6}$$

For $0 \leq t \leq t_N$, let $S_{j,t_N}(t)$ denote the number of mutations that accumulate in $[0, t]$ and are found in $j \geq 1$ cells at time t_N. We write $S_j(t_N) := S_{j,t_N}(t_N)$ for the site frequency spectrum at time t_N. Say a cell division occurs in an infinitesimal time interval $[t, t + \Delta t]$. The division results in w mutations on average, each assigned to one of the two daughter cells, and the clone started by this cell has size $j \geq 1$ cells at time t_N with probability $p_j(t_N - t) + O(\Delta t)$. We wish to show that on the event $\{Z_0(t_N) > 0\}$ of survival of the population up until time t_N, the expected number of mutations that accumulate in $[t, t + \Delta t]$ and are found in $j \geq 1$ cells at time t_N is

$$\mathbb{E}\big[\big(S_{j,t_N}(t + \Delta t) - S_{j,t_N}(t)\big) 1_{\{Z_0(t_N)>0\}}\big] = w r_0 \Delta t \cdot e^{\lambda_0 t} \cdot p_j(t_N - t) + o(\Delta t), \tag{B.7}$$

where we use that $\mathbb{E}[Z_0(t)] = e^{\lambda_0 t}$. It will then follow from (B.7) that

$$\begin{aligned}&\mathbb{E}\big[S_{j,t_N}(t+\Delta t) - S_{j,t_N}(t) \,\big|\, Z_0(t_N) > 0\big] \\ &= wr_0 e^{\lambda_0 t}\Delta t \cdot p_j(t_N - t)/(1 - p_0(t_N)) + o(\Delta t). \end{aligned} \tag{B.8}$$

Note that (B.7) clearly holds for the semideterministic model in which the tumor bulk grows deterministically at rate λ_0, mutant clones arise at stochastic rate wr_0, and mutant clones grow stochastically. It is not obvious that (B.7) also holds for our fully stochastic model, since including the event $\{Z_0(t_N) > 0\}$ of survival up until time t_N should presumably affect the expected population size at time $t \leq t_N$. The key is to observe that if a mutation occurs on a cell division at time $t \leq t_N$ and ends up in $j \geq 1$ cells at time t_N, the population is automatically alive at time t_N. The relevant survival event in (B.7) is therefore $\{Z_0(t) > 0\}$, and the relevant population size factor is $\mathbb{E}[Z_0(t)1_{\{Z_0(t)>0\}}] = \mathbb{E}[Z_0(t)] = e^{\lambda_0 t}$. Another potential concern in establishing (B.7) for our model is that we allow multiple mutations to occur per cell division. To not distract further from the main calculations, we assume that the reader is willing to accept (B.7) as true for the moment, and we provide a detailed mathematical argument for this expression at the end of the proof.

Using (B.8), we can integrate over time to obtain

$$\mathbb{E}\big[S_j(t_N)|Z_0(t_N) > 0\big] = (1 - p_0(t_N))^{-1} \cdot \int_0^{t_N} wr_0 e^{\lambda_0 t} p_j(t_N - t)dt. \tag{B.9}$$

Focusing on the integral, we write

$$\begin{aligned}&\int_0^{t_N} wr_0 e^{\lambda_0 t} p_j(t_N - t)dt \\ &= (w/q_0) \cdot \int_0^{t_N} \tfrac{q_0^2 e^{\lambda_0 t_N} e^{\lambda_0 t}}{(e^{\lambda_0 t_N} - p_0 e^{\lambda_0 t})^2} \cdot \big(\tfrac{e^{\lambda_0 t_N} - e^{\lambda_0 t}}{e^{\lambda_0 t_N} - p_0 e^{\lambda_0 t}}\big)^{j-1} \cdot \lambda_0 e^{\lambda_0 t} dt \\ &= wq_0 e^{\lambda_0 t_N} \cdot \int_0^{t_N} \tfrac{e^{\lambda_0 t}}{(e^{\lambda_0 t_N} - p_0 e^{\lambda_0 t})^2} \cdot \big(\tfrac{e^{\lambda_0 t_N} - e^{\lambda_0 t}}{e^{\lambda_0 t_N} - p_0 e^{\lambda_0 t}}\big)^{j-1} \cdot \lambda_0 e^{\lambda_0 t} dt.\end{aligned}$$

Set $L := e^{\lambda_0 t_N}$. Using the substitution $x := e^{\lambda_0 t}$, $dx = \lambda_0 e^{\lambda_0 t}$, we obtain

$$\int_0^{t_N} wr_0 e^{\lambda_0 t} p_j(t_N - t)dt = wq_0 L \cdot \int_1^L \tfrac{x}{(L - p_0 x)^2} \cdot \big(\tfrac{L - x}{L - p_0 x}\big)^{j-1} dx.$$

We again change variables, this time $y := (L-x)/(L-p_0x)$, in which case

$$\begin{aligned} x &= L(1-y)/(1-p_0y), \\ dx &= -\big(q_0L/(1-p_0y)^2\big)dy, \\ L-p_0x &= q_0L/(1-p_0y), \end{aligned}$$

and $y = (L-1)/(L-p_0) = 1 - q_0/(L-p_0)$ for $x = 1$ and $y = 0$ for $x = L$, which implies

$$\int_0^{t_N} wr_0e^{\lambda_0 t}p_j(t_N-t)dt = wL \cdot \int_0^{1-q_0/(L-p_0)}(1-p_0y)^{-1}(1-y)y^{j-1}dy. \qquad \text{(B.10)}$$

We now apply (B.9) and (B.6) to see that

$$\mathbb{E}\big[S_j(t_N)|Z_0(t_N)>0\big] = w \cdot \tfrac{e^{\lambda_0 t_N}-p_0}{q_0} \cdot \int_0^{1-q_0/(e^{\lambda_0 t_N}-p_0)}(1-p_0y)^{-1}(1-y)y^{j-1}dy,$$

and the desired result (3.20) follows from the fact that $(e^{\lambda_0 t_N}-p_0)/q_0 = N$ by the definition of t_N in (3.18). It also follows that for fixed $j \geq 1$,

$$\mathbb{E}\big[S_j(t_N)|Z_0(t_N)>0\big] \sim wN \cdot \int_0^1(1-p_0y)^{-1}(1-y)y^{j-1}dy$$

as $N \to \infty$. To write the last expression as a sum, note that

$$(1-p_0y)^{-1} = \textstyle\sum_{k=0}^{\infty} p_0^k y^k,$$

which is valid for all $0 \leq p_0 < 1$ and $0 \leq y \leq 1$. It follows that

$$\begin{aligned} \int_0^1(1-p_0y)^{-1}(1-y)y^{j-1}dy &= \textstyle\sum_{k=0}^{\infty} p_0^k\big(\int_0^1(1-y)y^{j+k-1}dy\big) \\ &= \textstyle\sum_{k=0}^{\infty} \frac{p_0^k}{(j+k)(j+k+1)}. \end{aligned}$$

We conclude by establishing (B.7) above, which was

$$\mathbb{E}\big[(S_{j,t_N}(t+\Delta t) - S_{j,t_N}(t))1_{\{Z_0(t_N)>0\}}\big] = wr_0\Delta t \cdot e^{\lambda_0 t} \cdot p_j(t_N-t) + o(\Delta t).$$

We decompose according to population size at time t. Assume that $Z_0(t) = k$ with

$k \geq 1$, i.e. there are k cells at time t. Let $D_{t,\Delta t}$ denote the event that exactly one of the k cells divides in the infinitesimal time interval $[t, t+\Delta t]$, and enumerate the $k+1$ cells after the cell division as $Y_t^1, \ldots, Y_t^{k+1}$, where Y_t^1 and Y_t^2 are the two new cells. Let W denote the number of mutations that occur on the cell division, where W is a nonnegative integer-valued random variable with $\mathbb{E}[W] = w$, independent of $(Z_0(t))_{t\geq 0}$. For $\ell \geq 1$, let B_ℓ be i.i.d. with $\mathbb{P}(B_\ell = 1) = \mathbb{P}(B_\ell = 2) = 1/2$, independent of $(Z_0(t))_{t\geq 0}$ and W, and assign mutation number ℓ to cell number B_ℓ for $1 \leq \ell \leq W$. Finally, let $Y_t^m(s)$ be the number of descendants of cell Y_t^m at time $t+s$, with $Y_t^m(0) = 1$. With this notation, define

$$A_{j,k,\ell}(t) := \{Z_0(t) = k\} \cap D_{t,\Delta t} \cap \{\ell \leq W\} \cap \{Y_t^{B_\ell}(t_N - t) = j\} \cap \{Z_0(t_N) > 0\}.$$

This is the event that the tumor survives to time t_N, that it consists of k cells at time t, that exactly one of the k cells divides in $[t, t+\Delta t]$, that at least ℓ mutations occur on this division, and that mutation number ℓ is found in j cells at time t_N. The reason we are interested in this event is that we can write

$$\mathbb{E}\big[(S_{j,t_N}\big(t+\Delta t) - S_{j,t_N}(t)\big)1_{\{Z_0(t_N)>0\}}\big] = \textstyle\sum_{k=1}^{\infty}\sum_{\ell=1}^{\infty} \mathbb{P}\big(A_{j,k,\ell}(t)\big) + o(\Delta t),$$

where the $o(\Delta t)$ term captures the possibility of more than one cell division in $[t, t+\Delta t]$.

To compute $\mathbb{P}\big(A_{j,k,\ell}(t)\big)$, note first that

$$\mathbb{P}\big(A_{j,k,\ell}(t)\big) = \mathbb{P}\big(\{Z_0(t) = k\} \cap D_{t,\Delta t} \cap \{\ell \leq W\} \cap \{Y_t^{B_\ell}(t_N - t) = j\}\big),$$

since the survival event $\{Z_0(t_N) > 0\}$ is implied by the other events. By independence,

$$\mathbb{P}\big(A_{j,k,\ell}(t)\big) = \mathbb{P}(\ell \leq W) \cdot \mathbb{P}\big(\{Z_0(t) = k\} \cap D_{t,\Delta t} \cap \{Y_t^{B_\ell}(t_N - t) = j\}\big).$$

To analyze the latter probability, note that since B_ℓ is independent of $(Z_0(t))_{t\geq 0}$,

and $\big((Z_0(s))_{s\le t}, (Y_t^1(s))_{s\ge 0}\big) \overset{d}{=} \big((Z_0(s))_{s\le t}, (Y_t^2(s))_{s\ge 0}\big)$, we can write

$$\begin{aligned}&\mathbb{P}\big(\{Z_0(t)=k\}\cap D_{t,\Delta t}\cap\{Y_t^{B_\ell}(t_N-t)=j\}\big)\\&=\mathbb{P}\big(\{Z_0(t)=k\}\cap D_{t,\Delta t}\cap\{Y_t^{1}(t_N-t)=j\}\big).\end{aligned}$$

Using the Markov property, we can calculate the latter probability as

$$\begin{aligned}&\mathbb{P}\big(\{Z_0(t)=k\}\cap D_{t,\Delta t}\cap\{Y_t^{1}(t_N-t)=j\}\big)\\&=\mathbb{P}(Z_0(t)=k)\cdot\mathbb{P}(D_{t,\Delta t}|Z_0(t)=k)\cdot\mathbb{P}\big(Y_t^1(t_N-t)=j\big|Z_0(t)=k,D_{t,\Delta t}\big)\\&=\mathbb{P}(Z_0(t)=k)\cdot e^{-kr_0\Delta t}kr_0\Delta t\cdot(p_j(t_N-t)+O(\Delta t)).\end{aligned}$$

Combining the above, we obtain

$$\begin{aligned}&\mathbb{E}\big[\big(S_{j,t_N}(t+\Delta t)-S_{j,t_N}(t)\big)1_{\{Z_0(t_N)>0\}}\big]\\&=\textstyle\sum_{k=1}^\infty\sum_{\ell=1}^\infty\mathbb{P}\big(A_{j,k,\ell}(t)\big)+o(\Delta t)\\&=r_0p_j(t_N-t)\Delta t\cdot\big(\textstyle\sum_{\ell=1}^\infty\mathbb{P}(W\ge\ell)\big)\cdot\big(\textstyle\sum_{k=1}^\infty k\mathbb{P}(Z_0(t)=k)\big)+o(\Delta t)\\&=wr_0\Delta t\cdot e^{\lambda_0 t}\cdot p_j(t_N-t)+o(\Delta t),\end{aligned}$$

where we use $\sum_{\ell=1}^\infty\mathbb{P}(W\ge\ell)=\mathbb{E}[W]=w$ and $\sum_{k=1}^\infty k\mathbb{P}(Z_0(t)=k)=\mathbb{E}[Z_0(t)]=e^{\lambda_0 t}$. This concludes the proof.

(2) Let $(X_n)_{n\ge0}$ denote the discrete-time jump process embedded in $(Z_0(t))_{t\ge0}$ that only keeps track of changes in population size. More precisely, if σ_n is the time of the n-th jump of $(Z_0(t))_{t\ge0}$ for $n\ge1$, then $X_0=1$ and $X_n=Z_0(\sigma_n)$ for $n\ge1$. Since cells divide at rate r_0 and die at rate d_0, $(X_n)_{n\ge0}$ is a simple random walk, absorbed at 0, which moves up with probability $a:=r_0/(r_0+d_0)=1/(1+p_0)$ and down with probability $b=1-a=p_0/(1+p_0)$. Since we are only interested in what happens until the population either goes extinct or reaches level N, we treat N as an absorbing state. Define

$$T_k:=\inf\{n\ge0:X_n=k\},\quad 0\le k\le N, \tag{B.11}$$

as the (discrete) time at which the random walk first hits level k, with $\inf\varnothing=\infty$.

Let $\mathbb{P}_j$ denote the probability measure of $(X_n)_{n\geq 0}$ when started at $X_0 = j$. By the gambler's ruin formula,

$$\mathbb{P}_j(T_k < T_0) = (1 - p_0^j)/(1 - p_0^k), \quad 0 \leq j \leq k. \tag{B.12}$$

For $1 \leq k \leq N-1$, let $\Lambda_{k,k+1}$ denote the number of transitions from k to $k+1$,

$$\Lambda_{k,k+1} := \textstyle\sum_{j=0}^{\infty} 1_{\{Z_j = k, Z_{j+1} = k+1\}},$$

and let Λ_k denote the number of visits to k,

$$\Lambda_k := \textstyle\sum_{j=0}^{\infty} 1_{\{Z_j = k\}}.$$

By the strong Markov property and (B.12), we can write

$$\mathbb{E}_1[\Lambda_k] = \mathbb{P}_1(T_k < T_0) \cdot \mathbb{E}_k[\Lambda_k] = \tfrac{q_0}{1-p_0^k} \cdot \mathbb{E}_k[\Lambda_k].$$

When the chain leaves state k, it moves up with probability $1/(1+p_0)$ and down with probability $p_0/(1+p_0)$. Starting from $k+1$, the probability that the chain does not return to k (probability it is absorbed at N) is $q_0/(1-p_0^{N-k})$ by (B.12), and starting from $k-1$, the probability it does not return to k (probability it is absorbed at 0) is $1-(1-p_0^{k-1})/(1-p_0^k)$ again by (B.12). Thus, starting from k, Λ_k has the geometric distribution with support $\{1, 2, \ldots\}$ and success probability

$$\tfrac{1}{1+p_0} \cdot \tfrac{q_0}{1-p_0^{N-k}} + \tfrac{p_0}{1+p_0} \cdot \big(1 - \tfrac{1-p_0^{k-1}}{1-p_0^k}\big) = \tfrac{q_0(1-p_0^N)}{(1+p_0)(1-p_0^k)(1-p_0^{N-k})}.$$

It follows that

$$\begin{aligned} \mathbb{E}_1[\Lambda_k] &= \tfrac{(1+p_0)(1-p_0^{N-k})}{1-p_0^N}, \\ \mathbb{E}_1[\Lambda_{k,k+1}] &= \tfrac{1}{1+p_0} \cdot \mathbb{E}_1[\Lambda_k] = \tfrac{1-p_0^{N-k}}{1-p_0^N}. \end{aligned} \tag{B.13}$$

For $1 \leq k \leq N-1$, define $T_{k,k+1}^i$ as the (discrete) time of the i-th transition from

k to $k+1$ inductively by

$$T^i_{k,k+1} := \inf\{n > T^{i-1}_{k,k+1} : X_{n-1} = k, X_n = k+1\}, \quad i \geq 1,$$

with $T^0_{k,k+1} := 0$ and $\inf \varnothing = \infty$. A transition from k to $k+1$ in $(X_n)_{n\geq 0}$ occurs due to one of the k cells in the original process $(Z_0(t))_{t\geq 0}$ dividing. Assume $W_{i,k}$ mutations occur on the i-th such transition, where $W_{i,k}$ are i.i.d. nonnegative integer-valued random variables with $\mathbb{E}[W_{i,k}] = w$, independent of $(X_n)_{n\geq 0}$. Enumerate the cells at time $T^i_{k,k+1}$ as $Y^1_{i,k}, \ldots, Y^{k+1}_{i,k}$. By the same argument as laid out in part (1) above, we can assume that each mutation is assigned to the first cell. Let $Y^m_{i,k}(n)$ be the number of descendants of cell $Y^m_{i,k}$ at time step $T^i_{k,k+1} + n$, with $Y^m_{i,k}(0) = 1$. Then define the event

$$A_{j,k,i,\ell} := \big\{T_N < T_0,\ T^i_{k,k+1} < \infty,\ \ell \leq W_{i,k},\ Y^1_{i,k}(T_N - T^i_{k,k+1}) = j\big\}.$$

This is the event that the random walk eventually hits level N, that it transitions at least i times from k to $k+1$ before doing so, that at least ℓ mutations occur on the i-th such transition, and that the ℓ-th mutation is found in j cells at level N. We can then write

$$\begin{aligned}\mathbb{E}[S_j(\tau_N)|\tau_N < \infty] &= \big(\mathbb{P}_1(T_N < T_0)\big)^{-1} \cdot \textstyle\sum_{k=1}^{N-1}\sum_{i=1}^{\infty}\sum_{\ell=1}^{\infty} \mathbb{P}_1\big(A_{j,k,i,\ell}\big) \\ &= \tfrac{1-p_0^N}{q_0} \cdot \textstyle\sum_{k=1}^{N-1}\sum_{i=1}^{\infty}\sum_{\ell=1}^{\infty} \mathbb{P}_1\big(A_{j,k,i,\ell}\big).\end{aligned}$$

To compute $\mathbb{P}_1\big(A_{j,k,i,\ell}\big)$, note first that by independence,

$$\begin{aligned}&\mathbb{P}_1\big(T_N < T_0, T^i_{k,k+1} < \infty, \ell \leq W_{i,k}, Y^1_{i,k}(T_N - T^i_{k,k+1}) = j\big) \\ &= \mathbb{P}\big(\ell \leq W_{i,k}\big) \cdot \mathbb{P}_1(T_N < T_0, T^i_{k,k+1} < \infty, Y^1_{i,k}(T_N - T^i_{k,k+1}) = j\big).\end{aligned}$$

By the strong Markov property,

$$\begin{aligned}&\mathbb{P}_1(T_N < T_0, T^i_{k,k+1} < \infty, Y^1_{i,k}(T_N - T^i_{k,k+1}) = j\big) \\ &= \mathbb{P}_1(T_N < T_0, Y^1_{i,k}(T_N - T^i_{k,k+1}) = j | T^i_{k,k+1} < \infty\big) \cdot \mathbb{P}_1(T^i_{k,k+1} < \infty) \\ &= \mathbb{P}_{k+1}\big(T_N < T_0, Y^1(T_N) = j\big) \cdot \mathbb{P}_1(T^i_{k,k+1} < \infty),\end{aligned}$$

where we restart the chain at the stopping time $T^i_{k,k+1}$ with $k+1$ cells enumerated as $Y^1, \ldots, Y^{k+1}$. Define $h^{(j,N-j)}_{(1,k)} := \mathbb{P}_{k+1}(T_N < T_0, Y^1(T_N) = j)$ for the moment, i.e. the probability that starting with one cell carrying a particular mutation and k cells without it, j cells carry the mutation when the population reaches size N. We can then write

$$\mathbb{P}_1\big(T_N < T_0, T^i_{k,k+1} < \infty, Y^1_{i,k}(T_N - T^i_{k,k+1}) = j\big) = h^{(j,N-j)}_{(1,k)} \cdot \mathbb{P}_1(T^i_{k,k+1} < \infty).$$

Combining the above, and using $\mathbb{E}_1[\Lambda_{k,k+1}] = (1 - p_0^{N-k})/(1 - p_0^N)$ by (B.13), we obtain

$$\begin{aligned}
&\mathbb{E}[S_j(\tau_N)|\tau_N < \infty] \\
&= \tfrac{1-p_0^N}{q_0} \cdot \textstyle\sum_{k=1}^{N-1} \big(\sum_{i=1}^{\infty} \big(\sum_{\ell=1}^{\infty} \mathbb{P}(W_{i,k} \geq \ell)\big) \mathbb{P}_1(T^i_{k,k+1} < \infty)\big) h^{(j,N-j)}_{(1,k)} \\
&= \tfrac{1-p_0^N}{q_0} \cdot \textstyle\sum_{k=1}^{N-1} w \cdot \mathbb{E}_1[\Lambda_{k,k+1}] \cdot h^{(j,N-j)}_{(1,k)} \\
&= (w/q_0) \cdot \textstyle\sum_{k=1}^{N-1} (1 - p_0^{N-k}) \cdot h^{(j,N-j)}_{(1,k)},
\end{aligned}$$

which is the desired result.

It remains to determine how the probabilities $h^{(j,N-j)}_{(1,k)}$ for $1 \leq j \leq N$ can be computed. As in the proof of (1) of Proposition 3.1, one can view $h^{(r,N-r)}_{(\ell,m)}$ as the probability of absorption in state $(r, N-r)$, starting from state (ℓ, m), for a Markov chain on the state space $\mathcal{S} := \{(\ell, m) : \ell, m \geq 0 \text{ and } \ell + m \leq N\}$, where ℓ is the number of cells carrying a particular mutation and m is the number of cells without it. The difference is that now, cells can die, so population level changes can be both up and down. The transition probabilities are therefore more complex in this case,

and given by

$$
\begin{aligned}
(\ell,m) \to (\ell+1,m) \quad &\text{w.p.} \quad \ell r_0/((\ell+m)(r_0+d_0)) \\
&\qquad\quad = \ell/((\ell+m)(1+p_0)), \\
(\ell,m) \to (\ell-1,m) \quad &\text{w.p.} \quad \ell d_0/((\ell+m)(r_0+d_0)) \\
&\qquad\quad = \ell p_0/((\ell+m)(1+p_0)), \\
(\ell,m) \to (\ell,m+1) \quad &\text{w.p.} \quad m r_0/((\ell+m)(r_0+d_0)) \\
&\qquad\quad = m/((\ell+m)(1+p_0)), \\
(\ell,m) \to (\ell,m-1) \quad &\text{w.p.} \quad m d_0/((\ell+m)(r_0+d_0)) \\
&\qquad\quad = m p_0/((\ell+m)(1+p_0)),
\end{aligned}
$$

for (ℓ,m) with $\ell,m \geq 0$ and $0 < \ell+m < N$. The states $(0,0)$ and $(\ell, N-\ell)$ with $0 \leq \ell \leq N$ are absorbing. A diagram of the Markov chain is shown in Figure B.2.

For given $(r,s) \in A := \{(0,0)\} \cup \{(r,s) : r,s \geq 0 \text{ and } r+s = N\}$ and (ℓ,m) with $\ell,m \geq 0$ and $0 < \ell+m < N$, by conditioning on the first transition out of state (ℓ,m), we can derive the following recursion for $h^{(r,s)}_{(\ell,m)}$:

$$
(\ell+m)(1+p_0)h^{(r,s)}_{(\ell,m)} = \ell h^{(r,s)}_{(\ell+1,m)} + \ell p_0 h^{(r,s)}_{(\ell-1,m)} + m h^{(r,s)}_{(\ell,m+1)} + m p_0 h^{(r,s)}_{(\ell,m-1)}.
$$

By the gambler's ruin formula (B.12), the boundary conditions are

$$
\begin{aligned}
&h^{(N,0)}_{(\ell,0)} = 1 - h^{(0,0)}_{(\ell,0)} = (1-p_0^\ell)/(1-p_0^N), \quad 0 \leq \ell \leq N, \\
&h^{(r,s)}_{(\ell,0)} = 0, \quad 0 \leq \ell \leq N,\ (r,s) \notin \{(0,0),(N,0)\}, \\
&h^{(0,N)}_{(0,m)} = 1 - h^{(0,0)}_{(0,m)} = (1-p_0^m)/(1-p_0^N), \quad 0 \leq m \leq N, \\
&h^{(r,s)}_{(0,m)} = 0, \quad 0 \leq m \leq N,\ (r,s) \notin \{(0,0),(0,N)\}, \\
&h^{(r,N-r)}_{(\ell,N-\ell)} = \delta_{r,\ell}, \quad 1 \leq \ell \leq N-1, \\
&h^{(r,s)}_{(\ell,N-\ell)} = 0, \quad 1 \leq \ell \leq N-1,\ (r,s) \in \{(0,0),(N,0),(0,N)\}.
\end{aligned}
\tag{B.14}
$$

This is the desired linear system. □

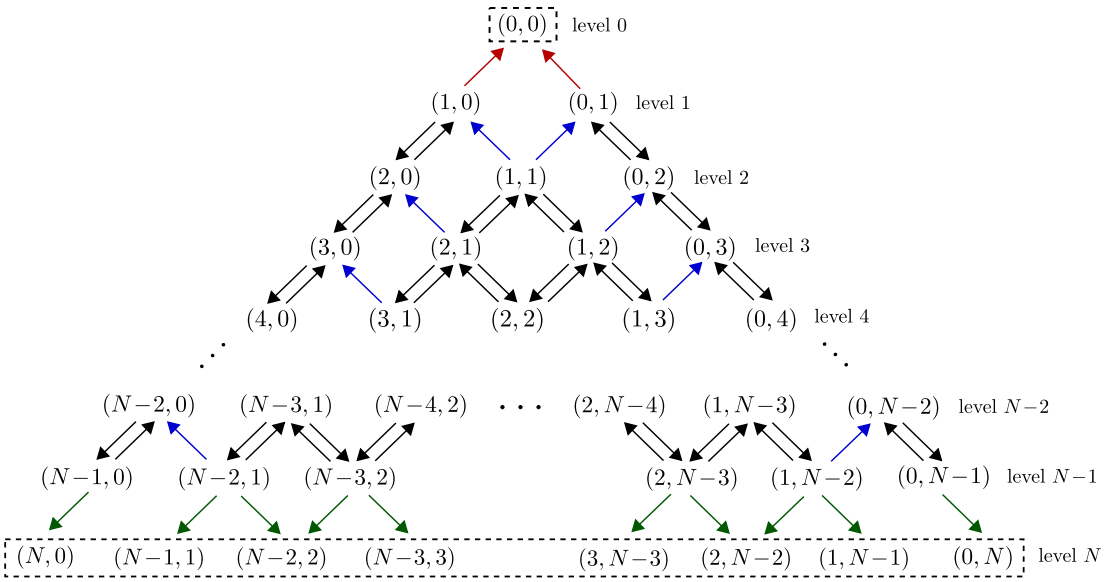

Figure B.2: A diagram of the discrete-time Markov chain on $\mathcal{S} := \{(\ell, m) : \ell, m \geq 0, \ell + m \leq N\}$, where ℓ is the number of cells carrying a particular mutation and m is the number of cells without it. Contrary to the chain for the skeleton process of Proposition 3.1 (Figure B.1), we now incorporate cell death, which means both that population level changes can be up and down, and that we add states of the form (ℓ, m) with $\ell = 0$ or $m = 0$. The states $(0, 0)$ and $\{(\ell, N - \ell)\}_{0 \leq \ell \leq N}$ are absorbing (dashed boxes), and the states $\{(\ell, m) : \ell, m \geq 1, \ell + m < N\}$, $\{(\ell, 0)\}_{1 \leq \ell \leq N-1}$ and $\{(0, m)\}_{1 \leq m \leq N-1}$ form their respective communicating classes. Colored arrows indicate transitions out of communicating classes.

B.3 Results for continuous mutation accumulation

In this section, we discuss how to derive the expected SFS of the skeleton and the total population under the continuous model of mutation accumulation (see Section 3.2.3). This requires minor modifications to the proofs of Propositions 3.1 and 3.2.

B.3.1 Skeleton spectrum

Under continuous mutation accumulation, the expected fixed-time spectrum of the skeleton is given by

$$\mathbb{E}[\tilde{S}_j(\tilde{t}_N)] = (\nu/\lambda_0)N \cdot \int_0^{1-1/N} (1-y)y^{j-1}dy,$$

which is the same as (3.12) of Proposition 3.1 with w/q_0 replaced by ν/λ_0, the effective mutation rate in the continuous-time model. We can use the same proof as in part (1) of Appendix B.1, simply replacing the mutation rate wr_0 by ν and noting that $q_0 = \lambda_0/r_0$. However, the expected fixed-size spectrum of the skeleton becomes

$$\mathbb{E}[\tilde{S}_j(\tilde{\tau}_N)] = \begin{cases} (\nu/\lambda_0)N \cdot 1/(j(j+1)) - (\nu/\lambda_0)\delta_{1,j}, & 1 \leq j \leq N-1, \\ \nu/\lambda_0, & j = N. \end{cases} \tag{B.15}$$

In the continuous model, mutations occur at rate ν per unit time, and the effective type-2 cell divisions occur at rate λ_0 per unit time. Thus, for $1 \leq k \leq N-1$, the number of mutations that accumulate on skeleton population size level k, prior to the type-2 division that changes levels to $k+1$, has the geometric distribution with support $\{0, 1, 2, \ldots\}$ and success probability

$$k\lambda_0/(k\lambda_0 + k\nu) = \lambda_0/(\lambda_0 + \nu).$$

The expected number of mutations per level is therefore

$$(\lambda_0 + \nu)/\lambda_0 - 1 = \nu/\lambda_0,$$

which applies to all levels k with $1 \leq k \leq N-1$. In particular, there are ν/λ_0 $(= w/q_0)$ clonal mutations in the continuous model, as opposed to $wp_0/q_0 = w/q_0 - w$ clonal mutations in the discrete model. Recall that in the latter model, mutations coincide with cell divisions, and clonal mutations come from the type-1 divisions that occur before the first type-2 division in the process. The first type-2 division adds w mutations, but it also changes levels, so mutations occurring on this division are not clonal. In the continuous model, all mutations occur in between cell divisions, which is why there is no such boundary effect. Similarly, in the discrete model, the very last type-2 division that changes the skeleton size to N adds w mutations on average that each ends up in one cell. Since this does not occur in the continuous model, the extra term $-\nu/\lambda_0$ $(= -w/q_0)$ that appears for $j = 1$ in (B.15) differs from the extra term $-wp_0/q_0 = -w/q_0 + w$ in the discrete model by w. These key differences between mutation accumulation in the two models are diagramed in Figure B.3.

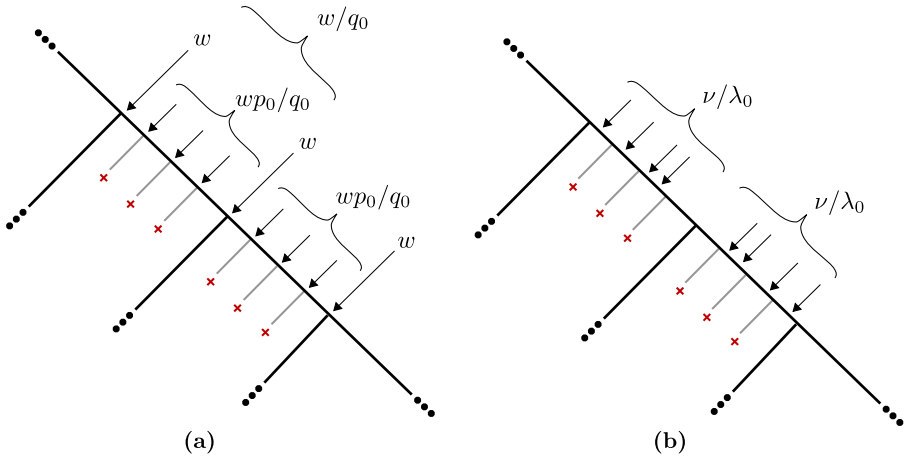

Figure B.3: Mutation accumulation on the skeleton in the discrete vs. continuous model of mutation. **(a)** In the discrete model, mutations coincide with cell divisions. On average, wp_0/q_0 mutations accumulate on type-1 divisions in between two type-2 divisions, and w mutations are added on each type-2 division. This results in $wp_0/q_0 + w = w/q_0$ mutations on average per skeleton population size level, for all but the first level. **(b)** In the continuous model, all ν/λ_0 $(= w/q_0)$ mutations per level accumulate in between cell divisions. The differences between the discrete and continuous model result in slightly different behavior at the boundary values $j = 1$ and $j = N$ in the fixed-size spectrum (3.14) of Proposition 3.1 and fixed-size result (B.15) for the continuous model.

B.3.2 Total population spectrum

Under continuous mutation accumulation, the expected fixed-time spectrum of the total population is given by

$$\mathbb{E}[S_j(t_N)|Z_0(t_N) > 0] = (\nu/r_0)N \cdot \int_0^{1-1/N} (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy,$$

which is the same as (3.20) of Proposition 3.2 with w replaced by ν/r_0. We can use the same proof as in part (1) of Appendix B.2, replacing the mutation rate wr_0 by ν. However, the expected fixed-size spectrum of the total population becomes, for $1 \leq j \leq N$,

$$\mathbb{E}[S_j(\tau_N)|\tau_N < \infty] = (\nu/\lambda_0) \cdot \sum_{k=1}^{N-1} (1 - p_0^{N-k}) \cdot h_{(1,k-1)}^{(j,N-j)}.$$

In the continuous model, mutations no longer coincide with changes in population level, so instead of counting level changes $\Lambda_{k,k+1}$ as we did in the proof of part (2) of Proposition 3.2, we need to compute the expected time spent at level k. We already know from (B.13) that the number of visits to level k in the embedded discrete-time chain $(X_n)_{n\geq 0}$ has expected value

$$\mathbb{E}_1[\Lambda_k] = \tfrac{(1+p_0)(1-p_0^{N-k})}{1-p_0^N}.$$

During each visit to state k in the discrete-time chain, the time spent at population level k in the continuous-time process $(Z_0(t))_{t\geq 0}$ is exponentially distributed with rate $k(r_0+d_0) = kr_0(1+p_0)$. It follows that the mean time spent on level k for $1 \leq k \leq N-1$ is

$$\tfrac{1}{kr_0(1+p_0)} \cdot \tfrac{(1+p_0)(1-p_0^{N-k})}{1-p_0^N} = \tfrac{1-p_0^{N-k}}{kr_0(1-p_0^N)}.$$

Since mutations occur at rate $k\nu$ per unit time on level k, we obtain

$$\begin{aligned}\mathbb{E}[S_j(\tau_N)|\tau_N < \infty] &= \tfrac{1-p_0^N}{q_0} \cdot \textstyle\sum_{k=1}^{N-1} k\nu \cdot \tfrac{1-p_0^{N-k}}{kr_0(1-p_0^N)} \cdot h_{(1,k-1)}^{(j,N-j)} \\ &= (\nu/\lambda_0) \cdot \textstyle\sum_{k=1}^{N-1}(1-p_0^{N-k}) \cdot h_{(1,k-1)}^{(j,N-j)}, \end{aligned} \tag{B.16}$$

where we use that $q_0 = \lambda_0/r_0$. We now have $h_{(1,k-1)}^{(j,N-j)}$ in the sum instead of $h_{(1,k)}^{(j,N-j)}$ in (3.22) of Proposition 3.2 since mutations no longer coincide with population level changes. The first equality in (B.16) can be obtained rigorously following the same line of reasoning as in the proof of part (2) of Proposition 3.2.

B.4 Proof of Proposition 3.3

In this section, we prove Proposition 3.3 on the total mutational burden of the tumor, both under the fixed-time and fixed-size spectrum. The result follows from Proposition 3.2 using simple calculations.

Proof of Proposition 3.3. (1) Define $M_j(t) := \sum_{k\geq j} S_k(t)$ as the cumulative number of mutations found in $\geq j$ cells at time t. For fixed $j \geq 1$, by (3.20) of Proposition 3.2 and Fubini's theorem, the expected cumulative fixed-time spectrum can be written

as

$$\begin{aligned}
\mathbb{E}[M_j(t_N)|Z_0(t_N)>0] &= wN\cdot\textstyle\sum_{k=j}^{\infty}\big(\int_0^{1-1/N}(1-p_0y)^{-1}(1-y)y^{k-1}dy\big)\\
&= wN\cdot\textstyle\int_0^{1-1/N}(1-p_0y)^{-1}(1-y)\big(\sum_{k=j}^{\infty}y^{k-1}\big)dy\\
&= wN\cdot\textstyle\int_0^{1-1/N}(1-p_0y)^{-1}y^{j-1}dy\\
&= wN\cdot\textstyle\sum_{k=0}^{\infty}p_0^k\big(\int_0^{1-1/N}y^{j+k-1}dy\big)\\
&= wN\cdot\textstyle\sum_{k=0}^{\infty}\frac{p_0^k}{j+k}(1-\frac{1}{N})^{j+k}. \qquad \text{(B.17)}
\end{aligned}$$

To obtain the desired result, set $j=1$ in (B.17) and use that $\sum_{k=1}^{\infty}x^k/k=-\log(1-x)$.

(2) By (3.22) of Proposition 3.2,

$$\begin{aligned}
\mathbb{E}[M_1(\tau_N)|\tau_N<\infty] &= \textstyle\sum_{j=1}^{N}\big((w/q_0)\cdot\sum_{k=1}^{N-1}(1-p_0^{N-k})\cdot h_{(1,k)}^{(j,N-j)}\big)\\
&= (w/q_0)\cdot\textstyle\sum_{k=1}^{N-1}(1-p_0^{N-k})\cdot\big(\sum_{j=1}^{N}h_{(1,k)}^{(j,N-j)}\big),
\end{aligned}$$

and the result follows from the fact that $\sum_{(r,s)\in A}h_{(1,k)}^{(r,s)}=1$. □

B.5 Fixed-time vs. fixed-size total population spectrum

Here, we present a simple heuristic argument for why the fixed-size spectrum of the total population can be approximated by the fixed-time spectrum on $j\ll N$ when N is sufficiently large. As we discussed in Section 3.4.5 of the main text, conditional on the nonextinction event Ω_∞, the tumor eventually grows at exponential rate λ_0. If s is the time it takes to go from population level k to population level N, and we assume that s can be treated as deterministic, we can write $ke^{\lambda_0 s}=N$ i.e. $e^{\lambda_0 s}=N/k$, following e.g. Iwasa et al. [81]. We can then make the approximation $h_{(1,k)}^{(j,N-j)}\approx p_j(s)$, where $h_{(1,k)}^{(j,N-j)}$ is defined as in the proof of part (2) of Proposition 3.2, and $(p_j(s))_{j\geq 0}$ is the size-distribution at time s for a single-cell derived clone. Applying (B.5) with $e^{\lambda_0 s}=N/k$, we obtain

$$h_{(1,k)}^{(j,N-j)}\approx p_j(s)=\tfrac{q_0^2Nk}{(N-p_0k)^2}\cdot\big(\tfrac{N-k}{N-p_0k}\big)^{j-1}.$$

Then, observing that $1 - p_0^{N-k} \approx 1$ when $k \ll N$, we can write

$$(w/q_0) \cdot \sum_{k=1}^{N-1} (1 - p_0^{N-k}) \cdot h_{(1,k)}^{(j,N-j)} \approx (w/q_0) \cdot \int_1^N \tfrac{q_0^2 N k}{(N-p_0 k)^2} \cdot \left(\tfrac{N-k}{N-p_0 k}\right)^{j-1} dk.$$

Using the substitution $y := (N-k)/(N-p_0 k)$ and writing $N - p_0 \approx N$, this becomes the fixed-time spectrum (3.20) of Proposition 3.2.

B.6 Derivation of expressions (3.24) and (3.25)

Here, we establish the asymptotic expressions (3.24) and (3.25) in the main text. To establish (3.24), fix $0 < p_0 < 1$ and set

$$f_j(k) := p_0^k \cdot \tfrac{j(j+1)}{(j+k)(j+k+1)}, \quad j \geq 1,\ k \geq 0.$$

Clearly, $f_j(k) \leq p_0^k$ for all $j \geq 1$ and $k \geq 0$. Since $\sum_{k=0}^{\infty} p_0^k = 1/q_0 < \infty$, it follows from the dominated convergence theorem that

$$\lim_{j\to\infty} \sum_{k=0}^{\infty} f_j(k) = 1/q_0,$$

from which it follows that

$$wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{(j+k)(j+k+1)} \sim (w/q_0)N \cdot 1/(j(j+1)), \quad j \to \infty.$$

To establish (3.25), fix $j \geq 1$ and set

$$f_{p_0}(k) := \tfrac{p_0^k}{(j+k)(j+k+1)}, \quad 0 < p_0 < 1,\ k \geq 0.$$

Clearly, $f_{p_0}(k) \leq 1/((j+k)(j+k+1))$ for all $0 < p_0 < 1$ and $k \geq 0$. Since

$$\sum_{k=0}^{\infty} \tfrac{1}{(j+k)(j+k+1)} = \sum_{k=0}^{\infty} \left(\tfrac{1}{j+k} - \tfrac{1}{j+k+1}\right) = 1/j < \infty,$$

it follows from the dominated convergence theorem that

$$\lim_{p_0\to 1} \sum_{k=0}^{\infty} f_{p_0}(k) = 1/j,$$

from which it follows that

$$wN \cdot \sum_{k=0}^{\infty} \frac{p_0^k}{(j+k)(j+k+1)} \sim wN \cdot 1/j, \quad p_0 \to 1.$$

B.7 Derivation of expression (3.30)

Here, we establish expression (3.30) in the main text. By (3.20) of Proposition 3.2 and Fubini's theorem, we can write

$$\begin{aligned}
&\mathbb{E}[S_1(t_N)|Z_0(t_N) > 0] \\
&= wN \cdot \int_0^{1-1/N} (1-p_0 y)^{-1}(1-y)dy \\
&= wN \cdot \sum_{k=0}^{\infty} p_0^k \big(\int_0^{1-1/N} y^k (1-y)dy \big) \\
&= wN \cdot \sum_{k=0}^{\infty} p_0^k \big(\tfrac{1}{k+1}(1-\tfrac{1}{N})^{k+1} - \tfrac{1}{k+2}(1-\tfrac{1}{N})^{k+2} \big) \\
&= wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{k+1}(1-\tfrac{1}{N})^{k+1} - wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{k+2}(1-\tfrac{1}{N})^{k+2}. \qquad \text{(B.18)}
\end{aligned}$$

Using that $\sum_{k=1}^{\infty} x^k/k = -\log(1-x)$, the former term can be computed as

$$wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{k+1}(1-\tfrac{1}{N})^{k+1} = -wN \cdot (1/p_0)\log(q_0 + p_0/N),$$

and the latter term can be computed as

$$\begin{aligned}
wN \cdot \sum_{k=0}^{\infty} \tfrac{p_0^k}{k+2}(1-\tfrac{1}{N})^{k+2} &= wN \cdot (1/p_0^2) \sum_{k=0}^{\infty} \tfrac{p_0^{k+2}}{k+2}(1-\tfrac{1}{N})^{k+2} \\
&= wN \cdot (1/p_0^2)\big(-\log(q_0 + p_0/N) - p_0(1-1/N) \big).
\end{aligned}$$

Combining with (B.18), we obtain

$$\mathbb{E}[S_1(t_N)|Z_0(t_N) > 0] = wN \cdot (1/p_0)(1 - 1/N + (q_0/p_0)\log(q_0 + p_0/N)),$$

the desired result.

B.8 Laws of large numbers

Here, we present simple calculations in support of the conjectured laws of large numbers (3.32) and (3.33) of the main text. As stated in the main text, conditional on the nonextinction event Ω_∞, we have $Z_0(t) \sim Ye^{\lambda_0 t}$ as $t \to \infty$ almost surely, where Y follows the exponential distribution with mean $1/q_0$ (Theorem 1 of Durrett [50]). For the fixed-time spectrum, the number of mutations that accumulate in $[0, t_N]$ and are found in $j \geq 1$ cells at time t_N is then approximately

$$S_j(t_N) \approx \int_0^{t_N} wr_0 \cdot Ye^{\lambda_0 t} \cdot p_j(t_N - t)dt.$$

From (B.10) in the proof of Proposition 3.2, we know that

$$\begin{aligned}
&\int_0^{t_N} wr_0 e^{\lambda_0 t} p_j(t_N - t)dt \\
&= we^{\lambda_0 t_N} \cdot \int_0^{1-q_0/(e^{\lambda_0 t_N} - p_0)} (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy \\
&\sim wq_0 N \cdot \int_0^1 (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy, \quad N \to \infty,
\end{aligned}$$

where we use that $e^{\lambda_0 t_N} = q_0 N + p_0$ by the definition of t_N in (3.18). This implies that

$$S_j(t_N) \approx q_0 Y \cdot wN \cdot \int_0^1 (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy$$

for large N. Since Y has the exponential distribution with mean $1/q_0$, $q_0 Y$ has the exponential distribution with mean 1. This suggests (3.32) in the main text.

For the fixed-size spectrum, note that if N is large, then at time $\tau_N - t$, we can write $Z_0(\tau_N - t) \approx Ne^{-\lambda_0 t}$. The number of mutations that accumulate in $[0, \tau_N]$ and are found in $j \geq 1$ cells at time τ_N is then approximately

$$\begin{aligned}
S_j(\tau_N) &\approx \int_0^{\tau_N} wr_0 \cdot Ne^{-\lambda_0 t} \cdot p_j(t)dt \\
&= Ne^{-\lambda_0 \tau_N} \cdot \int_0^{\tau_N} wr_0 e^{\lambda_0 t} p_j(\tau_N - t)dt.
\end{aligned}$$

Again using (B.10) from the proof of Proposition 3.2, we can write

$$\int_0^{\tau_N} wr_0 e^{\lambda_0 t} p_j(\tau_N - t)dt = we^{\lambda_0 \tau_N} \cdot \int_0^{1-q_0/(e^{\lambda_0 \tau_N} - p_0)} (1 - p_0 y)^{-1}(1 - y)y^{j-1}dy,$$

from which it follows that

$$\begin{aligned} S_j(\tau_N) &\approx wN \cdot \textstyle\int_0^{1-q_0/(e^{\lambda_0\tau_N}-p_0)}(1-p_0y)^{-1}(1-y)y^{j-1}dy \\ &\approx wN \cdot \textstyle\int_0^1(1-p_0y)^{-1}(1-y)y^{j-1} \end{aligned}$$

for N large. This suggests (3.33) in the main text.

We finally mention that it is straightforward to prove a law of large numbers for a simplified version of our model, where the tumor bulk grows deterministically ($Z_0(t) = e^{\lambda_0 t}$), mutant clones arise at stochastic rate wr_0, and mutant clones grow stochastically. To state the result, let $\hat{S}_j(\tilde{t}_N)$ denote the number of mutations found in $j \geq 1$ cells at time $\tilde{t}_N$ under the simplified model, where $\tilde{t}_N$ is given by (3.10), i.e. $e^{\lambda_0\tilde{t}_N} = N$. We want to show that

$$\hat{S}_j(\tilde{t}_N) \sim wN \cdot \textstyle\int_0^1(1-p_0y)^{-1}(1-y)y^{j-1}dy \tag{B.19}$$

as $N \to \infty$ almost surely. Note that the limit is a constant since we assume deterministic growth of the tumor bulk. For $0 \leq t \leq \tilde{t}_N$, let $\hat{N}_{j,\tilde{t}_N}(t)$ denote the number of mutant clones created in $[0,t]$ that have size $j \geq 1$ at time $\tilde{t}_N$. Then $(\hat{N}_{j,\tilde{t}_N}(t))_{0\leq t\leq \tilde{t}_N}$ is an inhomogeneous Poisson process with rate function $\hat{\lambda}(t) = wr_0e^{\lambda_0 t}p_j(\tilde{t}_N - t)$ and mean function

$$\hat{m}(t) = \textstyle\int_0^t \hat{\lambda}(s)ds, \quad 0 \leq t \leq \tilde{t}_N.$$

Set $\hat{N}_j(\tilde{t}_N) := \hat{N}_{j,\tilde{t}_N}(\tilde{t}_N)$. By (B.10) in the proof of Proposition 3.2, and the fact that $e^{\lambda_0\tilde{t}_N} = N$,

$$\begin{aligned} \hat{m}(\tilde{t}_N) &= wN \cdot \textstyle\int_0^{1-q_0/(N-p_0)}(1-p_0y)^{-1}(1-y)y^{j-1}dy \\ &\sim wN \cdot \textstyle\int_0^1(1-p_0y)^{-1}(1-y)y^{j-1}dy, \quad N \to \infty. \end{aligned}$$

Then, by a simple Poisson concentration inequality, see Theorem 1 of Cannone [26],

$$\begin{aligned} \mathbb{P}\big(|\hat{N}_j(\tilde{t}_N)/\hat{m}(\tilde{t}_N) - 1| > (\hat{m}(\tilde{t}_N))^{-1/3}\big) &= \mathbb{P}\big(|\hat{N}_j(\tilde{t}_N) - \hat{m}(\tilde{t}_N)| > (\hat{m}(\tilde{t}_N))^{2/3}\big) \\ &\leq 2\exp\big(-(\hat{m}(\tilde{t}_N))^{1/3}/\big(2(1+(\hat{m}(\tilde{t}_N))^{-1/3})\big)\big). \end{aligned}$$

Since $\hat{m}(\tilde{t}_N)$ is of order N as $N \to \infty$, it follows from the Borel-Cantelli lemma that

$\hat{N}_j(\tilde{t}_N)/\hat{m}(\tilde{t}_N) \to 1$ as $N \to \infty$ almost surely. Since $\hat{S}_j(\tilde{t}_N) = \hat{N}_j(\tilde{t}_N)$, we have the result.

Appendix C

Appendix to Chapter 4

C.1 Parametrization

To parametrize the model, we rely on Sharma et al.'s [139] investigation of an *EGFR*-mutant non-small-cell lung cancer (NSCLC) population (PC9) treated with 2 μM erlotinib, with fresh drug added every 3 days, and Hata et al.'s [78] study of the evolution of resistance conferred by the $EGFR^{\mathrm{T790M}}$ gatekeeper mutation in PC9 cells.

Most of the discussion in the main text concerns behavior in the presence of an anti-cancer agent, so we start with those dynamics. Following Hata et al. [78], we assume that drug-tolerant type-1 cells give birth at rate $r_1 = 0.0162$ per hour and die at rate $d_1 = 0.015$ per hour, and that stably resistant type-2 cells give birth at rate $r_2 = 0.04$ and die at rate $d_2 = 0.0015$. Although these rates apply to treatment with 300 nM gefitinib as opposed to 2 μM erlotinib, viability curves presented in [139] indicate that it is reasonable to assume similar rates for the case of 2 μM erlotinib treatment.

In Sharma et al. [139], the authors report that under 9 days of continuous anti-cancer treatment, almost all cells in the population die, while persister cells corresponding to around $0.27\% \pm 0.21\%$ of the original population survive. Assuming exponential decay of the drug-sensitive type-0 population, we estimate a net birth rate of $\lambda_0 = -0.04$ per hour for type-0 cells and set $r_0 = 0.04$ and $d_0 = 0.08$, where we assume without loss of generality that the anti-cancer agent increases the death rate of type-0 cells without affecting their birth rate (the birth rate of type-0 cells in the absence of the anti-cancer agent is $r_0 = 0.04$ as is discussed below). Note that if we assume instead that the drug

decreases the birth rate of type-0 cells, with the net birth rate $\lambda_0 = -0.04$ unchanged, the type-0 population will decay at the same net rate as before, but the mutation rate $r_0\xi_0$ of type-0 cells will decrease proportionally to the decrease in r_0. We can therefore view variation in the effect of the drug on the birth rate r_0 as equivalent to variation in the mutation rate ξ_0.

The rates μ of epimutation and ν of reversion are more difficult to estimate, since it is not clear in [139] how many of the type-1 persister cells that survive the first 9 days of treatment are present at the onset. However, the authors do conclude from experiments applying the histone deacetylase inhibitor TSA as pretreatment to erlotinib that type-1 cells do emerge *de novo* during treatment. For our baseline parameter regime, we will assume that around half the cells that survive the first 9 days of treatment are present at the onset and set $\mu = 4 \cdot 10^{-5}$ per hour accordingly, which implies that an epimutation occurs once in every 1,000 divisions of a type-0 cell. In determining ν, we assume that cells transition more freely out of the resistant state than into it and set $\nu = 4 \cdot 10^{-4}$, which is consistent with observed phenotypic switching dynamics between persister cells and normal cells in *Escherichia coli* bacterial populations [120], and between stem-like and non-stem-like cells in breast cancer [77].

The point mutation rate per nucleotide per cell division has been estimated as $5 \cdot 10^{-10}$ [50]. To obtain the rate of mutations that confer resistance to anti-cancer therapy, this number needs to be multiplied by the number of resistance-conferring point mutations. Following e.g. [78] and [15], we assume a baseline rate of $\xi_0 = \xi_1 = 10^{-7}$ per cell division, although a reasonable range can be anywhere from 10^{-9} to 10^{-5}, which is the range suggested by [50] for mutations leading to cancer. This assumption on ξ_0 and ξ_1 translates into a mutation rate of $r_0\xi_0 = 4 \cdot 10^{-9}$ per hour for type-0 cells and $r_1\xi_1 = 1.62 \cdot 10^{-9}$ per hour for type-1 cells.

As for the rate η of epigenetic reprogramming, we note that measurements provided by Bintu et al. [11] of the dynamics of epigenetic silencing under recruitment of chromatin regulators suggest that it may be natural to assume that η is around 1-2 orders of magnitude lower than μ. We will apply $\eta = 4 \cdot 10^{-7}$ or $\eta = 4 \cdot 10^{-8}$ as our baseline rate depending on the context, which is 2-3 orders of magnitude smaller than μ, and 1-2 orders of magnitude larger than the mutation rates $r_0\xi_0$ and $r_1\xi_1$. As is the case for the mutation rate, one can expect significant variation in the epigenetic reprogramming

	Presence of drug
r_0	0.04
d_0	0.08
r_1	0.0162
d_1	0.015
μ	$4 \cdot 10^{-5}$
ν	$4 \cdot 10^{-4}$
η	$4 \cdot 10^{-7}$ or $4 \cdot 10^{-8}$
$\xi_0 = \xi_1$	10^{-7}
r_2	0.04
d_2	0.0015

Table C.1: Baseline parameter regime in the presence of an anti-cancer agent. All rates are measured per hour, except the mutation rates ξ_0 and ξ_1, which are measured per cell division.

rate depending on the cancer type, the anti-cancer agent being applied and the concentration of this agent. We also note that 'stable epigenetic resistance' may not even be a well-defined concept if epigenetic reprogramming occurs in a multi-stage or continuous fashion, with each stage leading to an increasingly stable phenotype, which is either less likely to revert back to sensitivity once removed from drug or does so on a longer time scale.

In the main text, we show all results assuming an initial population size of 10^6 cells, following the *in vitro* experiments conducted by Sharma et al. We note that 10^6 refers to the *effective population size*, i.e. those tumor cells that are capable of undergoing phenotypic switching, which may only apply to a subset of tumor cells [113]. We also note that the initial population size should be viewed in context of the mutation rates $r_0\xi_0 = 4 \cdot 10^{-9}$ and $r_1\xi_1 = 1.62 \cdot 10^{-9}$ and the reprogramming rates $\eta = 4 \cdot 10^{-7}$ or $\eta = 4 \cdot 10^{-8}$, since the relationship between these three parameters largely determines the dynamics of resistance acquisition.

The only section in the main text where tumor dynamics in the absence of an anti-cancer agent are considered is Section 4.4.4. Since most of the data referenced above concerns behavior in the presence of drug, we will mostly assume that the dynamics on and off drug are identical, with the important exception of the birth and death rate of

type-0 cells, which we assume to be $r_0 = 0.04$ and $d_0 = 0.0015$ off drug, following Hata et al. [78]. The resulting parameter values (Table C.2) therefore reflect the qualitative setting where type-0 cells proliferate rapidly in the absence of drug but die rapidly in its presence, and type-1 cells proliferate slowly both on and off drug. The assumption that type-1 cells are at a selective disadvantage off drug is not essential to our results, as is discussed in Appendix C.8 below.

	Absence of drug
r_0	0.04
d_0	0.0015
r_1	0.0162
d_1	0.015
μ	$4 \cdot 10^{-5}$
ν	$4 \cdot 10^{-4}$
η	0
$\xi_0 = \xi_1$	10^{-7}
r_2	0.04
d_2	0.0015

Table C.2: Baseline parameter regime in the absence of an anti-cancer agent. All rates are measured per hour, except the mutation rates ξ_0 and ξ_1, which are measured per cell division.

We finally note that according to the above dynamics, type-1 cells will constitute around 0.11% of the population during long-term expansion in the absence of drug and at the start of treatment. In the main text, we generally consider both the case where no type-1 cell is present at the start of anti-cancer treatment ($m = 0$), and the case where type-1 cells constitute 0.1% of the population at the onset ($m \gg 0$).

C.2 Two-type Markovian branching processes

Here, we derive expression (4.3) in the main text and describe some of its properties. For the case $\eta = 0$ and $\xi_0 = \xi_1 = 0$, our model reduces to a two-type continuous-time Markovian branching process [6]. Let $X(t) = (X_0(t), X_1(t))$ denote such a process, where the two types are designated as type-0 and type-1. Associated with $X(t)$ is the

mean matrix $\mathbf{M}(t) = \{m_{ij}(t) : (i,j) \in \{0,1\}^2\}$, defined by

$$m_{ij}(t) = E[X_j(t)|X(0) = \mathbf{e}_i], \quad (i,j) \in \{0,1\}^2,$$

where $\mathbf{e}_i$ denotes the unit vector with 1 in the i-th coordinate, and the *infinitesimal generator* $\mathbf{A} = \{a_{ij} : (i,j) \in \{0,1\}^2\}$, which satisfies

$$\mathbf{M}(t) = \exp(\mathbf{A}t), \quad t \geq 0.$$

We can interpret $m_{ij}(t)$ as the mean number of type-j particles alive at time t, given that the process is started by a single type-i cell, and a_{ij} as the infinitesimal rate at which a type-i cell produces a type-j cell. We assume that $a_{12}, a_{21} > 0$, i.e. the rates of switching between types are strictly positive. For further discussion on the above matrices, see [6].

Let $\phi_0^{(n,m)}(t)$ and $\phi_1^{(n,m)}(t)$ denote the mean number of type-0's and type-1's alive at time t, assuming initial conditions $(X_0(0), X_1(0)) = (n,m)$, where n and m are nonnegative integers with $n + m > 0$. Using the mean matrix, we can easily compute these means as

$$[\phi_0^{(n,m)}(t) \;\; \phi_1^{(n,m)}(t)] = [n \;\; m]\, \mathbf{M}(t). \tag{C.1}$$

Note that the infinitesimal generator $\mathbf{A}$ possesses distinct real eigenvalues $\rho < \sigma$ given by

$$\frac{a_{11} + a_{22} \pm \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}}}{2}, \tag{C.2}$$

which follows from our assumption that $a_{12} > 0$ and $a_{21} > 0$. The eigenvalues of $\mathbf{M}(t) = \exp(\mathbf{A}t)$ are then easily obtained as $e^{\rho t}$ and $e^{\sigma t}$. If we define $\delta := (a_{11} - \rho)/a_{21}$ and $\gamma := (\sigma - a_{11})/a_{21}$, it is easily established that $\mathbf{v} = [1 \;\; \delta]$ and $\mathbf{w} = [1 \;\; \gamma]$ are left eigenvectors of $\mathbf{A}$ with respect to ρ and σ, respectively. Decomposing $[n \;\; m]$ in terms of

v and **w**, we can then compute the means in (C.1) explicitly as

$$\begin{aligned}\phi_0^{(n,m)}(t) &= \frac{n\delta+m}{\delta+\gamma}e^{\sigma t} + \frac{n\gamma-m}{\delta+\gamma}e^{\rho t},\\ \phi_1^{(n,m)}(t) &= \frac{\gamma(n\delta+m)}{\delta+\gamma}e^{\sigma t} - \frac{\delta(n\gamma-m)}{\delta+\gamma}e^{\rho t}.\end{aligned} \tag{C.3}$$

If we define $\alpha := (n\gamma - m)/(\delta+\gamma)$ and $\beta := (n\delta+m)/(\delta+\gamma)$, we can simplify these expressions further to

$$\begin{aligned}\phi_0^{(n,m)}(t) &= \beta e^{\sigma t} + \alpha e^{\rho t},\\ \phi_1^{(n,m)}(t) &= \gamma\beta e^{\sigma t} - \delta\alpha e^{\rho t}.\end{aligned} \tag{C.4}$$

Note that although we do not show it explicitly, the constants α and β depend on the initial conditions (n,m), whereas γ, δ, σ and ρ are completely determined by the infinitesimal generator **A**.

We next establish that that $a_{11}-\rho > 0$ and $\sigma - a_{11} > 0$, which will imply that $\gamma > 0$, $\delta > 0$ and $\beta > 0$ in (C.3) and (C.4). By rewriting (C.3) as

$$\begin{aligned}\phi_0^{(n,m)}(t) &= \frac{n}{\delta+\gamma}(\delta e^{\sigma t} + \gamma e^{\rho t}) + \frac{m}{\delta+\gamma}(e^{\sigma t} - e^{\rho t}),\\ \phi_1^{(n,m)}(t) &= \frac{n\delta\gamma}{\delta+\gamma}(e^{\sigma t} - e^{\rho t}) + \frac{m}{\delta+\gamma}(\gamma e^{\sigma t} + \delta e^{\rho t}),\end{aligned}$$

and noting that $\rho < \sigma$, it will also follow that the expected number of type-0's and type-1's is strictly positive for all $t > 0$, i.e. $X(t)$ is a positive-regular process [6].

To show that $a_{11}-\rho > 0$ and $\sigma - a_{11} > 0$, we first note that since ρ is an eigenvalue for **A**, we have $(a_{11}-\rho)(a_{22}-\rho) = a_{12}a_{21}$. By our assumption that $a_{12}, a_{21} > 0$, the terms $a_{11}-\rho$ and $a_{22}-\rho$ must then be nonzero and have the same sign. Using (E.5),

we get since $a_{12}, a_{21} > 0$:

$$\begin{aligned}\rho &= \frac{a_{11} + a_{22} - \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}}}{2} \\ &< \frac{a_{11} + a_{22} - |a_{11} - a_{22}|}{2} \\ &= \begin{cases} a_{11} \text{ if } a_{11} \leq a_{22} \\ a_{22} \text{ if } a_{11} > a_{22}. \end{cases}\end{aligned}$$

This implies that $a_{11} - \rho > 0$ or $a_{22} - \rho > 0$ depending on the relationship between a_{11} and a_{22}. Since $a_{11} - \rho$ and $a_{22} - \rho$ have the same sign, one being positive implies that the other one is as well, so in both cases, we have $a_{11} - \rho > 0$ and $a_{22} - \rho > 0$. A similar argument confirms that $\sigma - a_{11} > 0$.

Now let

$$\phi^{(n,m)}(t) := \phi_0^{(n,m)}(t) + \phi_1^{(n,m)}(t)$$

denote the expected total number of cells (type-0 and type-1) alive at time t, starting from n type-0 cells and m type-1 cells. By (C.4), we can write

$$\phi^{(n,m)}(t) = \beta(1+\gamma)e^{\sigma t} + \alpha(1-\delta)e^{\rho t}. \tag{C.5}$$

We then obtain the following expressions for the expected size of a clone, at time t, started by a single cell of each type:

$$\begin{aligned}\phi^{(1,0)}(t) &= \frac{\delta(1+\gamma)}{\delta+\gamma}e^{\sigma t} + \frac{\gamma(1-\delta)}{\delta+\gamma}e^{\rho t}, \\ \phi^{(0,1)}(t) &= \frac{1+\gamma}{\delta+\gamma}e^{\sigma t} - \frac{1-\delta}{\delta+\gamma}e^{\rho t}.\end{aligned} \tag{C.6}$$

We conclude by discussing large-t asymptotics. Since $\sigma > \rho$, the long-run behavior will be dominated by the former term in (C.3), (C.4), (C.5) and (C.6). In particular, we

can write

$$\begin{aligned}\phi_0^{(n,m)}(t) &= \beta e^{\sigma t} + o(e^{\sigma t}),\\ \phi_1^{(n,m)}(t) &= \gamma\beta e^{\sigma t} + o(e^{\sigma t}),\end{aligned} \tag{C.7}$$

and

$$\begin{aligned}\phi^{(1,0)}(t) &= \frac{\delta(1+\gamma)}{\delta+\gamma}e^{\sigma t} + o(e^{\sigma t}),\\ \phi^{(0,1)}(t) &= \frac{1+\gamma}{\delta+\gamma}e^{\sigma t} + o(e^{\sigma t}),\end{aligned} \tag{C.8}$$

where $o(e^{\sigma t})$ denotes a function that satisifies $o(e^{\sigma t})/e^{\sigma t} \to 0$ as $t \to \infty$. Note that from (C.7), it is clear that γ is the long-run ratio between type-1's and type-0's in the population, and from (C.8), we see that δ is the long-run ratio between mean clone size started by a single type-0 cell vs. a single type-1 cell. This gives an intuitive interpretation of δ and γ in (C.3), (C.4), (C.5) and (C.6), while α and β depend on n, m, δ and γ through $\alpha = (n\gamma - m)/(\delta+\gamma)$ and $\beta = (n\delta + m)/(\delta+\gamma)$.

C.3 Extinction probabilities

We continue to assume the absence of permanent resistance mechanisms ($\eta = 0$ and $\xi_0 = \xi_1 = 0$), which implies that our model still reduces to a two-type branching process $(X_0(t), X_1(t))$. Define

$$\begin{aligned}p_0 &:= P(X_0(t) + X_1(t) = 0 \text{ for some } t | X_0(0) = 1, X_1(0) = 0),\\ p_1 &:= P(X_0(t) + X_1(t) = 0 \text{ for some } t | X_0(0) = 0, X_1(0) = 1),\end{aligned}$$

as the extinction probabilities of clones started by a single cell of each type. Note that if we start the process with a single type-0 cell, the initial event in the process will be a (i) cell division with probability $r_0/(r_0 + d_0 + \mu)$, (ii) cell death with probability $d_0/(r_0 + d_0 + \mu)$, and a switch to type-1 with probability $\mu/(r_0 + d_0 + \mu)$. If the initial event is a cell division, an additional type-0 cell will be created and the process goes extinct if and only if both cells go extinct, which occurs with probability p_0^2. If the initial

event is a switch between types, the new type-1 cell will go extinct with probability p_1. Therefore, by conditioning on whether the initial event is a division, death or switch between types, we can derive the following conditions for p_0 and p_1:

$$p_0 = \frac{r_0}{r_0 + d_0 + \mu} \cdot p_0^2 + \frac{d_0}{r_0 + d_0 + \mu} \cdot 1 + \frac{\mu}{r_0 + d_0 + \mu} \cdot p_1,$$

which yields

$$(r_0 + d_0 + \mu) \cdot p_0 = r_0 p_0^2 + d_0 + \mu p_1.$$

By the same argument for a process started with a single type-1 cell, we obtain the following nonlinear system for p_0 and p_1:

$$\begin{aligned} (r_0 + d_0 + \mu) \cdot p_0 &= r_0 p_0^2 + d_0 + \mu p_1, \\ (r_1 + d_1 + \nu) \cdot p_1 &= r_1 p_1^2 + d_1 + \nu p_0. \end{aligned} \tag{C.9}$$

When $\sigma > 0$ and the process is supercritical, we have $p_0 < 1$ and $p_1 < 1$ which are uniquely determined by these two equations, while when $\sigma < 0$, we have $p_0 = 1$ and $p_1 = 1$ [6].

Note that the classification into supercritical and subcritical is independent of the initial conditions $(X_0(0), X_1(0)) = (n, m)$ (a change in initial conditions does not affect the extinction probability of a single-cell derived clone). The initial conditions do however affect the survival probability of the population as a whole. Indeed, the survival probability for a population starting with n type-0 cells and m type-1 cells is

$$1 - p_0^n p_1^m, \tag{C.10}$$

since the population goes extinct if and only if each individual type-0 cell and each individual type-1 cell goes extinct, which occurs with probability $p_0^n p_1^m$ by independence.

C.4 Conditions for supercriticality

Here, we derive condition (4.7) in the main text. We again let $(X_0(t), X_1(t))$ denote the stochastic process corresponding to the model in the absence of permanent resistance mechanisms ($\eta = 0$ and $\xi_0 = \xi_1 = 0$).

Lemma C.1. *Assume $\lambda_0 < 0$, $\lambda_1 > 0$, $\mu > 0$ and $\nu > 0$.*

(1) $(X_0(t), X_1(t))$ is subcritical if $\nu\lambda_0 + \mu\lambda_1 < \lambda_0\lambda_1$, critical if $\nu\lambda_0 + \mu\lambda_1 = \lambda_0\lambda_1$ and supercritical if $\nu\lambda_0 + \mu\lambda_1 > \lambda_0\lambda_1$.

(2) A sufficient condition for supercriticality is $\lambda_1 \geq \nu$.

Proof. It is trivial to check, by determining the sign of the larger eigenvalue σ in (4.4) in the main text, that

(i) $(X_0(t), X_1(t))$ is subcritical if $\lambda_0 + \lambda_1 < \mu + \nu$ and $\nu\lambda_0 + \mu\lambda_1 < \lambda_0\lambda_1$.

(ii) $(X_0(t), X_1(t))$ is critical if $\lambda_0 + \lambda_1 < \mu + \nu$ and $\nu\lambda_0 + \mu\lambda_1 = \lambda_0\lambda_1$.

(iii) $(X_0(t), X_1(t))$ is supercritical if $\lambda_0 + \lambda_1 \geq \mu + \nu$ or $\nu\lambda_0 + \mu\lambda_1 > \lambda_0\lambda_1$.

What remains to show is that the first condition is redundant in each case, and that (iii) is satisfied when $\lambda_1 \geq \nu$.

Note first that if $\lambda_1 \geq \nu$, we have

$$\nu\lambda_0 + \mu\lambda_1 - \lambda_0\lambda_1 = \lambda_0(\nu - \lambda_1) + \mu\lambda_1 \geq \mu\lambda_1 > 0.$$

since $\lambda_0 < 0$, $\mu > 0$ and $\lambda_1 > 0$, from which supercriticality follows. This establishes (2). Then note that since $\lambda_1 \geq \nu$ implies $\nu\lambda_0 + \mu\lambda_1 > \lambda_0\lambda_1$, and we have $\nu - \lambda_0 > 0$, we get

$$\nu\lambda_0 + \mu\lambda_1 \leq \lambda_0\lambda_1 \quad \Rightarrow \quad \lambda_1 < \nu \quad \Rightarrow \quad \lambda_1 < \mu + \nu - \lambda_0,$$

which shows that the first condition in (i) and (ii) above is implied by the second condition. Using what we have proven, we additionally have

$$\lambda_0 + \lambda_1 \geq \mu + \nu \quad \Rightarrow \quad \lambda_1 \geq \nu \quad \Rightarrow \quad \nu\lambda_0 + \mu\lambda_1 - \lambda_0\lambda_1 > 0,$$

showing that the second condition in (iii) above contains the first one. □

C.5 Alternate version of the model

In our baseline model, we assume that epimutations and reversions can occur at any time during the cell cycle. If we instead want a model where these events only occur at cell divisions, we proceed as follows, assuming again the absence of permanent resistance mechanisms ($\eta = 0$ and $\xi_0 = \xi_1 = 0$).

Let u denote the probability of epimutation at each division of a type-0 cell and v denote the probability of reversion at each division of a type-1 cell. A type-0 cell then divides into two type-0 cells at rate $r_0(1-u)$ and it produces one type-0 and one type-1 cell at rate $r_0 u$. Similarly, a type-1 cell divides into two type-1 cells at rate $r_1(1-v)$ and it produces one type-0 and one type-1 cell at rate $r_1 v$. These dynamics can be captured by the infinitesimal generator

$$\mathbf{A} = \begin{bmatrix} r_0(1-u) - d_0 & r_0 u \\ r_1 v & r_1(1-v) - d_1 \end{bmatrix} = \begin{bmatrix} \lambda_0 - r_0 u & r_0 u \\ r_1 v & \lambda_1 - r_1 v \end{bmatrix}.$$

Note that this infinitesimal generator has the same form as the generator (E.4) for our baseline model if we define $\mu := r_0 u$ and $\nu := r_1 v$. This implies that the mean behavior of this new model is identical to the mean behavior of our original model, and all expressions of Appendix C.2 apply by making the above substitution. The models are distinct, however, on a sample-path basis, which is reflected e.g. in distinct extinction probabiltities. These probabilities are now determined by the equations

$$\begin{aligned} (r_0 + d_0) \cdot p_0 &= r_0(1-u)p_0^2 + d_0 + r_0 u p_0 p_1, \\ (r_1 + d_1) \cdot p_1 &= r_1(1-v)p_1^2 + d_1 + r_1 v p_0 p_1, \end{aligned}$$

which in general yield different solutions to (C.9).

C.6 Threshold value for μ

In the main text, we examine the survival probability of a tumor under continuous anti-cancer therapy, assuming no transiently resistant cell is present at the start of treatment ($m = 0$). We observe a threshold value μ' which is the minimal epimutation rate that

guarantees long-term survival of the resistant population. Here, we are interested in deriving an expression for this threshold value.

Let $Y_0(t)$ denote the number of type-0 cells still alive t time units into anti-cancer treatment, and let $Y_1(t)$ denote the number of type-1 cells that have been produced through type-0 epimutations up until time t. Since we are only interested in what happens during the initial stages of treatment, and the type-1 population is assumed to be small compared to the initial type-0 population, we ignore reversions from type-1 to type-0 and thus assume that the type-0 population decays at an exponential rate $|\lambda_0|$, with initial population size $Y_0(0) = n$.

Note that given $Y_0(s)$ for $0 \leq s \leq t$, the type-0 population produces a type-1 cell at rate $\mu Y_0(s)$ at time s. Therefore,

$$E[Y_1(t)|Y_0(s), s \leq t] = \int_0^t \mu Y_0(s) ds,$$

from which we conclude

$$E[Y_1(t)] = \int_0^t \mu n e^{-|\lambda_0|s} ds = \frac{n\mu}{|\lambda_0|}(1 - e^{-|\lambda_0|t}).$$

Note that $E[Y_1(t)] \to n\mu/|\lambda_0|$ as $t \to \infty$, so approximately

$$n\mu/|\lambda_0| \tag{C.11}$$

type-1 cells are created during the initial stages of treatment, during which most of the type-0 population is eradicated.

Now let $u \ll 1$. The number of type-1 cells, N_1, needed to guarantee that the resistant population survives with probability $1 - u$ satisfies

$$1 - (d_1/r_1)^{N_1} \geq 1 - u,$$

where we use that d_1/r_1 is the extinction probability of a type-1 clone, assuming no

reversions to type-0 [50]. We therefore conclude that

$$N_1 \geq \frac{\log u}{\log(d_1/r_1)} \tag{C.12}$$

(recall that $r_1 > d_1$ by assumption). To determine the threshold epimutation rate μ' that guarantees survival probability of at least $1-u$ under treatment, we combine (C.11) and (C.12) to get

$$\mu' \approx \frac{|\lambda_0| \log u}{n \log(d_1/r_1)}. \tag{C.13}$$

C.7 Time to stable resistance

We are now interested in analyzing the time at which the first successful type-2 cell emerges in the population during continuous anti-cancer therapy. Let $Z_0(s)$ and $Z_1(s)$ denote the number of type-0 and type-1 cells alive at time s, now assuming the presence of permanent resistance mechanisms ($\eta > 0$ and $\xi_0, \xi_1 > 0$). Assume the initial conditions $(Z_0(0), Z_1(0)) = (n, m)$ and that no permanently resistant (type-2) cell is present at the onset. Let η^{eff} be the 'effective' reprogramming rate, i.e. the rate at which epigenetic reprogramming produces a successful type-2 cell (a type-2 cell that gives rise to a clone that does not go extinct), and let ξ_0^{eff} and ξ_1^{eff} be the 'effective' mutation rates. Then $\eta^{\text{eff}} = \eta\lambda_2/r_2$, $\xi_0^{\text{eff}} = \xi_0\lambda_2/r_2$ and $\xi_1^{\text{eff}} = \xi_1\lambda_2/r_2$, since type-2 cells form a single-type binary branching process with extinction probability $d_2/r_2 = 1 - \lambda_2/r_2$ [50].

Let τ denote the time of first occurrence of a successful type-2 cell. Note that at any time s, the total rate at which cells acquire a resistance-conferring mutation is $r_0\xi_0^{\text{eff}}Z_0(s) + r_1\xi_1^{\text{eff}}Z_1(s)$, and the total rate of epigenetic reprogramming is $\eta^{\text{eff}}Z_1(s)$. If we condition on the history of the type-0 and type-1 population up until some time t, the number of type-2 cells that have emerged by time t is then Poisson distributed with mean

$$\int_0^t (\xi_0^{\text{eff}} r_0 Z_0(s) + \xi_1^{\text{eff}} r_1 Z_1(s) + \eta^{\text{eff}} Z_1(s))ds,$$

which implies that

$$\begin{aligned}
&P(\tau > t \,|\, (Z_0(s), Z_1(s))_{s \le t}) \\
&= \exp\left(- \int_0^t (\xi_0^{\text{eff}} r_0 Z_0(s) + \xi_1^{\text{eff}} r_1 Z_1(s) + \eta^{\text{eff}} Z_1(s)) ds \right),
\end{aligned}$$

since $\{\tau > t\}$ is the event that no type-2 cell has emerged by time t. Taking expectations, we arrive at

$$\begin{aligned}
&P(\tau > t) \\
&= E\left[\exp\left(- \int_0^t (\xi_0^{\text{eff}} r_0 Z_0(s) + \xi_1^{\text{eff}} r_1 Z_1(s) + \eta^{\text{eff}} Z_1(s)) ds \right) \right].
\end{aligned}$$

We now argue that $Z_0(s)$ and $Z_1(s)$ in the exponent can be well-approximated by their means. Assuming the initial condition $(Z_0(0), Z_1(0)) = (n, m)$, we can write

$$(Z_0(s), Z_1(s)) = \sum_{j=1}^{n} \left(Z_{0,j}^{(0)}(s), Z_{1,j}^{(0)}(s) \right) + \sum_{j=1}^{m} \left(Z_{0,j}^{(1)}(s), Z_{1,j}^{(1)}(s) \right),$$

where $Z_{i,j}^{(k)}(t)$ represents the number of type-i cells present at time t descended from a single type-k cell and $\{Z_{i,j}^{(k)}(t) : j \in \{1, 2, \ldots\}\}$ is an i.i.d. sequence of such random variables. Given sufficiently large n and m, we know by the law of large numbers that the approximation $(Z_0(s), Z_1(s)) \approx (\phi_0(s), \phi_1(s))$ should be justified. To examine fluctuations around the mean, we note that by the central limit theorem,

$$\begin{aligned}
\sum_{j=1}^{n} \left(Z_{0,j}^{(0)}(s), Z_{1,j}^{(0)}(s) \right) &\approx n \left(E[Z_{0,1}^{(0)}(s)], E[Z_{1,1}^{(0)}(s)] \right) \\
&\quad + \sqrt{n} \left(W_1 \sqrt{\text{Var}\big(Z_{0,1}^{(0)}(s)\big)}, W_2 \sqrt{\text{Var}\big(Z_{1,1}^{(0)}(s)\big)} \right),
\end{aligned}$$

where (W_1, W_2) is a mean-zero bivariate Gaussian. By equation (21) of Section V.7 of [6] we know that there exists a $C > 0$ such that

$$\sqrt{\text{Var}\big(Z_{i,1}^{(0)}(s)\big)} \le C E[Z_{i,1}^{(0)}(s)], \quad s > 0.$$

We can therefore write

$$\sum_{j=1}^{n}\left(Z_{0,j}^{(0)}(s), Z_{1,j}^{(0)}(s)\right) \approx n\left(E[Z_{0,1}^{(0)}(s)], E[Z_{1,1}^{(0)}(s)]\right)(1+C/\sqrt{n}).$$

A similar result holds for the descendants of the type-1 cells. Ignoring the second-order term, we arrive at

$$\begin{aligned} &P(\tau > t) \\ &\approx \exp\left(-\int_0^t (\xi_0^{\text{eff}} r_0 \phi_0(s) + \xi_1^{\text{eff}} r_1 \phi_1(s) + \eta^{\text{eff}} \phi_1(s)) ds\right). \end{aligned} \tag{C.14}$$

To test the quality of this approximation, we show in Figure C.1 a comparison between (C.14) and simulation results for as few as $n = 10^4$ starting cells. To calculate the mean number of type-0 and type-1 cells alive at time s, $\phi_0(s)$ and $\phi_1(s)$ in (C.14), we can apply (4.3) in the main text. Technically, the birth rates r_0 and r_1 should be replaced by $r_0(1-\xi_0^{\text{eff}})$ and $r_1(1-\xi_1^{\text{eff}})$, and the death rate d_1 should be replaced by $d_1+\eta^{\text{eff}}$, to reflect the introduction of permanent resistance mechanisms, but assuming $\xi \ll 1$ and $\eta \ll d_1$, we can safely ignore this minor complication. Using (4.3), we can therefore calculate

$$\begin{aligned} \int_0^t \phi_0(s) ds &= -\frac{\beta}{\sigma}(1-e^{\sigma t}) - \frac{\alpha}{\rho}(1-e^{\rho t}), \\ \int_0^t \phi_1(s) ds &= -\frac{\gamma\beta}{\sigma}(1-e^{\sigma t}) + \frac{\delta\alpha}{\rho}(1-e^{\rho t}), \end{aligned} \tag{C.15}$$

which gives an explicit expression for (C.14) assuming $\sigma \neq 0$ and $\rho \neq 0$. Note that the expressions in (C.15) can be interpreted as the 'total mass' of type-0 and type-1 cells, respectively, up until time t.

In the main text, we apply (C.14) and (C.15) to evaluate interval treatment strategies, where the anti-cancer agent is intermittently applied and removed. This can create a problem, since our calculations of η^{eff}, ξ_0^{eff} and ξ_1^{eff} assume that the rates r_2 and d_2 of birth and death of type-2 cells do not change over time. However, since we make the assumption in our baseline parameter regime that r_2 and d_2 are identical on and

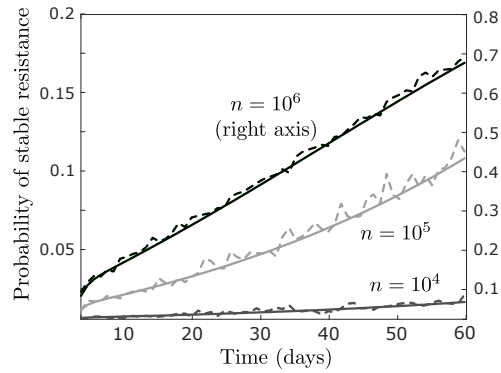

Figure C.1: Comparison of expression (C.14) (whole curves) with simulation results (dotted curves) in the baseline parameter regime, assuming a starting population of $n = 10^4$ (bottom curve; left axis), $n = 10^5$ (middle curve; left axis) or $n = 10^6$ cells (top curve; right axis). Simulation results are based on 2000 runs of the process. Other parameter values are $r_0 = 0.04$ (per hour), $d_0 = 0.08$, $r_1 = 0.0162$, $d_1 = 0.015$, $\mu = 4 \cdot 10^{-5}$, $\nu = 4 \cdot 10^{-4}$, $\eta = 4 \cdot 10^{-7}$, $\xi_0 = \xi_1 = 10^{-7}$ (per cell division), and $m = 0$.

off the anti-cancer agent, we do not have to make any adjustments to (C.14) to use it to evaluate interval treatment strategies. In fact, since we assume $r_2 = 0.04$ and $d_2 = 0.0015$ with $\lambda_2/d_2 \approx 1$, we can effectively take $\eta^{\text{eff}} = \eta$, $\xi_0^{\text{eff}} = \xi_0$ and $\xi_1^{\text{eff}} = \xi_1$.

To determine whether stable resistance develops at all during continuous anti-cancer therapy, we consider the event $\{\tau = \infty\}$. It is possible to estimate its probability by taking $t \to \infty$ in (C.14), but then we first have to condition on whether the overall population of type-0 and type-1 cells goes extinct or not, since $\phi_0(s)$ and $\phi_1(s)$ blow up as $s \to \infty$ for any supercritical process, irrespective of the likelihood that the population survives. However, it is more straightforward to proceed as follows. Let p_0, p_1 and p_2 denote the extinction probabilities of clones started by a single cell of each type. By conditioning on the initial event, we derive the following conditions for these probabilities:

$$\begin{aligned}
&(r_0 + d_0 + \mu) \cdot p_0 = r_0(1 - \xi_0)p_0^2 + d_0 + \mu p_1 + r_0\xi_0 p_0 p_2, \\
&(r_1 + d_1 + \nu + \eta) \cdot p_1 = r_1(1 - \xi_1)p_1^2 + d_1 + \nu p_0 + r_1\xi_1 p_1 p_2 + \eta p_2, \\
&p_2 = d_2/r_2,
\end{aligned} \tag{C.16}$$

where we use that fact that type-2 cells form a single-type binary branching process

with extinction probability d_2/r_2. Note that these equations reduce to (C.9) if we set $\xi_0 = \xi_1 = 0$ and $\eta = 0$. We can then calculate

$$P(\tau = \infty) = p_0^n p_1^m, \tag{C.17}$$

which gives the probability that stable resistance does not develop at all during continuous administration of anti-cancer treatment. The probability that stable resistance does develop is then

$$P(\tau < \infty) = 1 - p_0^n p_1^m. \tag{C.18}$$

C.8 Robustness analysis

Survival probability

Table C.3 shows how survival probability, both at the indidividual cell level and for the tumor population as a whole, responds to changes in r_0, d_0, r_1 and d_1, pertinent to the discussion of Section 4.4.1.

Threshold value for μ

Table C.4 shows how the threshold value for μ, denoted as μ' and given by (4.11), responds to to changes in r_1 and d_1, pertinent to the discussion of Section 4.4.2.

Evaluation of combination treatment strategies

We now discuss various possible changes to the assumptions underlying Figure 4.6 in the main text. In this section, we let μ_d and ν_d denote the rates of epimutation and reversion when the epigenetic drug is being applied to distinguish them from the rates μ and ν off the epigenetic drug.

Tumor size and rates of mutation and reprogramming

In the main text, we assume an initial population size of $n = 10^6 \cdot 0.999$ and $m = 10^6 \cdot 0.001$ cells, an epigenetic reprogramming rate of $\eta = 4 \cdot 10^{-7}$ per hour, and a

r_0	d_0	r_1	d_1	$1-p_0$	$1-p_1$	$1-p_0^n p_1^m$
0.04	0.08	0.0162	0.015	0.0049%	4.94%	100%
+10%				0.0055%	4.94%	100%
	+10%			0.0041%	4.94%	100%
+10%	+10%			0.0045%	4.94%	100%
		+10%		0.0136%	13.58%	100%
			+10%	0%	0%	0%
		+10%	+10%	0.0052%	5.16%	100%
+50%				0.0099%	4.94%	100%
	+50%			0.0025%	4.94%	100%
+50%	+50%			0.0033%	4.94%	100%
		+50%		0.0366%	36.63%	100%
			+50%	0%	0%	0%
		+50%	+50%	0.0058%	5.76%	100%

Table C.3: Response of the survival probability of (i) a clone started by a single type-0 cell $(1-p_0)$, (ii) a clone started by a single type-1 cell $(1-p_1)$, and (iii) the overall population of $n = 10^6$ type-0 and $m = 0$ type-1 cells $(1-p_0^n p_1^m)$, to changes in r_0, r_1, d_0 and d_1. The probabilities p_0 and p_1 can be calculated using (C.9). In the top row, we show the baseline parameter values for r_0, d_0, r_1 and d_1, and the survival probabilities given these values. In the remaining lines, we adjust one or two parameters at a time as indicated and display survival probabilities given these changes. All rates are measured per hour.

mutation rate of $\xi_0 = \xi_1 = 10^{-7}$ per cell division. If we assume that the rate of epigenetic reprogramming is increased by an order of magnitude to $\eta = 4 \cdot 10^{-6}$, the relative attractiveness of Schedule B over the other schedules becomes even more pronounced than in the baseline case (Fig C.2a). If we also increase the mutation rate an order of magnitude to $\xi_0 = \xi_1 = 10^{-6}$, it becomes more difficult to eradicate the tumor, but the relative attractiveness of the schedules remains unchanged (Fig C.2b).

If we increase the population size one order of magnitude to $n = 10^7 \cdot 0.999$ and $m = 10^7 \cdot 0.001$, assuming the baseline values $\eta = 4 \cdot 10^{-7}$ and $\xi_0 = \xi_1 = 10^{-7}$ for the rates of reprogramming and mutation (figure not shown), the effects will be similar to those shown in Figure C.2b, where the population size is as in the baseline case but both rate parameters are increased an order of magnitude. Thus, the relationship between

r_1	d_1	μ'
0.0162	0.015	$3.59 \cdot 10^{-6}$
+10%		$1.60 \cdot 10^{-6}$
	+10%	$+\infty$
+10%	+10%	$3.59 \cdot 10^{-6}$
+50%		$0.57 \cdot 10^{-6}$
	+50%	$+\infty$
+50%	+50%	$3.59 \cdot 10^{-6}$

Table C.4: Response of the threshold value μ', as given by (4.11), to changes in r_1 and d_1. We use $+\infty$ to indicate that there is no value for μ for which tumor survival is guaranteed, which is the case when the population is subcritical. All rates are measured per hour.

initial population size and the rate of mutation and epigenetic reprogramming is more important than the absolute value of each in isolation, and any effect of increasing the population size can be negated by a proportional decrease in the these rates. We also note that if the anti-cancer agent being applied reduces the birth rate r_0 of type-0 cells, this will have the same effect as reducing the mutation rate ξ_0 (see discussion in Appendix C.1 above).

For an even larger initial tumor size of $n = 10^8 \cdot 0.999$ and $m = 10^8 \cdot 0.001$, a resistance-conferring mutation will be guaranteed to arise no matter which schedule is applied (figure not shown), although we recover the same behavior as for the $n = 10^7 \cdot 0.999$ and $m = 10^7 \cdot 0.001$ case if the change in size is accompanied by a smaller reprogramming rate, $\eta = 4 \cdot 10^{-8}$, and a smaller mutation rate, $\xi_0 = \xi_1 = 10^{-8}$. However, even if the initial population size is $n = 10^8 \cdot 0.999$ and $m = 10^8 \cdot 0.001$, with unchanged rates of mutation and reprogramming, there can still be benefits to applying a combination of an anti-cancer agent and an epigenetic drug, since Schedule B can prevent resistance conferred through epigenetic reprogramming (Fig C.2c). In fact, if we assume that the epigenetic drug is stronger than in the main text, so that it increases the reversion rate ν by three orders of magnitude and decreases the epimutation rate to the same extent, then Schedule B is able to prevent resistance through epigenetic reprogramming with high probability, whereas resistance is effectively guaranteed under any other schedule (Fig C.2d).

Different epigenetic drug action

In the main text, we assume that the epigenetic drug increases the reversion rate ν by two orders of magnitude, and decreases the epimutation rate μ to the same extent. In Figures C.2e and C.2f, we show results first assuming that the change in each is one order of magnitude (Fig C.2e), and then three orders of magnitude (Fig C.2f). We note that Schedule B continues to be the most attractive schedule, and in Figure C.2e, the difference between Schedule B and Schedules C,D becomes more pronounced than for the case shown in the main text.

Faster or slower switching dynamics

In the main text, we assume an epimutation rate of $\mu = 4 \cdot 10^{-5}$ in the absence of the epigenetic drug, in accordance with our baseline parameter regime. In Figures C.3a and C.3b, we show results assuming $\mu = 4 \cdot 10^{-4}$ (Fig C.3a) and $\mu = 4 \cdot 10^{-6}$ (Fig C.3b), with the epimutation rate μ_d in the presence of the epigenetic drug obtained by reducing μ by two orders of magnitude in each case. As before, the qualitative dynamics remain similar, with the difference between Schedule B and the other schedules even more pronounced in Figure C.3a than for the baseline case.

Different treatment block sizes

In the main text, we assume that each treatment block is 2 days. In Figures C.3c and C.3d, we show results assuming that each block is 1 day (Fig C.3c) or 3 days (Fig C.3d), and observe similar results as before.

Resistant phenotype is not at a selective disadvantage off drug

In our baseline regime, we assume that the resistant type-1 phenotype is at a selective disadvantage to drug-sensitive type-0 cells in the absence of the anti-cancer agent. Figure C.3e reveals that our results do not depend on this assumption, as if we assume $r_1 = r_0 = 0.04$ and $d_1 = d_0 = 0.0015$ off the anti-cancer agent, Schedule B continues to perform the best.

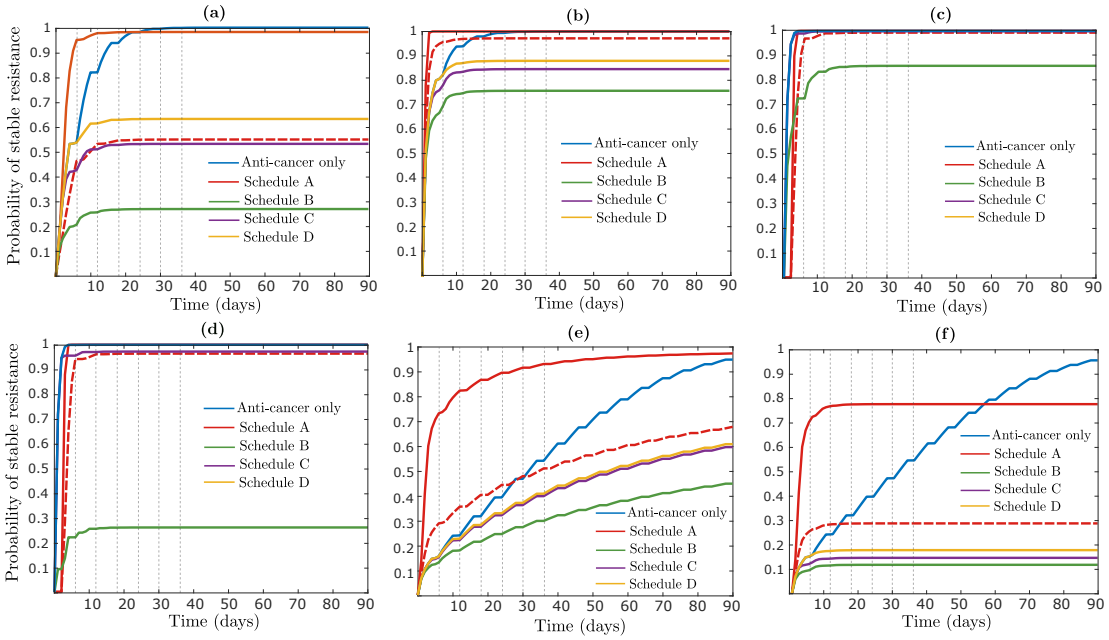

Figure C.2: Time-evolution of the probability of successful type-2 emergence for the treatment schedules examined in Figure 4.6 in the main text. Parameters behave as in Figure 4.6 except in (a), $\eta = 4 \cdot 10^{-6}$ (per hour), in (b), $\eta = 4 \cdot 10^{-6}$ and $\xi_0 = \xi_1 = 10^{-6}$ (per cell division), in (c), $n = 10^8 \cdot 0.999$, $m = 10^8 \cdot 0.001$ and $\xi_0 = \xi_1 = 0$, in (d), $n = 10^8 \cdot 0.999$, $m = 10^8 \cdot 0.001$, $\xi_0 = \xi_1 = 0$, $\mu_d = \mu \cdot 10^{-3}$ and $\nu_d = \nu \cdot 10^3$, in (e), $\mu_d = \mu \cdot 10^{-1}$ and $\nu_d = \nu \cdot 10^1$, and in (f), $\mu_d = \mu \cdot 10^{-3}$ and $\nu_d = \nu \cdot 10^3$.

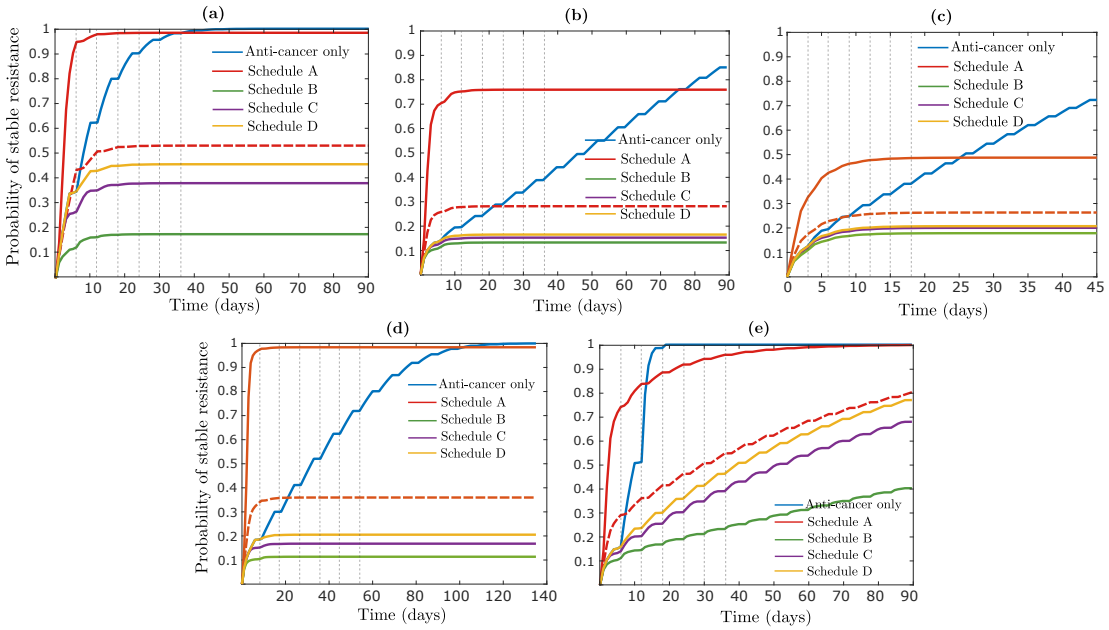

Figure C.3: Time-evolution of the probability of successful type-2 emergence for the treatment schedules examined in Figure 4.6 in the main text. Parameters behave as in Figure 4.6 except in (a), $\mu = 4{\cdot}10^{-4}$ in the absence of epigenetic drug, in (b), $\mu = 4{\cdot}10^{-6}$ in the absence of epigenetic drug, in (c), the size of each treatment block is 1 day as opposed to 2 days, in (d), the size of each block is 3 days, and in (e), $r_1 = 0.04$ (per hour) and $d_1 = 0.0015$ in the absence of the anti-cancer agent.

Appendix D

Appendix to Chapter 5

D.1 Review of existing methods

At the single-cell-level, phenotypic switching has commonly been modeled by a discrete-time Markov chain with $K \geq 2$ states, where K is the number of phenotypes. In each time step, a cell in state j transitions to state $k \neq j$ with probability p_{jk}, and it remains in state j with probability $p_{jj} = 1 - \sum_{k \neq j} p_{jk}$. The transition probabilities are collected into the $K \times K$ *transition matrix* $\mathbf{P} = (p_{jk})$. The evolution of the Markov chain is determined by $\mathbf{P}$ and the *initial distribution* $\mathbf{q} = (q_1, \ldots, q_K)$, where q_j is the probability that a cell starts in state j. If we let $\mathbf{q}^{(\ell)}$ denote the cell state distribution after $\ell \geq 1$ time steps, then $\mathbf{q}^{(\ell)} = \mathbf{q}\mathbf{P}^{\ell}$.

Say we conduct K cell line experiments starting with N cells in each experiment and known initial cell state distributions $\mathbf{q}_1, \ldots, \mathbf{q}_K$. The initial distributions are collected into a $K \times K$ matrix $\mathbf{Q}$, where $\mathbf{q}_i$ is the i-th row vector. Each experiment is run for $\ell \geq 1$ time steps, at which point the fraction of cells in each state is recorded. Let $f_{ij}^{(\ell)}$ be the observed fraction of cells in state j under the i-th initial condition. The observations at the ℓ-th time step under the i-th initial condition are collected into a vector $\mathbf{f}_i^{(\ell)} = \left(f_{i1}^{(\ell)}, \ldots, f_{iK}^{(\ell)}\right)$, and the observations at the ℓ-th time step are collected into a $K \times K$ matrix $\mathbf{F}^{(\ell)} = \left(f_{ij}^{(\ell)}\right)$. If there are multiple replicates $r = 1, \ldots, R$, we let $\mathbf{F}^{(\ell),r}$ denote the data from the r-th replicate.

Now, assume that the starting population N is large, that there is no cell division or cell death, and that each cell switches between states according to the above Markov

model. In this case, by the law of large numbers, the model-predicted distribution between cell states $\mathbf{QP}^\ell$ after ℓ time steps can be approximated by the experimentally observed cell-state fractions $\mathbf{F}^{(\ell)}$. If we equate these two matrices, we can obtain an estimate $\mathbf{P}_\ell$ of the transition matrix $\mathbf{P}$ by inverting the matrix $\mathbf{Q}$ of initial distributions and taking an ℓ-th matrix root, $\mathbf{P}_\ell = \left(\mathbf{Q}^{-1}\mathbf{F}^{(\ell)}\right)^{1/\ell}$. Here, we assume that $\mathbf{Q}$ is invertible, which is e.g. the case when experiments are started with isolated subpopulations.

This simple estimation idea was applied by Gupta et al. [77] to investigate phenotypic switching between stem-like, basal and luminal cell states in breast cancer, using data from a single time point. A multiple-time-point version has since been implemented in the R package CellTrans [24]. Say that cell state fractions are experimentally observed at time steps $m_1, \ldots, m_L$ for $L \geq 1$. CellTrans first computes an estimate $\mathbf{P}_{m_\ell}$ of the transition matrix for each time step as above, and then returns a final estimate as the average across time steps:

$$\widehat{\mathbf{P}} := (1/L) \textstyle\sum_{\ell=1}^{L} \mathbf{P}_{m_\ell}. \tag{D.1}$$

CellTrans also involves a regularization step to ensure that $\widehat{\mathbf{P}}$ is stochastic. CellTrans is used on publicly available datasets in [24] and it has been applied more recently in [34, 45, 153].

Cell populations in culture typically change in size over time. CellTrans is applied under the assumption that all phenotypes grow at the same rate, and that the evolution of cell state fractions can be captured by a constant-sized population model for that case. Both Gupta et al. [77] and Su et al. [144] have applied an augmented version of the above model, which is intended to capture proliferation differences between the phenotypes, specifically in the context of anti-cancer treatment response. In the augmented model, a type-j cell first grows deterministically to a population of size Λ_{jj}, and a fraction p_{jk} of the population then switches to type-k. The growth factors Λ_{jj} are collected into a diagonal proliferation matrix $\mathbf{\Lambda}$, and the multiple $\mathbf{\Lambda P}$ is used to predict the distribution between cell states as opposed to $\mathbf{P}$. In Gupta et al. [77] and Su et al. [144], the matrix $\mathbf{\Lambda}$ is found by randomly sampling candidate parameter values and selecting the values that best fit the experimental data.

TRANSCOMPP [82] is a more systematic version of the estimation method from

Gupta et al. [77] and Su et al. [144]. In TRANSCOMPP, the diagonal proliferation matrix $\mathbf{\Lambda}$ and the transition matrix $\mathbf{P}$ are estimated by minimizing the sum of squared errors between the model prediction and the data,

$$\min_{\mathbf{\Lambda},\mathbf{P}} \sum_{i=1}^{I} \sum_{\ell=1}^{L} \sum_{r=1}^{R} \left\| \mathbf{f}_i^{(m_\ell),r} - \operatorname{diag}\big(\mathbf{q}_i \big(\mathbf{\Lambda}\mathbf{P}\big)^{m_\ell} \mathbf{1}^T\big)^{-1} \mathbf{q}_i \big(\mathbf{\Lambda}\mathbf{P}\big)^{m_\ell} \right\|^2. \tag{D.2}$$

In the TRANSCOMPP paper [82], the normalizing factor $\operatorname{diag}\big(\mathbf{q}_i \big(\mathbf{\Lambda}\mathbf{P}\big)^{m_\ell} \mathbf{1}^T\big)^{-1}$ is not shown in the statement (4) of their optimization problem, but it is indicated in the surrounding text. TRANSCOMPP is implemented in MATLAB, and it includes a stochastic resampling procedure for estimating the distributions of the transition probability estimates.

All of the above estimation methods are based on discrete-time models. In modeling switching between HER2+ and HER2− states in breast cancer, Li and Thirumalai [105] employ a deterministic model in continuous time. Their model assumes symmetric and asymmetric cell divisions, which is mathematically equivalent to assuming symmetric cell divisions and switching between types. Li and Thirumalai assume equal rates of asymmetric division for the two cell types, and they show that if experiments are started with isolated subpopulations, the slopes of the cell fraction trajectories at time 0 can be used to estimate these rates. They also show that the equilibrium proportion between types can be used to extract the symmetric division rate difference between types. Li and Thirumalai use the proportion between cells in the parental population as an estimate of the equilibrium proportion. We have made use of these insights in our identifiability analysis in Section 5.6.2 in the main text.

Finally, in their investigation of epithelial to mesenchymal transition in breast cancer, Devaraj and Bose [42] employ a similar model to [105] with two important differences. First of all, they assume that the rates of switching and growth are time-dependent. Second, they include a separate state for dead cells and estimate a death rate for each cell type. In their experiments, Devaraj and Bose measure both the phenotypic proportions and the number of alive and dead cells in the population at each time point. On the modeling side, they derive difference equations for the change in the number of cells in each state between time points. They then propose a multi-objective optimization problem to estimate their model parameters, which minimizes

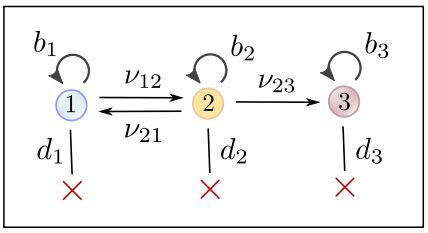

Figure D.1: To demonstrate that our estimation framework is applicable to reducible switching models, we consider a three-type model with a reversible transition between type-1 and type-2, and an irreversible transition from type-2 to type-3. This model is applicable e.g. to epigenetic gene silencing under the recruitment of chromatin regulators [11] and to epigenetically-driven drug resistance in cancer [75].

the least squares error between the model predictions and the data, while ensuring that parameters do not vary too drastically between time periods.

D.2 Estimation for reducible switching dynamics

In the main text, we have assumed that the switching dynamics are irreducible, meaning that it is possible to switch between any pair of phenotypes, possibly through a finite number of intermediate types. In this section, we show how our framework can be applied to the case of reducible switching dynamics. For simplicity, we will consider one particular model shown in Figure D.1. This model has been applied e.g. to the dynamics of epigenetic gene silencing under recruitment of chromatin regulators [11] and the evolution of epigenetically-driven drug resistance in cancer, where drug-sensitive cells (type-1) first acquire a transiently resistant phenotype (type-2) and then evolve to stable epigenetic resistance (type-3) [75].

Say that experiments are conducted from isolated initial conditions, and say first that cell number data is collected. For the model in Figure D.1, the distribution of the data vector $\mathbf{n}_{3,\ell,r}$ is degenerate, since $n_{3,\ell,r,j} = 0$ for $j = 1, 2$. As a result, the covariance matrix $\mathbf{\Sigma}^{(3)}(t_\ell)$ is singular for all $\ell = 1, \ldots, L$, and the likelihood function in (5.6) is not defined. To resolve this issue, we set $\mathbf{C}_1 = \mathbf{C}_2 = \mathbf{I}$ and $\mathbf{C}_3 = \mathbf{e}_3^T$, where $\mathbf{e}_3$ is the $1 \times K$ third unit vector. By Proposition 5.1, $\mathbf{n}_{3,\ell,r}\mathbf{C}_3 = n_{3,\ell,r,3}$ has a normal distribution,

which is nondegenerate. We therefore modify the likelihood function in (5.6) to

$$\begin{aligned}
&\mathcal{L}_{\text{num}}\big(\boldsymbol{\theta}_{\text{num}}\big|(\mathbf{n}_{i,\ell,r})_{i,\ell,r}\big)\\
&= \textstyle\prod_{i=1}^{3}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^K\det\big(\mathbf{C}_i^T\big(N_i\boldsymbol{\Sigma}^{(i)}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{num}}\big)\mathbf{C}_i\big)\Big)^{-1/2}\\
&\quad\cdot\exp\big(-\tfrac{1}{2}\big(\mathbf{n}_{i,\ell,r}-N_i\mathbf{m}^{(i)}(t_\ell)\big)\mathbf{C}_i\big(\mathbf{C}_i^T\big(N_i\boldsymbol{\Sigma}^{(i)}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{num}}\big)\mathbf{C}_i\big)^{-1}\\
&\qquad\qquad\mathbf{C}_i^T\big(\mathbf{n}_{i,\ell,r}-N_i\mathbf{m}^{(i)}(t_\ell)\big)^T\big).
\end{aligned}$$

From this likelihood function, MLE estimates and confidence intervals can be computed as described in Section 5.4.3, where we restrict the set of feasible parameters $\boldsymbol{\Theta}_{\text{num}}$ so that $\nu_{13}=\nu_{31}=\nu_{32}=0$. By our analysis in Section 5.6.1, all model parameters are structurally identifiable for this example. To accommodate model structures such as the one discussed here, the above modified likelihood function is implemented in our MATLAB codes (Appendix D.3). By taking $\mathbf{C}_i=\mathbf{I}$ for each $i=1,\ldots,I$, we recover the original likelihood function in (5.6).

If cell fraction data is collected, there is no value in conducting experiments starting only from type-3 cells. We therefore use the likelihood function

$$\begin{aligned}
&\mathcal{L}_{\text{frac}}\big(\boldsymbol{\theta}_{\text{frac}}\big|(\mathbf{f}_{i,\ell,r})_{i,\ell,r}\big)\\
&= \textstyle\prod_{i=1}^{2}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^{K-1}\det\big(\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{(i)}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\mathbf{B}\big)\Big)^{-1/2}\\
&\quad\cdot\exp\Big(-\tfrac{1}{2}\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{(i)}(t_\ell)\big)\mathbf{B}\big(\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{(i)}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\,\mathbf{B}\big)^{-1}\mathbf{B}^T\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{(i)}(t_\ell)\big)^T\Big),
\end{aligned}$$

where we only include experiments started by type-1 and type-2 cells, respectively. By our analysis in Section 5.6.2, the switching rates ν_{12}, ν_{21} and ν_{23}, and the net birth rate differences $\lambda_2-\lambda_1$ and $\lambda_3-\lambda_2$, are structurally identifiable in this case. An example of a model structure where it becomes necessary to modify the above likelihood function is given in Appendix D.3.

D.3 Implementation in MATLAB

In this section, we discuss how our estimation framework is implemented in MATLAB.

D.3.1 Cell number data

The first step in the implementation for cell number data is to compute simple parameter estimates for the switching rates $\boldsymbol{\nu}$ and the net birth rates $\boldsymbol{\lambda}$ based on a deterministic population model. This model is obtained by ignoring the stochastic terms in the statistical model (5.5), i.e. by equating the data vector $\mathbf{n}_{i,\ell,r}$ with the mean prediction of (5.5):

$$\mathbf{n}_{i,\ell,r} = N_i \mathbf{m}^{\mathbf{f}_i}(t_\ell) = \mathbf{n}_i \mathbf{M}(t_\ell). \tag{D.3}$$

Let $\mathbf{N}$ be the $I \times K$ matrix with the initial conditions $\mathbf{n}_i$ as row vectors, and let $\mathbf{N}_{\ell,r}$ be the $I \times K$ matrix with the data vectors $\mathbf{n}_{i,\ell,r}$ as row vectors. We can then write (D.3) in matrix form as

$$\mathbf{N}_{\ell,r} = \mathbf{N}\mathbf{M}(t_\ell) = \mathbf{N} \exp(t_\ell \mathbf{A}). \tag{D.4}$$

Assuming $\mathbf{N}$ has rank K, we can solve for $\mathbf{A}$ in (D.4) by first multiplying both sides by $\mathbf{N}^T$, then multiplying both sides by the inverse of $\mathbf{N}^T\mathbf{N}$, and finally taking a matrix logarithm. We can thus obtain an estimate for the infinitesimal generator $\mathbf{A}$,

$$\mathbf{A}^*_{\ell,r} := (1/t_\ell) \log\left(\left(\mathbf{N}^T\mathbf{N}\right)^{-1}\mathbf{N}^T\,\mathbf{N}_{\ell,r}\right).$$

We then compute a final estimate $\mathbf{A}^*$ by averaging across time points and replicates:

$$\mathbf{A}^* := (1/(LR)) \textstyle\sum_{\ell=1}^{L} \sum_{r=1}^{R} \mathbf{A}^*_{\ell,r}. \tag{D.5}$$

From $\mathbf{A}^*$, we can obtain estimates of the switching rates $\boldsymbol{\nu}$ and the net birth rates $\boldsymbol{\lambda}$.

As indicated in Appendix D.2, we implement the following likelihood function in our

codes:

$$\begin{aligned}
&\mathcal{L}_{\text{num}}\big(\boldsymbol{\theta}_{\text{num}}\big|(\mathbf{n}_{i,\ell,r})_{i,\ell,r}\big)\\
&= \textstyle\prod_{i=1}^{I}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^K\det\big(\mathbf{C}_i^T\big(N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{num}}\big)\mathbf{C}_i\big)\Big)^{-1/2}\\
&\quad\cdot\exp\big(-\tfrac{1}{2}\big(\mathbf{n}_{i,\ell,r}-N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)\big)\mathbf{C}_i\big(\mathbf{C}_i^T\big(N_i\boldsymbol{\Sigma}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{num}}\big)\mathbf{C}_i\big)^{-1}\\
&\qquad\qquad\mathbf{C}_i^T\big(\mathbf{n}_{i,\ell,r}-N_i\mathbf{m}^{\mathbf{f}_i}(t_\ell)\big)^T\big).
\end{aligned}$$

For each $i = 1, \ldots, I$, $\mathbf{C}_i$ is a $K \times J_i$ matrix for some $1 \leq J_i \leq K$, which can be used to reduce the dimension of the data vector $\mathbf{n}_{i,\ell,r}$ when necessary. This option can e.g. be useful for models with reducible switching dynamics, see Appendix D.2.

From the above likelihood function, we compute a negative double log-likelihood as in (5.7), and solve the MLE problem (5.8) using the sequential quadratic programming (sqp) solver in MATLAB. For the optimization, one must supply an initial guess $\boldsymbol{\theta}_{\text{num}}^{(0)}$ for the parameter vector $\boldsymbol{\theta}_{\text{num}}$, and a set of feasible parameters $\boldsymbol{\Theta}_{\text{num}}$ of the form

$$\boldsymbol{\Theta}_{\text{num}} = \{\boldsymbol{\theta}_{\text{num}} : \mathbf{l} \leq \boldsymbol{\theta}_{\text{num}} \leq \mathbf{u}, \mathbf{G}\,\boldsymbol{\theta}_{\text{num}} \leq \mathbf{h}, \mathbf{G}_{\text{eq}}\,\boldsymbol{\theta}_{\text{num}} = \mathbf{h}_{\text{eq}}\}.$$

By default, we assume lower bounds of $\mathbf{0}$ for the switching rates $\boldsymbol{\nu}$ and the birth rates $\mathbf{b}$, and we impose the inequality constraint $\boldsymbol{\lambda} \leq \mathbf{b}$. The user is expected to provide lower bounds for the net birth rates $\boldsymbol{\lambda}$ and upper bounds for all parameters, and they have the option to provide further inequality or equality constraints as necessary. This provides the opportunity to impose constraints such as $\lambda_1 = \lambda_2$ (Section 5.8) or $\nu_{13} = \nu_{31} = \nu_{32} = 0$ (Appendix D.2).

For the initial guess $\boldsymbol{\theta}_{\text{num}}^{(0)}$, we use the simple estimates for $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}$ computed from (D.5). An initial guess for the birth rate b_i is generated as $|\lambda_i|/U$, where U is uniformly distributed between 0 and 1. The idea is that if $\lambda_i > 0$, then in the absence of phenotypic switching, the survival probability of a single-cell derived clone of type i is $q_i = \lambda_i/b_i$ [50]. Since we do not assume any information on q_i, we sample it uniformly between 0 and 1, and then use the initial guess for λ_i to compute an initial guess for b_i. If $|\lambda_i|/U$ falls outside the lower or upper bound provided for b_i, the initial guess is moved to one of the bounds.

In addition to being used to initialize the optimization, the initial guess $\boldsymbol{\theta}_{\text{num}}^{(0)}$ is used

to estimate the relative scales of the parameters $\boldsymbol{\nu}$, $\boldsymbol{\lambda}$ and $\mathbf{b}$. In particular, for the i-th coordinate of the initial guess, we define the corresponding scale variable

$$s_i^{(0)} := 10^{\lfloor \log_{10} |\theta_{\text{num},i}^{(0)}| \rfloor},$$

with $s_i^{(0)} := 1$ if $\theta_{\text{num},i}^{(0)} = 0$. For example, if the initial guesses are $\mathbf{b}^{(0)} = (1.5, 1.2)$ for the birth rates, $\boldsymbol{\lambda}^{(0)} = (0.3, 0.4)$ for the net birth rates, and $\boldsymbol{\nu}^{(0)} = (0.05, 0.002)$ for the switching rates, the corresponding scale variables are $(1, 1)$, $(0.1, 0.1)$ and $(0.01, 0.001)$, respectively. For a given parameter vector $\boldsymbol{\theta}_{\text{num}}$, we define the transformed vector

$$\widetilde{\boldsymbol{\theta}}_{\text{num}} := \boldsymbol{\theta}_{\text{num}} \oslash \mathbf{s}^{(0)},$$

where $\oslash$ denotes elementwise division. For the initial guesses $\mathbf{b}^{(0)} = (1.5, 1.2)$, $\boldsymbol{\lambda}^{(0)} = (0.3, 0.4)$ and $\boldsymbol{\nu}^{(0)} = (0.05, 0.002)$, the corresponding transformed values are $\widetilde{\mathbf{b}}^{(0)} = (1.5, 1.2)$, $\widetilde{\boldsymbol{\lambda}}^{(0)} = (3, 4)$ and $\widetilde{\boldsymbol{\lambda}}^{(0)} = (5, 2)$. With this transformation, all nonzero parameters take values in $[1, 10]$. When we solve the MLE problem (5.8), we treat $\widetilde{\boldsymbol{\theta}}_{\text{num}}$ as the parameter vector instead of $\boldsymbol{\theta}_{\text{num}}$, and solve

$$\min_{\widetilde{\boldsymbol{\theta}}_{\text{num}} \in \widetilde{\boldsymbol{\Theta}}_{\text{num}}} l_{\text{num}}\big(\widetilde{\boldsymbol{\theta}}_{\text{num}} \odot \mathbf{s}^{(0)}\big), \tag{D.6}$$

where $\widetilde{\boldsymbol{\Theta}}_{\text{num}}$ is the transformed set of feasible parameters. The parameter scaling is applied to ensure that all model parameters are of a similar magnitude in the optimization.

In most cases, we have found it sufficient to solve the optimization problem (D.6) once. However, in our codes, we provide an option to solve the problem multiple times, using (i) user-supplied initial guesses, (ii) initial guesses based on the simple estimates from (D.5), with new birth rates selected randomly each time, or (iii) randomly sampled initial guesses, using the parameter generation procedure described in Section D.4 below.

The optimization problems (5.9) for the endpoints of the confidence intervals are solved in a similar way, except the initial guess is taken to be the maximum likelihood estimate.

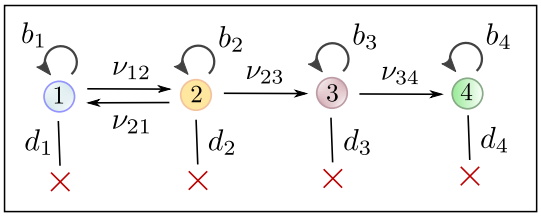

Figure D.2: An example of a four-type switching model where the likelihood function (5.12) for cell fraction data from the main text must be modified to avoid degeneracy issues. This model structure can e.g. arise in the context of epigenetically-driven drug resistance in cancer, where drug-sensitive (type-0) cells can acquire transient resistance (type-1), which then evolves gradually to stable resistance (type-4) in two steps [75].

D.3.2 Cell fraction data

The implementation for cell fraction data is similar with the following modifications. First of all, we parametrize the model in terms of λ_1 and the net birth rate differences $\boldsymbol{\lambda}^{[-1]}$ instead of the net birth rates $\boldsymbol{\lambda}$ (Section 5.5.2). Second, the initial guess for the MLE problem is based on solving the least squares problem (5.16). When solving (5.16), we first compute initial guesses for the switching rates $\boldsymbol{\nu}$ based on part (1) of Proposition 5.4, which shows how $\boldsymbol{\nu}$ can be estimated from the slopes of the mean functions $\mathbf{p}^{(j)}(t)$ at time zero. We approximate the slopes of $\mathbf{p}^{(j)}(t)$ at zero using experimentally observed cell fractions at the first time point. The initial guesses for the remaining parameters are set to 0. The simple problem (5.16) returns estimates for $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}^{[-1]}$, which we use as initial guesses for (5.14).

In our codes, we implement the following likelihood function for cell fraction data:

$$
\begin{aligned}
&\mathcal{L}_{\text{frac}}\big(\boldsymbol{\theta}_{\text{frac}}\big|(\mathbf{f}_{i,\ell,r})_{i,\ell,r}\big) \\
&= \textstyle\prod_{i=1}^{I}\prod_{\ell=1}^{L}\prod_{r=1}^{R}\Big((2\pi)^{K-1}\det\big(\mathbf{C}_i^T\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\mathbf{B}\mathbf{C}_i\big)\Big)^{-1/2} \\
&\quad\cdot\exp\Big(-\tfrac{1}{2}\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{\mathbf{f}_i}(t_\ell)\big)\,\mathbf{B}\mathbf{C}_i\,\big(\mathbf{C}_i^T\mathbf{B}^T\big(N_i^{-1}\mathbf{S}^{\mathbf{f}_i}(t_\ell)+\mathbf{E}_{i,\ell}^{\text{frac}}\big)\,\mathbf{B}\mathbf{C}_i\big)^{-1} \\
&\qquad\qquad\mathbf{C}_i^T\mathbf{B}^T\,\big(\mathbf{f}_{i,\ell,r}-\mathbf{p}^{\mathbf{f}_i}(t_\ell)\big)^T\Big).
\end{aligned}
$$

Recall from (5.12) that the matrix $\mathbf{B}$ is applied to reduce the data vector $\mathbf{f}_{i,\ell,r}$ to a $(K-1)$-dimensional nondegenerate vector. To accommodate reducible switching dynamics, the user is allowed to implement a further reduction in the data by specifying

a $(K-1) \times J_i$ matrix $\mathbf{C}_i$ for each initial condition i. This can for example be useful for the model displayed in Figure D.2, in which case we would take $I = 3$, $\mathbf{C}_1 = \mathbf{C}_2 = \mathbf{I}$ and $\mathbf{C}_3 = \mathbf{e}_3^T$, and we would restrict the set of feasible parameters $\boldsymbol{\Theta}_{\text{frac}}$ so that $\nu_{13} = \nu_{14} = \nu_{24} = \nu_{31} = \nu_{32} = \nu_{41} = \nu_{42} = \nu_{43} = 0$.

D.4 Generation of artifical data

Here, we discuss how the artificial data was generated for the numerical experiments in Section 5.7.3. First, to generate each parameter regime, we sampled the birth rates $\mathbf{b}$ and death rates $\mathbf{d}$ uniformly at random on $(0, 1)$, with the following caveats: The birth rates $\mathbf{b}$ and net birth rates $\boldsymbol{\lambda}$ were required to be larger than 0.01 in absolute value, and at least one of the net birth rates $\lambda_1, \ldots, \lambda_K$ was required to be positive. Each switching rate ν_{ij} was sampled as 10^{-3+2U}, where U is uniform between 0 and 1, meaning that it was sampled log-uniformly between 10^{-3} and 10^{-1}. The starting number of cells N_i was chosen as $N_i = 10^{-3}$, $N_i = 10^{-4}$ or $N_i = 10^{-5}$ for $i = 1, \ldots, K$ based on the order of magnitude of the smallest switching rate. The time points for data collection were selected as $t = 1, \ldots, 6$.

In Section 5.7.4, where the number of time points was doubled, the time points were taken as either $t = 0.5, 1, 1.5, 2, \ldots, 6$ or $t = 1, 2, 3, \ldots, 12$, depending on whether the new time points were added in between or after the previous time points.

Once the parameters were set, we performed stochastic simulations of the model in Section 5.2 to obtain the artificial datasets.

D.5 AIC and BIC weights

To evaluate model fit relative to model complexity in Section 5.8, we compute model weights based on both the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) [155]. For a statistical model with parameters $\boldsymbol{\theta}$ and negative double log-likelihood $l(\boldsymbol{\theta})$, the AIC and BIC are given by

$$\begin{aligned} \text{AIC} &= l\big(\widehat{\boldsymbol{\theta}}\big) + 2p, \\ \text{BIC} &= l\big(\widehat{\boldsymbol{\theta}}\big) + p\log(n), \end{aligned}$$

where $\widehat{\boldsymbol{\theta}}$ is the MLE estimate, p is the number of parameters in the statistical model, and n is the number of datapoints. When comparing two models, the model with the lower AIC or BIC is preferred, depending on which criterion is used. The BIC criterion generally favors simpler models, i.e. models with fewer parameters, to a greater extent than the AIC criterion. For ease of interpretation, we compute model weights based on the two criteria as follows. For the AIC, the weight for model k is given by

$$w\big(\mathrm{AIC}_k\big) = \frac{\exp\big(-\frac{1}{2}\mathrm{AIC}_k\big)}{\sum_m \exp\big(-\frac{1}{2}\mathrm{AIC}_m\big)}.$$

The weight $w(\mathrm{AIC}_k)$ indicates the strength of evidence for model k relative to the other candidate models. Weights for the BIC can be computed analogously [155].

D.6 Proof of Proposition 5.1

Proof of Proposition 5.1. First note that we can write

$$\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) = \textstyle\sum_{j=1}^{K} \mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t), \tag{D.7}$$

where $\big(\mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(s)\big)_{s\geq 0}$ for $j = 1, \ldots, K$ are independent branching processes started with $\lfloor N\alpha_j\rfloor$ cells of type-j, respectively. For each process, we can write

$$\mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t) = \textstyle\sum_{m=1}^{\lfloor N\alpha_j\rfloor} \mathbf{Z}^{(j),m}(t),$$

where $\big(\mathbf{Z}^{(j),m}(s)\big)_{s\geq 0}$ for $m = 1, \ldots, \lfloor N\alpha_j\rfloor$ are i.i.d. copies of the branching process $\big(\mathbf{Z}^{(j)}(s)\big)_{s\geq 0}$ started by a single type-j cell. Set

$$\mathbf{W}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t) := N^{-1/2}\big(\mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t) - N\alpha_j \mathbf{m}^{(j)}(t)\big). \tag{D.8}$$

Let $J \geq 1$ and let $\mathbf{C}$ be a $K \times J$ matrix. By the (multivariate) central limit theorem, as $N \to \infty$,

$$\mathbf{W}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t)\,\mathbf{C} \xrightarrow{d} \mathcal{N}\big(\mathbf{0}, \alpha_j \mathbf{C}^T \boldsymbol{\Sigma}^{(j)}(t)\,\mathbf{C}\big),$$

where $\mathbf{\Sigma}^{(j)}(t)$ is the covariance matrix for $\mathbf{Z}^{(j)}(t)$. We can then conclude from (D.7) that as $N \to \infty$,

$$N^{-1/2}\big(\mathbf{Z}^{\lfloor N\alpha \rfloor}(t)\,\mathbf{C} - N\mathbf{m}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\big) \xrightarrow{d} \mathcal{N}\big(\mathbf{0}, \mathbf{C}^T\mathbf{\Sigma}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\big).$$

It remains to derive the expression (5.4) for the covariance matrix $\mathbf{\Sigma}^{(j)}(t)$. To that end, let $\mathbf{D}^{(j)}(t)$ be the matrix of second factorial moments of $\mathbf{Z}^{(j)}(t)$,

$$D^{(j)}_{k\ell}(t) := \mathbb{E}\big[Z^{(j)}_k(t)\big(Z^{(j)}_\ell(t) - \delta_{k\ell}\big)\big],$$

where $\delta_{k\ell}$ is the Kronecker delta. Let $\mathbf{s} = (s_1, \ldots, s_K)$ be a K-dimensional vector of real numbers and set $h_j := b_j + d_j + \sum_{k\neq j} \nu_{jk}$ for $j = 1, \ldots, K$. Furthermore, let

$$u^{(j)}(\mathbf{s}) := b_j s_j^2 + d_j + \textstyle\sum_{k\neq j} \nu_{jk} s_k - h_j s_j, \quad \mathbf{0} \leq \mathbf{s} \leq \mathbf{1},$$

be the infinitesimal generating function for $\mathbf{Z}^{(j)}(t)$, and let

$$F^{(j)}(\mathbf{s}, t) := \mathbb{E}\big[\mathbf{s}^{\mathbf{Z}^{(j)}(t)}\big] = \mathbb{E}\Big[\textstyle\prod_{k=1}^K s_k^{Z^{(j)}_k(t)}\Big], \quad \mathbf{0} \leq \mathbf{s} \leq \mathbf{1},\ t \geq 0,$$

be the probability generating function for $\mathbf{Z}^{(j)}(t)$. With this notation, we can write the Kolmogorov forward equation for $\mathbf{Z}^{(j)}(t)$ as

$$\tfrac{\partial}{\partial t} F^{(j)}(\mathbf{s}, t) = \textstyle\sum_{i=1}^K u^{(i)}(\mathbf{s})\, \tfrac{\partial}{\partial s_i} F^{(j)}(\mathbf{s}, t).$$

Then, for $k, \ell = 1, \ldots, K$,

$$\begin{aligned}
&\tfrac{\partial}{\partial t}\big(\tfrac{\partial}{\partial s_k}\tfrac{\partial}{\partial s_\ell} F^{(j)}(\mathbf{s}, t)\big) \\
&= \textstyle\sum_{i=1}^K \big(\tfrac{\partial}{\partial s_k}\tfrac{\partial}{\partial s_\ell} u^{(i)}(\mathbf{s})\tfrac{\partial}{\partial s_i} F^{(j)}(\mathbf{s}, t) + \tfrac{\partial}{\partial s_\ell} u^{(i)}(\mathbf{s})\tfrac{\partial}{\partial s_k}\tfrac{\partial}{\partial s_i} F^{(j)}(\mathbf{s}, t) \\
&\qquad + \tfrac{\partial}{\partial s_k} u^{(i)}(\mathbf{s})\tfrac{\partial}{\partial s_\ell}\tfrac{\partial}{\partial s_i} F^{(j)}(\mathbf{s}, t) + u^{(i)}(\mathbf{s})\tfrac{\partial}{\partial s_k}\tfrac{\partial}{\partial s_\ell}\tfrac{\partial}{\partial s_i} F^{(j)}(\mathbf{s}, t)\big).
\end{aligned} \tag{D.9}$$

Now,

$$\frac{\partial}{\partial s_k} u^{(i)}(\mathbf{s}) = \begin{cases} 2b_i s_i - h_i, & k = i, \\ \nu_{ik}, & k \neq i, \end{cases}$$
$$\frac{\partial}{\partial s_k}\frac{\partial}{\partial s_\ell} u^{(i)}(\mathbf{s}) = \delta_{ki}\delta_{\ell i} 2b_i.$$

Let $\mathbf{A}$ be the infinitesimal generator and $\mathbf{M}(t)$ be the mean matrix as defined in Section 5.2.3. Since

$$\begin{aligned} a_{ik} &= \tfrac{\partial}{\partial s_k} u^{(i)}(\mathbf{s})|_{\mathbf{s}=\mathbf{1}}, \\ M_{jk}(t) &= \tfrac{\partial}{\partial s_k} F^{(j)}(\mathbf{s}, t)\big|_{\mathbf{s}=1}, \\ D^{(j)}_{k\ell}(t) &= \tfrac{\partial}{\partial s_k}\tfrac{\partial}{\partial s_\ell} F^{(j)}(\mathbf{s}, t)\big|_{\mathbf{s}=1}, \end{aligned}$$

and $u^{(i)}(\mathbf{1}) = 0$, we can conclude from (D.9) that

$$\begin{aligned} \tfrac{d}{dt} D^{(j)}_{k\ell}(t) &= \textstyle\sum_{i=1}^K \big(\delta_{ki}\delta_{\ell i} 2b_i M_{ji}(t) + a_{i\ell} D^{(j)}_{ki}(t) + a_{ik} D^{(j)}_{\ell i}(t)\big) \\ &= \textstyle\sum_{i=1}^K a_{ik} D^{(j)}_{i\ell}(t) + \sum_{i=1}^K D^{(j)}_{ki}(t) a_{i\ell} + \delta_{k\ell} 2b_k M_{jk}(t). \end{aligned}$$

In the second step, we use that $D^{(j)}_{i\ell}(t) = D^{(j)}_{\ell i}(t)$. This yields a Lyapunov matrix differential equation,

$$\tfrac{d}{dt}\mathbf{D}^{(j)}(t) = \mathbf{A}^T\mathbf{D}^{(j)}(t) + \mathbf{D}^{(j)}(t)\mathbf{A} + 2\,\mathrm{diag}\big(\mathbf{b} \odot \mathbf{m}^{(j)}(t)\big), \tag{D.10}$$

with initial condition $\mathbf{D}^{(j)}(0) = \mathbf{0}$. The solution is

$$\begin{aligned} \mathbf{D}^{(j)}(t) &= 2\exp(t\mathbf{A}^T)\big(\textstyle\int_0^t \exp(-\tau\mathbf{A}^T)\mathrm{diag}\big(\mathbf{b} \odot \mathbf{m}^{(j)}(\tau)\big)\exp(-\tau\mathbf{A})d\tau\big)\exp(t\mathbf{A}) \\ &= 2\textstyle\int_0^t (\mathbf{M}(t-\tau))^T \mathrm{diag}\big(\mathbf{b} \odot \mathbf{m}^{(j)}(\tau)\big)(\mathbf{M}(t-\tau))d\tau, \end{aligned}$$

and the expression (5.4) for $\mathbf{\Sigma}^{(j)}(t)$ follows from the fact that

$$\mathbf{\Sigma}^{(j)}(t) = \mathbf{D}^{(j)}(t) + \mathrm{diag}\big(\mathbf{m}^{(j)}(t)\big) - (\mathbf{m}^{(j)}(t))^T\mathbf{m}^{(j)}(t). \tag{D.11}$$

□

D.7 Proof of Proposition 5.2

Proof of Proposition 5.2. Recall from (D.7) in the proof of Proposition 5.1 that we can write

$$\mathbf{Z}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) = \sum_{j=1}^{K} \mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t),$$

where $\left(\mathbf{Z}^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(s)\right)_{s\geq 0}$ for $j = 1, \ldots, K$ are independent branching processes started with $\lfloor N\alpha_j\rfloor$ cells of type-j, respectively. Define

$$U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) := \sum_{k=1}^{K} Z_k^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) = \sum_{j=1}^{K}\sum_{k=1}^{K} Z_k^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t)$$

as the total population size at time t and note that

$$\Delta_i^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) = \frac{Z_i^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)}{U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)} = \frac{\sum_{j=1}^{K} Z_i^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t)}{U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)}.$$

We can therefore write

$$\begin{aligned}
&\sqrt{N}\left(\Delta_i^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) - p_i^{\boldsymbol{\alpha}}(t)\right)\\
&= \frac{\sqrt{N}}{U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)}\left(\sum_{j=1}^{K}\left((1-p_i^{\boldsymbol{\alpha}}(t))Z_i^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t) - p_i^{\boldsymbol{\alpha}}(t)\sum_{k\neq i} Z_k^{\lfloor N\alpha_j\rfloor \mathbf{e}_j}(t)\right)\right).
\end{aligned}$$

Note that by definition,

$$\mathbf{p}^{\boldsymbol{\alpha}}(t) = \left(\mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T\right)^{-1}\mathbf{m}^{\boldsymbol{\alpha}}(t) = \left(\sum_{j=1}^{K}\sum_{k=1}^{K}\alpha_j m_k^{(j)}(t)\right)^{-1}\sum_{j=1}^{K}\alpha_j \mathbf{m}^{(j)}(t).$$

It follows that

$$\begin{aligned}
&\sum_{j=1}^{K} N\alpha_j\left((1-p_i^{\boldsymbol{\alpha}}(t))m_i^{(j)}(t) - p_i^{\boldsymbol{\alpha}}(t)\sum_{k\neq i} m_k^{(j)}(t)\right)\\
&= N\left(\sum_{j=1}^{K}\alpha_j m_i^{(j)}(t) - p_i^{\boldsymbol{\alpha}}(t)\sum_{j=1}^{K}\sum_{k=1}^{K}\alpha_j m_k^{(j)}(t)\right)\\
&= 0.
\end{aligned}$$

We can therefore write

$$\begin{aligned}&\sqrt{N}\big(\Delta_i^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) - p_i^{\boldsymbol{\alpha}}(t)\big)\\ &= \tfrac{N}{U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)}\Big(\textstyle\sum_{j=1}^K \Big((1-p_i^{\boldsymbol{\alpha}}(t))W_i^{\lfloor N\alpha_j\rfloor\mathbf{e}_j}(t) - p_i^{\boldsymbol{\alpha}}(t)\textstyle\sum_{k\neq i} W_k^{\lfloor N\alpha_j\rfloor\mathbf{e}_j}(t)\Big)\Big),\end{aligned}$$

where the vector $\mathbf{W}^{\lfloor N\alpha_j\rfloor\mathbf{e}_j}(t)$ is defined as in (D.8). In vector form, this becomes

$$\sqrt{N}\big(\boldsymbol{\Delta}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t) - \mathbf{p}^{\boldsymbol{\alpha}}(t)\big) = \tfrac{N}{U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)}\textstyle\sum_{j=1}^K \mathbf{W}^{\lfloor N\alpha_j\rfloor\mathbf{e}_j}(t)\mathbf{Q}^{\boldsymbol{\alpha}}(t),$$

where $\mathbf{Q}^{\boldsymbol{\alpha}}(t)$ is defined as in (5.10). By the law of large numbers, $U^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)/N \to \mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T$ almost surely as $N\to\infty$. Let $J\geq 1$ and let $\mathbf{C}$ be a $K\times J$ matrix. By the (multivariate) central limit theorem,

$$\mathbf{W}^{\lfloor N\alpha_j\rfloor\mathbf{e}_j}(t)\,\mathbf{Q}^{\boldsymbol{\alpha}}(t)\,\mathbf{C} \xrightarrow{d} \mathcal{N}\Big(\mathbf{0}, \alpha_j\mathbf{C}^T\big(\mathbf{Q}^{\boldsymbol{\alpha}}(t)\big)^T\boldsymbol{\Sigma}^{(j)}(t)\mathbf{Q}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\Big).$$

Writing $\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t) = \sum_{j=1}^K \alpha_j\boldsymbol{\Sigma}^{(j)}(t)$, it finally follows from Slutsky's theorem that

$$\sqrt{N}\big(\boldsymbol{\Delta}^{\lfloor N\boldsymbol{\alpha}\rfloor}(t)\,\mathbf{C} - \mathbf{p}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\big) \xrightarrow{d} \mathcal{N}\Big(\mathbf{0}, \big(\mathbf{m}^{\boldsymbol{\alpha}}(t)\mathbf{1}^T\big)^{-2}\mathbf{C}^T\big(\mathbf{Q}^{\boldsymbol{\alpha}}(t)\big)^T\boldsymbol{\Sigma}^{\boldsymbol{\alpha}}(t)\mathbf{Q}^{\boldsymbol{\alpha}}(t)\,\mathbf{C}\Big).$$

It remains to show that $\mathbf{p}^{\boldsymbol{\alpha}}(t)$ can be written solely as a function of the switching rates $\boldsymbol{\nu}$ and the net birth rate differences $\boldsymbol{\lambda}^{[-j]}$ for any $j=1,\ldots,K$. To this end, we define

$$\mathbf{A}^{[-j]} := \mathbf{A} - \lambda_j\mathbf{I}, \tag{D.12}$$

where $\mathbf{I}$ is the $K\times K$ identity matrix, and

$$\mathbf{M}^{[-j]}(t) := \exp\big(t\mathbf{A}^{[-j]}\big) = \textstyle\sum_{k=0}^\infty (t^k/k!)\big(\mathbf{A}^{[-j]}\big)^k, \quad t\geq 0. \tag{D.13}$$

Note that $\mathbf{A}^{[-j]}$ and $\mathbf{M}^{[-j]}(t)$ only depend on $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}^{[-j]}$. It is easy to see that

$$\mathbf{M}(t) = e^{\lambda_j t}\mathbf{M}^{[-j]}(t),$$

from which it follows that

$$\begin{aligned}\mathbf{p}^{\boldsymbol{\alpha}}(t) &= \big(\boldsymbol{\alpha}\mathbf{M}(t)\mathbf{1}^T\big)^{-1}\boldsymbol{\alpha}\mathbf{M}(t)\\ &= \big(\boldsymbol{\alpha}\mathbf{M}^{[-j]}(t)\mathbf{1}^T\big)^{-1}\big(\boldsymbol{\alpha}\mathbf{M}^{[-j]}(t)\big), \quad t \geq 0. \end{aligned} \tag{D.14}$$

This completes the proof. □

D.8 Proof of Proposition 5.3

Proof of Proposition 5.3. (1) Since $\mathbf{M}(t) = \exp(t\mathbf{A}) = \sum_{k=0}^{\infty}(1/k!)t^k\mathbf{A}^k$, we have $\frac{d}{dt}\mathbf{M}(t) = \mathbf{A}\mathbf{M}(t)$. By taking $t = 0$ and noting that $\mathbf{M}(0) = \mathbf{I}$, we obtain

$$\tfrac{d}{dt}\mathbf{M}(t)\big|_{t=0} = \mathbf{A}.$$

If $\frac{d}{dt}\mathbf{M}(t)\big|_{t=0}$ is known, we can recover the switching rate ν_{jk} for $k \neq j$ by recalling that $a_{jk} = \nu_{jk}$. We can then recover λ_j for $j = 1, \ldots, K$ by recalling that $a_{jj} = \lambda_j - \sum_{k\neq j}\nu_{jk}$.

(2) Recall that $\mathbf{m}^{(j)}(t) = \mathbf{e}_j\mathbf{M}(t)$. By (D.11) in the proof of Proposition 5.1, we can write

$$\begin{aligned}&\tfrac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\\ &= \tfrac{d}{dt}\mathbf{D}^{(j)}(t) + \operatorname{diag}\big(\mathbf{e}_j\mathbf{A}\mathbf{M}(t)\big) - \mathbf{A}^T\big(\mathbf{M}(t)\big)^T\mathbf{e}_j^T\mathbf{e}_j\mathbf{M}(t) - \big(\mathbf{M}(t)\big)^T\mathbf{e}_j^T\mathbf{e}_j\mathbf{A}\mathbf{M}(t),\end{aligned}$$

where $\mathbf{D}^{(j)}(t)$ is the matrix of second factorial moments of $\mathbf{Z}^{(j)}(t)$. Next, by taking $t = 0$ in (D.10) and noting that $\mathbf{D}^{(j)}(0) = \mathbf{0}$ and $\mathbf{m}^{(j)}(0) = \mathbf{e}_j$ for all $j = 1, \ldots, K$, we see that

$$\tfrac{d}{dt}\mathbf{D}^{(j)}(t)\big|_{t=0} = 2b_j\mathbf{e}_j^T\mathbf{e}_j.$$

It follows that

$$\tfrac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\big|_{t=0} = 2b_j\mathbf{e}_j^T\mathbf{e}_j + \operatorname{diag}\big(\mathbf{e}_j\mathbf{A}\big) - \big(\mathbf{e}_j^T\mathbf{e}_j\mathbf{A}\big)^T - \mathbf{e}_j^T\mathbf{e}_j\mathbf{A}. \tag{D.15}$$

For each $j = 1, \ldots, K$, if the switching rates ν_{jk} for $k \neq j$ and the net birth rate λ_j are known, the birth rate b_j can be recovered from $\left(\frac{d}{dt}\mathbf{\Sigma}^{(j)}(t)\big|_{t=0}\right)_{jj}$ using this expression. □

D.9 Proof of Proposition 5.4

Proof of Proposition 5.4. We begin by establishing some notation. First, define $\mathbf{Q}^{(j)}(t) := \mathbf{Q}^{\mathbf{e}_j}(t)$ and $\mathbf{Q}^{(j)} := \mathbf{Q}^{(j)}(0) = \mathbf{I} - \mathbf{1}^T\mathbf{e}_j$, with $\mathbf{Q}^{\mathbf{e}_j}(t)$ defined as in (5.10). Also define

$$\mathbf{V} := \mathbf{A} - \operatorname{diag}(\boldsymbol{\lambda}) \tag{D.16}$$

as the infinitesimal generator $\mathbf{A}$ with the net birth rates $\boldsymbol{\lambda}$ removed from the diagonal. Let $\mathbf{v}^{(j)}$ denote the j-th row vector of $\mathbf{V}$ with coordinates $v_k^{(j)} = \nu_{jk}$ for $k \neq j$ and $v_j^{(j)} = -\sum_{k\neq j}\nu_{jk}$, and note that

$$\mathbf{v}^{(j)} = \mathbf{e}_j\mathbf{V} = \mathbf{e}_j\mathbf{A}^{[-j]}, \tag{D.17}$$

where $\mathbf{A}^{[-j]}$ is defined as in (D.12). Also note that $\mathbf{v}^{(j)}\mathbf{1}^T = 0$. In the proof, we will rely on the following basic facts:

$$\begin{aligned} \mathbf{e}_j\mathbf{Q}^{(j)} &= \mathbf{e}_j\big(\mathbf{I} - \mathbf{1}^T\mathbf{e}_j\big) = \mathbf{0}, \\ \mathbf{v}^{(j)}\mathbf{Q}^{(j)} &= \mathbf{v}^{(j)}\big(\mathbf{I} - \mathbf{1}^T\mathbf{e}_j\big) = \mathbf{v}^{(j)}. \end{aligned} \tag{D.18}$$

(1) Since $\mathbf{p}^{(j)}(t) = \big(\mathbf{e}_j\exp(t\mathbf{A})\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\exp(t\mathbf{A})\big)$, we can write

$$\begin{aligned} \tfrac{d}{dt}\mathbf{p}^{(j)}(t) = &\big(\mathbf{e}_j\exp(t\mathbf{A})\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\big) \\ &- \big(\mathbf{e}_j\exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\mathbf{1}^T\big)\big(\mathbf{e}_j\exp(t\mathbf{A})\big). \end{aligned} \tag{D.19}$$

Since $\exp(\mathbf{0}) = \mathbf{I}$, $\mathbf{e}_j\mathbf{1}^T = 1$ and $\mathbf{e}_j\mathbf{A}\mathbf{1}^T = \lambda_j$, we obtain by (D.17),

$$\tfrac{d}{dt}\mathbf{p}^{(j)}(t)\big|_{t=0} = \mathbf{e}_j(\mathbf{A} - \lambda_j\mathbf{I}) = \mathbf{e}_j\mathbf{A}^{[-j]} = \mathbf{v}^{(j)}. \tag{D.20}$$

Since the k-th coordinate of $\mathbf{v}^{(j)}$ is ν_{jk} for $k \neq j$, we can recover ν_{jk} from the k-th coordinate of $\frac{d}{dt}\mathbf{p}^{(j)}(t)\big|_{t=0}$.

(2) (i) Using (D.19), we begin by writing

$$\begin{aligned}\tfrac{d^2}{dt^2}\mathbf{p}^{(j)}(t) = {} & \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\mathbf{A}^2\exp(t\mathbf{A})\big) \\ & - \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\mathbf{1}^T\big)\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\big) \\ & + 2\big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-3}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\mathbf{1}^T\big)^2\big(\mathbf{e}_j\exp(t\mathbf{A})\big) \\ & - \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}^2\exp(t\mathbf{A})\mathbf{1}^T\big)\big(\mathbf{e}_j\exp(t\mathbf{A})\big) \\ & - \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\mathbf{1}^T\big)\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\big).\end{aligned}$$

Since $\exp(\mathbf{0}) = \mathbf{I}$, $\mathbf{e}_j\mathbf{1}^T = 1$, $\mathbf{e}_j\mathbf{A}\mathbf{1}^T = \lambda_j$, $\mathbf{v}^{(j)} = \mathbf{e}_j\mathbf{A}^{[-j]}$ and $\mathbf{Q}^{(j)} = \mathbf{I} - \mathbf{1}^T\mathbf{e}_j$,

$$\begin{aligned}\tfrac{d^2}{dt^2}\mathbf{p}^{(j)}(t)\big|_{t=0} &= 2\lambda_j\mathbf{e}_j\big(\lambda_j\mathbf{I} - \mathbf{A}\big) + \mathbf{e}_j\mathbf{A}^2\big(\mathbf{I} - \mathbf{1}^T\mathbf{e}_j\big) \\ &= -2\lambda_j\mathbf{v}^{(j)} + \mathbf{e}_j\mathbf{A}^2\mathbf{Q}^{(j)}.\end{aligned}$$

Recalling that $\mathbf{A} = \mathbf{A}^{[-j]} + \lambda_j\mathbf{I}$ by (D.12), we can write

$$\begin{aligned}\mathbf{e}_j\mathbf{A}^2 &= \mathbf{e}_j\big(\mathbf{A}^{[-j]}\big)^2 + 2\lambda_j\mathbf{e}_j\mathbf{A}^{[-j]} + \lambda_j^2\mathbf{e}_j \\ &= \mathbf{v}^{(j)}\mathbf{A}^{[-j]} + 2\lambda_j\mathbf{v}^{(j)} + \lambda_j^2\mathbf{e}_j. \qquad \text{(D.21)}\end{aligned}$$

Since $\mathbf{e}_j\mathbf{Q}^{(j)} = \mathbf{0}$ and $\mathbf{v}^{(j)}\mathbf{Q}^{(j)} = \mathbf{v}^{(j)}$ by (D.18), it follows that

$$\mathbf{e}_j\mathbf{A}^2\mathbf{Q}^{(j)} = \mathbf{v}^{(j)}\mathbf{A}^{[-j]}\mathbf{Q}^{(j)} + 2\lambda_j\mathbf{v}^{(j)}\mathbf{Q}^{(j)} = \mathbf{v}^{(j)}\mathbf{A}^{[-j]}\mathbf{Q}^{(j)} + 2\lambda_j\mathbf{v}^{(j)}, \quad \text{(D.22)}$$

which implies

$$\tfrac{d^2}{dt^2}\mathbf{p}^{(j)}(t)\big|_{t=0} = \mathbf{v}^{(j)}\mathbf{A}^{[-j]}\mathbf{Q}^{(j)}. \qquad \text{(D.23)}$$

It is straightforward to verify that for $i \neq j$,

$$\begin{aligned}&\big(\mathbf{v}^{(j)}\mathbf{A}^{[-j]}\mathbf{Q}^{(j)}\big)_i \\ &= \nu_{ji}(\lambda_i - \lambda_j) - \nu_{ij}\big(\textstyle\sum_{k\neq j}\nu_{jk}\big) - \nu_{ji}\big(\textstyle\sum_{\ell\neq i}\nu_{i\ell}\big) + \textstyle\sum_{m\neq j, m\neq i}\nu_{jm}\nu_{im}.\end{aligned}$$

If $\boldsymbol{\nu}$ and $\tfrac{d^2}{dt^2}\mathbf{p}^{(j)}(t)\big|_{t=0}$ are known, we can therefore use (D.23) to get an equation for $\lambda_i - \lambda_j$ of the form $\nu_{ji}(\lambda_i - \lambda_j) = C$ for some constant C.

If $\nu_{ji} \neq 0$, we immediately obtain the value of $\lambda_i - \lambda_j$. If $\nu_{ji} = 0$, then by our assumption of irreducibility, there exist integers $n_1, \ldots, n_k$ so that $\nu_{n_0 n_1}\nu_{n_1 n_2}\cdots\nu_{n_k n_{k+1}} > 0$, where $n_0 = j$ and $n_{k+1} = i$. For each $\ell = 0, \ldots, k$, we can use the fact that $\nu_{n_\ell n_{\ell+1}} > 0$ to obtain the value of $\lambda_{n_{\ell+1}} - \lambda_{n_\ell}$. Since $\lambda_{n_{k+1}} - \lambda_{n_0} = \sum_{\ell=0}^{k}(\lambda_{n_{\ell+1}} - \lambda_{n_\ell})$, we also obtain the value of $\lambda_i - \lambda_j$.

(ii) We know from (D.19) that

$$\begin{aligned}\tfrac{d}{dt}\mathbf{p}^{(j)}(t) &= \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\big)\\ &\quad - \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\exp(t\mathbf{A})\mathbf{1}^T\big)\big(\mathbf{e}_j\exp(t\mathbf{A})\big).\end{aligned}$$

We also know from (5.1) that

$$\lim_{t\to\infty} e^{-\sigma t}\exp\big(t\mathbf{A}\big) = \boldsymbol{\beta}^T\boldsymbol{\gamma},$$

where $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are positive vectors. It follows that as $t \to \infty$,

$$\begin{aligned}\tfrac{d}{dt}\mathbf{p}^{(j)}(t) &\to \big(\mathbf{e}_j\boldsymbol{\beta}^T\boldsymbol{\gamma}\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\boldsymbol{\gamma}\big) - \big(\mathbf{e}_j\boldsymbol{\beta}^T\boldsymbol{\gamma}\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\boldsymbol{\gamma}\mathbf{1}^T\big)\big(\mathbf{e}_j\boldsymbol{\beta}^T\boldsymbol{\gamma}\big)\\ &= \big(\mathbf{e}_j\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}}\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}}\big) - \big(\mathbf{e}_j\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}}\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}}\mathbf{1}^T\big)\big(\mathbf{e}_j\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}}\big),\end{aligned}$$

where $\overline{\boldsymbol{\gamma}}$ is the normalized version of $\boldsymbol{\gamma}$, see (5.2). Since $\mathbf{e}_j\boldsymbol{\beta}^T = \beta_j > 0$ and $\overline{\boldsymbol{\gamma}}\mathbf{1}^T = 1$, we obtain

$$\tfrac{d}{dt}\mathbf{p}^{(j)}(t) \to \beta_j^{-1}\big(\mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\,\overline{\boldsymbol{\gamma}} - \mathbf{e}_j\mathbf{A}\boldsymbol{\beta}^T\overline{\boldsymbol{\gamma}}\big) = \mathbf{0}. \tag{D.24}$$

On the other hand, by noting that $\mathbf{A}$ and $\exp(t\mathbf{A})$ commute, we can rewrite the expression (D.19) for $\frac{d}{dt}\mathbf{p}^{(j)}(t)$ as

$$\begin{aligned}\tfrac{d}{dt}\mathbf{p}^{(j)}(t) &= \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-1}\big(\mathbf{e}_j\exp(t\mathbf{A})\mathbf{A}\big)\\ &\quad - \big(\mathbf{e}_j \exp(t\mathbf{A})\mathbf{1}^T\big)^{-2}\big(\mathbf{e}_j\exp(t\mathbf{A})\mathbf{A}\mathbf{1}^T\big)\big(\mathbf{e}_j\exp(t\mathbf{A})\big).\end{aligned}$$

Since $\mathbf{A}\mathbf{1}^T = \boldsymbol{\lambda}^T$, $\mathbf{A} = \mathbf{V} + \operatorname{diag}(\boldsymbol{\lambda})$ by (D.16), $\overline{\boldsymbol{\gamma}}\operatorname{diag}(\boldsymbol{\lambda}) = \boldsymbol{\lambda}\operatorname{diag}\big(\overline{\boldsymbol{\gamma}}\big)$ and

$\overline{\gamma}\boldsymbol{\lambda}^T = \boldsymbol{\lambda}\overline{\gamma}^T$, we get as $t \to \infty$,

$$\tfrac{d}{dt}\mathbf{p}^{(j)}(t) \to \overline{\gamma}\mathbf{A} - \overline{\gamma}\mathbf{A}\mathbf{1}^T\overline{\gamma} = \overline{\gamma}\big(\mathbf{A} - \boldsymbol{\lambda}^T\overline{\gamma}\big) = \overline{\gamma}\,\mathbf{V} + \boldsymbol{\lambda}\,\mathrm{diag}\big(\overline{\gamma}\big) - \boldsymbol{\lambda}\overline{\gamma}^T\overline{\gamma}. \tag{D.25}$$

Combining (D.24) and (D.25), we obtain the following linear system for $\boldsymbol{\lambda}$:

$$\boldsymbol{\lambda}\big(\mathrm{diag}\big(\overline{\gamma}\big) - \overline{\gamma}^T\overline{\gamma}\big) = -\overline{\gamma}\,\mathbf{V}.$$

It is straightforward to verify that this system is solved by

$$\boldsymbol{\lambda} = \mathbf{a} + x\mathbf{1}, \quad x \in \mathbb{R},$$

for some vector $\mathbf{a}$, which can be used to extract $\boldsymbol{\lambda}^{[-1]}$.

(3) By the definition of $\mathbf{S}^{(j)}(t)$ in (5.10),

$$\begin{aligned}\tfrac{d}{dt}\mathbf{S}^{(j)}(t) = {} & \tfrac{d}{dt}\big(\mathbf{e}_j\mathbf{M}(t)\mathbf{1}^T\big)^{-2}\big(\mathbf{Q}^{(j)}(t)\big)^T\boldsymbol{\Sigma}^{(j)}(t)\mathbf{Q}^{(j)}(t) \\ & + \big(\mathbf{e}_j\mathbf{M}(t)\mathbf{1}^T\big)^{-2}\tfrac{d}{dt}\big(\mathbf{Q}^{(j)}(t)\big)^T\boldsymbol{\Sigma}^{(j)}(t)\mathbf{Q}^{(j)}(t) \\ & + \big(\mathbf{e}_j\mathbf{M}(t)\mathbf{1}^T\big)^{-2}\big(\mathbf{Q}^{(j)}(t)\big)^T\tfrac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\mathbf{Q}^{(j)}(t) \\ & + \big(\mathbf{e}_j\mathbf{M}(t)\mathbf{1}^T\big)^{-2}\big(\mathbf{Q}^{(j)}(t)\big)^T\boldsymbol{\Sigma}^{(j)}(t)\tfrac{d}{dt}\mathbf{Q}^{(j)}(t).\end{aligned} \tag{D.26}$$

Since $\boldsymbol{\Sigma}^{(j)}(0) = \mathbf{0}$ and $\mathbf{e}_j\mathbf{M}(0)\mathbf{1}^T = 1$, we obtain

$$\tfrac{d}{dt}\mathbf{S}^{(j)}(t)\big|_{t=0} = \big(\mathbf{Q}^{(j)}\big)^T\big(\tfrac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\big|_{t=0}\big)\mathbf{Q}^{(j)}.$$

From (D.15) in the proof of Proposition 5.3, we know that

$$\begin{aligned}\tfrac{d}{dt}\boldsymbol{\Sigma}^{(j)}(t)\big|_{t=0} &= 2b_j\mathbf{e}_j^T\mathbf{e}_j + \mathrm{diag}\big(\mathbf{e}_j\mathbf{A}\big) - \big(\mathbf{e}_j^T\mathbf{e}_j\mathbf{A}\big)^T - \mathbf{e}_j^T\mathbf{e}_j\mathbf{A} \\ &= \mathrm{diag}\big(\mathbf{e}_j\mathbf{A}^{[-j]}\big) - \big(\mathbf{e}_j^T\mathbf{e}_j\mathbf{A}\big)^T - \mathbf{e}_j^T\mathbf{e}_j\mathbf{A} + (2b_j + \lambda_j)\mathbf{e}_j^T\mathbf{e}_j,\end{aligned} \tag{D.27}$$

where in the second step, we write $\mathbf{A} = \mathbf{A}^{[-j]} + \lambda_j\mathbf{I}$. Since $\mathbf{e}_j\mathbf{Q}^{(j)} = \mathbf{0}$ and $\mathbf{e}_j\mathbf{A}^{[-j]} =$

$\mathbf{v}^{(j)}$, we obtain

$$\tfrac{d}{dt}\mathbf{S}^{(j)}(t)\big|_{t=0} = \big(\mathbf{Q}^{(j)}\big)^T \mathrm{diag}\big(\mathbf{e}_j\mathbf{A}^{[-j]}\big)\mathbf{Q}^{(j)} = \big(\mathbf{Q}^{(j)}\big)^T \mathrm{diag}\big(\mathbf{v}^{(j)}\big)\mathbf{Q}^{(j)}. \tag{D.28}$$

It is straightforward to verify that the (j,k)-th coordinate of $\big(\mathbf{Q}^{(j)}\big)^T \mathrm{diag}\big(\mathbf{v}^{(j)}\big)\mathbf{Q}^{(j)}$ is $-\nu_{jk}$. Thus, knowledge of the switching rates $\boldsymbol{\nu}$ follows immediately from knowledge of $\frac{d}{dt}\mathbf{S}^{(j)}(t)\big|_{t=0}$ for $j = 1, \ldots, K$, but no other parameters can be extracted. $\square$

Appendix E

Effect of phenotypic switching on the site frequency spectrum of neutral mutations during tumor progression

In this section, we derive the long-run site frequency spectrum of neutral mutations for tumors undergoing phenotypic switching. This is an extension of the results of Proposition 3.1 in Chapter 3 from a single-type model to a multi-type model.

E.1 Two-type model

We first consider a two-type branching process model, where cells can switch stochastically between two phenotypes, the type-0 phenotype and the type-1 phenotype. Each type-i cell divides into two type-i cells at rate $r_i \geq 0$ and dies at rate $d_i \geq 0$, with $\lambda_i := r_i - d_i > 0$ the net birth rate. A type-0 cell switches to type-1 at rate $\mu_0 > 0$, and a type-1 cell switches to type-0 at rate $\mu_1 > 0$. Let $Z_i(t)$ denote the number of type-i cells at time t and set $\mathbf{Z}(t) := (Z_0(t), Z_1(t))$. We define the event of nonextinction of

the population as

$$\Omega_\infty := \{Z_0(t) + Z_1(t) > 0 \text{ for all } t > 0\}, \tag{E.1}$$

and the extinction probability for a clone started by a single type-j cell as

$$p_j := \mathbb{P}(\Omega_\infty^c | \mathbf{Z}(0) = \mathbf{e}_j) = \mathbb{P}(Z_0(t) + Z_1(t) = 0 \text{ for some } t > 0 | \mathbf{Z}(0) = \mathbf{e}_j), \tag{E.2}$$

where $\mathbf{e}_j$ is the unit vector with 1 in the j-th coordinate. By conditioning on whether the initial event is a division, death or switch between types, we see that p_0 and p_1 must satisfy

$$\begin{aligned} (r_0 + d_0 + \mu_0)p_0 &= r_0 p_0^2 + d_0 + \mu_0 p_1, \\ (r_1 + d_1 + \mu_1)p_1 &= r_1 p_1^2 + d_1 + \mu_1 p_0, \end{aligned} \tag{E.3}$$

and it can be shown that p_0 and p_1 with $0 \leq p_0, p_1 < 1$ are uniquely determined by this system [6]. During each cell division of a type-i cell, we assume that w_i mutations accumulate on average.

The evolution of the process $(\mathbf{Z}(t))_{t \geq 0}$ is determined by its infinitesimal generator $\mathbf{A} = (a_{ij})_{i,j \in \{0,1\}}$. For our model, the infinitesimal generator is

$$\mathbf{A} = \begin{bmatrix} \lambda_0 - \mu_0 & \mu_0 \\ \mu_1 & \lambda_1 - \mu_1 \end{bmatrix}, \tag{E.4}$$

which has eigenvalues $0 < \rho < \sigma$ given by

$$\left.\begin{matrix} \sigma \\ \rho \end{matrix}\right\} = \frac{(\lambda_0 - \mu_0) + (\lambda_1 - \mu_1) \pm \sqrt{((\lambda_0 - \mu_0) - (\lambda_1 - \mu_1))^2 + 4\mu_0\mu_1}}{2}. \tag{E.5}$$

Define

$$\begin{aligned} \delta &:= \big((\lambda_0 - \mu_0) - \rho\big)/\mu_1 > 0, \\ \gamma &:= \big(\sigma - (\lambda_0 - \mu_0)\big)/\mu_1 > 0, \end{aligned} \tag{E.6}$$

and set $\alpha_0 := (1+\delta)/(\delta+\gamma) > 0$, $\alpha_1 := \gamma\alpha_0 > 0$, $\beta_0 := \delta/(1+\delta) > 0$ and $\beta_1 = (1/\delta)\beta_0 > 0$. It is easy to see that $\boldsymbol{\alpha} = (\alpha_0, \alpha_1)$ is a left eigenvector and $\boldsymbol{\beta} = (\beta_0, \beta_1)$ is a right eigenvector of $\mathbf{A}$ corresponding to σ, and that these vectors satisfy

$$\beta_0 + \beta_1 = 1 \quad \text{and} \quad \alpha_0\beta_0 + \alpha_1\beta_1 = 1. \tag{E.7}$$

Let $\phi_i^{(j)}(t) := \mathbb{E}[Z_i(t)|\mathbf{Z}(0) = \mathbf{e}_j]$ be the mean number of type-i cells at time t when the process is started by a single type-j cell. Set $\boldsymbol{\phi}^{(j)}(t) := \big(\phi_0^{(j)}(t), \phi_1^{(j)}(t)\big)$. It can be shown that in the long run,

$$\phi_i^{(j)}(t) \sim \beta_j\alpha_i e^{\sigma t}, \quad t \to \infty, \tag{E.8}$$

where $f(t) \sim g(t)$ means $f(t)/g(t) \to 1$ as $t \to \infty$. The existence of positive eigenvectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ that satisfy (E.7) and (E.8) is guaranteed by an application of the Perron-Frobenius theorem to the *mean matrix* $\mathbf{M}(t) = \exp(t\mathbf{A})$, see Section V.7 of Ahtreya & Ney [6]. Moreover, the asymptotic exponential growth of (E.8) holds on a sample-path basis: By Theorem 2 of Section V.7 of [6], there exists a nonnegative random variable W so that

$$\lim_{t\to\infty} \mathbf{Z}(t)e^{-\sigma t} = W\boldsymbol{\alpha} = (W\alpha_0, W\alpha_1) \tag{E.9}$$

almost surely. The distribution of W depends on the initial condition, with

$$\begin{aligned}
&\mathbb{P}(W > 0|\mathbf{Z}(0) = \mathbf{e}_j) = 1 - p_j, \\
&\mathbb{E}[W|\mathbf{Z}(0) = \mathbf{e}_j] = \beta_j.
\end{aligned}$$

The ratio α_1/α_0 is the long-run ratio between type-1 and type-0 cells in the population, both in the mean and on a sample-path basis. The ratio β_1/β_0 is the long-run ratio between the expected total sizes of clones started by a single type-1 cell vs. a single type-0 cell. We can determine these ratios explicitly as $\alpha_1/\alpha_0 = \gamma$ and $\beta_1/\beta_0 = 1/\delta$ with γ and δ given by (E.7).

E.2 Two-type skeleton

When the tumor does not die out, we can decompose the cells alive at time t into skeleton cells, i.e. cells with an infinite line of descent, and finite-family cells (Section 3.2.2). Let $\widehat{Z}_i(t)$ denote the number of type-i skeleton cells at time t and set $\widehat{\mathbf{Z}}(t) := (\widehat{Z}_0(t), \widehat{Z}_1(t))$. It follows from the same argument as in Section 3.3.1 that type-i skeleton cells divide into one skeleton cell and one finite-family cell at rate $2r_ip_i$ and into two skeleton cells at rate r_iq_i. Note that in this case, $2r_ip_i$ and r_iq_i are not $2d_i$ and λ_i, respectively, since p_i is no longer given by d_i/r_i. By a similar argument, a type-0 skeleton cell switches to a type-1 skeleton cell at rate $\mu_0(1-p_1)/(1-p_0) = (q_1/q_0)\mu_0$, and it switches back at rate $\mu_1(1-p_0)/(1-p_1) = (q_0/q_1)\mu_1$. This implies that the skeleton $(\widehat{\mathbf{Z}}(t))_{t\geq 0}$ is also a two-type branching process with no deaths and infinitesimal generator

$$\widehat{\mathbf{A}} = \begin{bmatrix} r_0q_0 - (q_1/q_0)\mu_0 & (q_1/q_0)\mu_0 \\ (q_0/q_1)\mu_1 & r_1q_1 - (q_0/q_1)\mu_1 \end{bmatrix}.$$

We can simplify the diagonal elements by noting that (E.3) can be rewritten as

$$r_0p_0 = d_0 + \mu_0(p_1 - p_0)/(1-p_0),$$

which implies

$$r_0q_0 - (q_1/q_0)\mu_0 = r_0(1-p_0) - \mu_0(1-p_1)/(1-p_0) = \lambda_0 - \mu_0.$$

We can therefore rewrite the infinitesimal generator in the simpler form

$$\widehat{\mathbf{A}} = \begin{bmatrix} \lambda_0 - \mu_0 & (q_1/q_0)\mu_0 \\ (q_0/q_1)\mu_1 & \lambda_1 - \mu_1 \end{bmatrix}. \tag{E.10}$$

The diagonal elements equal those for the generator $\mathbf{A}$ of the original process $(\mathbf{Z}(t))_{t\geq 0}$, so the net rate at which type-i cells produce cells of the same type is the same as for the original process. However, the switching rates, which we denote $\widehat{\mu}_0 = (q_1/q_0)\mu_0$ and $\widehat{\mu}_1 = (q_0/q_1)\mu_1$, are modified based on a simple robustness ratio between the two types. For example, if type-1 cells are more likely to initiate successful clones, $q_1 > q_0$,

the skeleton switches more frequently into the more robust state, and it also stays there longer. In other words, the ratio $\widehat{\mu}_0/\widehat{\mu}_1$ between switching rates is amplified by a square factor $(q_1/q_0)^2$ over the original process. However, the multiple of the switching rates $\widehat{\mu}_0\widehat{\mu}_1 = \mu_0\mu_1$ is preserved from the original process, which implies that $\widehat{\mathbf{A}}$ and $\mathbf{A}$ have the same eigenvalues.

Let $\widehat{\phi}_i^{(j)}(t) := \mathbb{E}[\widehat{Z}_i(t)|\widehat{\mathbf{Z}}(0) = \mathbf{e}_j]$ be the mean number of type-i skeleton cells at time t when the process is started by a single type-j skeleton cell. Since the skeleton is a two-type branching process with infinitesimal generator $\widehat{\mathbf{A}}$, we can define $\widehat{\delta}$, $\widehat{\gamma}$ and the vectors $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ analogously to (E.6), replacing μ_1 by $\widehat{\mu}_1 = (q_0/q_1)\mu_1$, which satisfy

$$\widehat{\beta}_0 + \widehat{\beta}_1 = 1 \quad \text{and} \quad \widehat{\alpha}_0\widehat{\beta}_0 + \widehat{\alpha}_1\widehat{\beta}_1 = 1. \tag{E.11}$$

and

$$\widehat{\phi}_i^{(j)}(t) \sim \widehat{\beta}_j\widehat{\alpha}_i e^{\sigma t}, \quad t \to \infty.$$

We emphasize that since $\widehat{\mathbf{A}}$ and $\mathbf{A}$ have the same eigenvalues, the long-run growth rate σ is the same for the skeleton as for the original process. This is consistent with the single-type case, where we observed that the skeleton had the same net growth rate λ_0 as the original process (Section 3.3.1). As for the original process, there exists a random variable $\widehat{W}$ so that

$$\lim_{t\to\infty} \widehat{\mathbf{Z}}(t)e^{-\sigma t} = \widehat{W}\boldsymbol{\alpha} = (\widehat{W}\widehat{\alpha}_0, \widehat{W}\widehat{\beta}_0), \tag{E.12}$$

where in this case, $\mathbb{P}(\widehat{W} > 0|\widehat{\mathbf{Z}}(0) = \mathbf{e}_j) = 1$, i.e. the limiting random variable is positive with probability 1 since skeleton cells do not go extinct, and $\mathbb{E}[\widehat{W}|\mathbf{Z}(0) = \mathbf{e}_j] = \widehat{\beta}_j$.

To relate the skeleton vectors $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\boldsymbol{\beta}}$ to the original vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, we note that by (E.6), $\widehat{\delta} = (q_1/q_0)\delta$ and $\widehat{\gamma} = (q_1/q_0)\gamma$. It follows that

$$\begin{aligned}
\widehat{\alpha}_1/\widehat{\alpha}_0 &= (q_1/q_0)\gamma = (q_1/q_0)(\alpha_1/\alpha_0),\\
\widehat{\beta}_1/\widehat{\beta}_0 &= (q_0/q_1)(1/\delta) = (q_0/q_1)(\beta_1/\beta_0).
\end{aligned}$$

Not surprisingly, if $q_1 > q_0$, the long-run ratio between the two cell types shifts to the

more robust type-1 phenotype on the skeleton relative to the original process. It is less intuitive that the long-ratio between clones started by type-1 and type-0 cells is smaller than for the original process. This has to do with the conditioning on nonextinction. The fact that type-0 clones are less robust in the original process means that the conditioning skews the clone-size distribution toward abnormally large clones. Conversely, the sizes of the more robust type-1 clones will more represent typical sizes for the unconditioned process. Importantly, since changes to the ratios $\widehat{\alpha}_1/\widehat{\alpha}_0$ and $\widehat{\beta}_1/\widehat{\beta}_0$ relative to the original process are complementary,

$$(\widehat{\alpha}_0\widehat{\beta}_0)/(\widehat{\alpha}_1\widehat{\beta}_1) = (\alpha_0\beta_0)/(\alpha_1\beta_1).$$

which, using $\widehat{\alpha}_0\widehat{\beta}_0 + \widehat{\alpha}_1\widehat{\beta}_1 = 1$ by (E.11), implies

$$\widehat{\alpha}_0\widehat{\beta}_0 = \widehat{\alpha}_0\widehat{\beta}_0/(\widehat{\alpha}_0\widehat{\beta}_0 + \widehat{\alpha}_1\widehat{\beta}_1) = \alpha_0\beta_0/(\alpha_0\beta_0 + \alpha_1\beta_1) = \alpha_0\beta_0. \tag{E.13}$$

Similarly, $\widehat{\alpha}_1\widehat{\beta}_1 = \alpha_1\beta_1$. These remarkably simple relationships between the skeleton and the original process are a consequence of the simple structure of the skeleton generator $\widehat{\mathbf{A}}$.

E.3 Long-run skeleton spectrum

In the long run, the skeleton grows at an exponential rate σ on a sample-path basis by (E.12), with the cell types maintained at the ratio $\widehat{\alpha}_1/\widehat{\alpha}_0 = (q_1/q_0)\gamma$. We now analyze the SFS of neutral mutations on the skeleton assuming that the skeleton grows deterministically according to its long-run mean, $\widehat{Z}_i(t) = \widehat{\beta}_j\widehat{\alpha}_i e^{\sigma t}$, where the process is started by a single type-j cell. Let $\widehat{M}(t)$ be the expected number of mutations that accumulate on the skeleton in $[0, t]$ under this assumption. By the same calculations as in (3.6), the mutation rate per type-i skeleton cell per unit time is $w_i r_i$, so

$$d\widehat{M}/ds = w_0 r_0 \cdot \widehat{\beta}_j\widehat{\alpha}_0 e^{\sigma t} + w_1 r_1 \cdot \widehat{\beta}_j\widehat{\alpha}_1 e^{\sigma t}, \quad 0 \le s \le t.$$

which implies

$$\widehat{M}(t) = (\widehat{\beta}_j/\sigma)(\widehat{\alpha}_0 w_0 r_0 + \widehat{\alpha}_1 w_1 r_1) \cdot (e^{\sigma t} - 1), \quad t \geq 0. \tag{E.14}$$

Suppose a neutral mutation occurs in a skeleton cell at time t. The probability that this mutation occurs in a type-0 cell is

$$\widehat{\eta} := \frac{w_0 r_0 \cdot \widehat{\beta}_j \widehat{\alpha}_0 e^{\sigma t}}{w_0 r_0 \cdot \widehat{\beta}_j \widehat{\alpha}_0 e^{\sigma t} + w_1 r_1 \cdot \widehat{\beta}_j \widehat{\alpha}_1 e^{\sigma t}} = \frac{\widehat{\alpha}_0 w_0 r_0}{\widehat{\alpha}_0 w_0 r_0 + \widehat{\alpha}_1 w_1 r_1}.$$

If the mutation occurs in a type-i cell, its long-run frequency in the population is $(\widehat{\beta}_i/\widehat{\beta}_j)e^{-\sigma t}$. To see why, note that the type-i cell starts a clone that has size $\widehat{\beta}_i(\widehat{\alpha}_0 + \widehat{\alpha}_1)e^{\sigma(T-t)}$ at time $T > t$, while the total population has size $\widehat{\beta}_j(\widehat{\alpha}_0 + \widehat{\alpha}_1)e^{\sigma T}$ at time T. The long-run frequency f of a mutation occurring at time t is therefore

$$f = \widehat{\eta} \cdot (\widehat{\beta}_0/\widehat{\beta}_j)e^{-\sigma t} + (1 - \widehat{\eta}) \cdot (\widehat{\beta}_1/\widehat{\beta}_j)e^{-\sigma t} = \frac{1}{\widehat{\beta}_j} \cdot \frac{\widehat{\alpha}_0 \widehat{\beta}_0 w_0 r_0 + \widehat{\alpha}_1 \widehat{\beta}_1 w_1 r_1}{\widehat{\alpha}_0 w_0 r_0 + \widehat{\alpha}_1 w_1 r_1} \cdot e^{-\sigma t}. \tag{E.15}$$

Now let $\widehat{M}(f)$ denote the expected number of mutations with frequency $\geq f$ in the long run. Substituting into (E.14) and replacing the constant term 1 with the $t = 0$ case in (E.15), we obtain the spectrum

$$\widehat{M}(f) = \big(\widehat{\alpha}_0 \widehat{\beta}_0 \cdot w_0 r_0/\sigma + \widehat{\alpha}_1 \widehat{\beta}_1 \cdot w_1 r_1/\sigma\big) \cdot (1/f - 1), \quad 0 \leq f \leq 1.$$

Recalling that $\widehat{\alpha}_i \widehat{\beta}_i = \alpha_i \beta_i$ by (E.13), we can rewrite the result as

$$\widehat{M}(f) = \big(\alpha_0 \beta_0 \cdot w_0 r_0/\sigma + \alpha_1 \beta_1 \cdot w_1 r_1/\sigma\big) \cdot (1/f - 1), \quad 0 \leq f \leq 1. \tag{E.16}$$

To compare with the corresponding single-type result (3.37), define $\pi_i := \sigma/r_i$ and $\theta_i := \alpha_i \beta_i$ and write

$$\widehat{M}(f) = (\theta_0 \cdot w_0/\pi_0 + \theta_1 \cdot w_1/\pi_1) \cdot (1/f - 1), \quad 0 \leq f \leq 1, \tag{E.17}$$

where $0 < \theta_0, \theta_1 < 1$ and $\theta_0 + \theta_1 = 1$ are normalized weights by (E.7). Note that we have written the skeleton spectrum entirely in terms of parameters of the original process:

The birth rates r_i, the mutation rates w_i and the long-run growth parameters σ, $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$. Recall that for the single-type case (3.37), we had

$$\widehat{M}(f) = (w/q_0) \cdot (1/f - 1), \quad 0 \leq f \leq 1,$$

where $q_0 := \lambda_0/r_0$ was the survival probability of a single-cell derived clone. When the two types of our model have identical growth characteristics, $r_0 = r_1$ and $d_0 = d_1$, and there is no switching between types, $\mu_0 = 0$ and $\mu_1 = 0$, expression (E.17) reduces to (3.37), so our two-type spectrum is a generalization of the $1/f$ single-type result. Here, π_0 and π_1 can no longer be interpreted as probabilities since σ may be larger than r_0 or r_1. However, $w/\pi_i = wr_i/\sigma$ can still be interpreted as the "effective" mutation rate. The slope in (E.17) always lies between the rates w_0/π_0 and w_1/π_1, and the weights θ_0 and θ_1 depend on (i) the proportion of type-0 and type-1 cells in the population in the long run through α_i, and (ii) the long-run proportion of sizes of single-cell clones derived from each type through β_i, which together determine the frequency of an arbitrary mutation occurring at some time t. If mutations occur at continuous rates $\nu_0 > 0$ and $\nu_1 > 0$ throughout the lifetime of each cell, the result becomes

$$\widehat{M}(f) = (\theta_0 \cdot \nu_0/\sigma + \theta_1 \cdot \nu_1/\sigma) \cdot (1/f - 1), \quad 0 \leq f \leq 1, \tag{E.18}$$

which is a restatement of (E.16) with $\nu_i = wr_i$. Note that our assumption that the process was started by a single type-j cell does not enter anywhere into (E.17). As a result, (E.17) holds for any initial condition, when skeleton grows deterministically according to its long-run mean. If we think again in terms of the stochastic model, we can deduce that the $1/f$ spectrum in (E.17) should hold for any mutation that occurs once the skeleton has entered its long-run growth stable phase, i.e. to mutations that occur after the skeleton has reached a certain size.

E.4 Extension to more than two types

We conclude by deriving an extension of our $1/f$ result (E.17) to a switching model with more than two cell types. Here, the cell types are $0, 1, \ldots, k$ with $k \geq 2$, the growth parameters are r_i, d_i, λ_i with $\lambda_i > 0$, and the switching rates from type-i to type-j are

$\mu_{ij} \geq 0$. By the same reasoning as in Section E.2, the infinitesimal generator $\mathbf{A} = (a_{ij})$ is given by

$$a_{ii} = \lambda_i - \textstyle\sum_{j \neq i} \mu_{ij},$$
$$a_{ij} = \mu_{ij}, \quad i \neq j,$$

and the mean matrix $\mathbf{M}(t) = (m_{i,j})$ is again given by $\mathbf{M}(t) := \exp(t\mathbf{A})$. We assume that there exists $t > 0$ so that $m_{ij}(t) > 0$ for all $0 \leq i, j \leq k$, which ensures that any type is accessible from any other type via switching, possibly through an intermediate type. This condition was implicit in our assumption that $\mu_0 > 0$ and $\mu_1 > 0$ for the two-type model.

Let $\phi_i^{(j)}(t)$ be the mean number of type-i cells at time t when the process is started by a single type-j cell, and letσ be the largest eigenvalue of $\mathbf{A}$. As for the two-type model, there exist positive eigenvectors $\boldsymbol{\alpha} = (\alpha_i)$ and $\boldsymbol{\beta} = (\beta_i)$ with

$$\textstyle\sum_{i=0}^{k} \alpha_i = 1 \quad \text{and} \quad \sum_{i=0}^{k} \alpha_i \beta_i = 1 \tag{E.19}$$

so that

$$\phi_i^{(j)} \sim \beta_j \alpha_i e^{\sigma t}, \quad t \to \infty,$$

see Section V.7 of Athreya & Ney [6]. Let p_i denote the extinction probability of a clone started by a single type-i cell, and set $q_i := 1 - p_i$ as the survival probability. The generator $\widehat{\mathbf{A}}$ for the skeleton process is given by

$$\widehat{a}_{ii} = \lambda_0 - \textstyle\sum_{j \neq i} \mu_{ij},$$
$$\widehat{a}_{ij} = (q_j/q_i)\mu_{ij},$$

by the same argument as in Section E.2. To determine the eigenvalues of $\widehat{\mathbf{A}}$, we note that the characteristic polynomial for $\widehat{\mathbf{A}}$ can be written as

$$\det(\widehat{\mathbf{A}} - x\mathbf{I}) = \textstyle\sum_{\varphi} \operatorname{sgn}(\varphi) \widehat{b}_{0\varphi(0)} \widehat{b}_{1\varphi(1)} \cdots \widehat{b}_{k\varphi(k)},$$

where the sum is taken over permutations φ on $\{0, 1, \ldots, k\}$ and the factors in the sum are $\widehat{b}_{ii} = \widehat{a}_{ii} - x$ and $\widehat{b}_{ij} = \widehat{a}_{ij}$ for $i \neq j$. Write the characteristic polynomial for the

generator $\mathbf{A}$ of the original process in the same way as

$$\det(\mathbf{A} - x\mathbf{I}) = \textstyle\sum_{\varphi} \operatorname{sgn}(\varphi) b_{0\varphi(0)} b_{1\varphi(1)} \cdots b_{k\varphi(k)},$$

where $b_{ii} = a_{ii} - x$ and $b_{ij} = a_{ij}$ for $i \neq j$. Since each permutation φ can be written as a product of disjoint cycles, it follows from the structure of $\widehat{\mathbf{A}}$ that

$$\widehat{b}_{0\varphi(0)} \widehat{b}_{1\varphi(1)} \cdots \widehat{b}_{k\varphi(k)} = b_{0\varphi(0)} b_{1\varphi(1)} \cdots b_{k\varphi(k)}.$$

Indeed, if $\varphi(i) = i$, then $\widehat{b}_{i\varphi(i)} = b_{i\varphi(i)}$ since $\widehat{\mathbf{A}}$ and $\mathbf{A}$ agree along the diagonal, and if φ contains the cycle $(i_1, i_2, \ldots, i_r)$, then

$$\begin{aligned}
\widehat{b}_{i_1 i_2} \widehat{b}_{i_2 i_3} \cdots \widehat{b}_{i_k i_1} &= (q_{i_2}/q_{i_1}) \mu_{i_1 i_2} \cdot (q_{i_3}/q_{i_2}) \mu_{i_2 i_3} \cdots (q_{i_1}/q_{i_k}) \mu_{i_k i_1} \\
&= \mu_{i_1 i_2} \mu_{i_2 i_3} \cdots \mu_{i_k i_1} \\
&= b_{i_1 i_2} b_{i_2 i_3} \cdots b_{i_k i_1}.
\end{aligned}$$

We can then conclude that

$$\det(\widehat{\mathbf{A}} - x\mathbf{I}) = \det(\mathbf{A} - x\mathbf{I}),$$

so $\widehat{\mathbf{A}}$ and $\mathbf{A}$ have the same eigenvalues. In particular, the long-run growth rate of the skeleton is equal to the long-run growth rate σ of the original process, and we can write

$$\widehat{\phi}_i^{(j)} \sim \widehat{\beta}_j \widehat{\alpha}_i e^{\sigma t}, \quad t \to \infty,$$

where $\widehat{\phi}_i^{(j)}$ is the mean number of skeleton cells at time t when the process is started by a single type-j skeleton cell. As before, $\boldsymbol{\alpha} = (\widehat{\alpha}_i)_i$ and $\boldsymbol{\beta} = (\widehat{\beta}_i)_i$ are left and right eigenvectors of $\widehat{\mathbf{A}}$ corresponding to σ, normalized as in (E.19). It is easy to show that $(\alpha_i q_i)_i$ is a left eigenvector and $(\beta_i/q_i)_i$ is a right eigenvector of $\widehat{\mathbf{A}}$ corresponding to σ. Since σ is a simple eigenvalue of $\widehat{\mathbf{A}}$ by the Perron-Frobenius theorem, we can conclude that for $i \neq j$,

$$\widehat{\alpha}_i/\widehat{\alpha}_j = (q_i/q_j) \cdot (\alpha_i/\alpha_j) \quad \text{and} \quad \widehat{\beta}_i/\widehat{\beta}_j = (q_j/q_i) \cdot (\beta_i/\beta_j).$$

It follows that $(\widehat{\alpha}_i\widehat{\beta}_i)/(\widehat{\alpha}_j\widehat{\beta}_j) = (\alpha_i\beta_i)/(\alpha_j\beta_j)$ for $i \neq j$. We can then use the same argument as for the two-type case (Section E.2) to show that for all i,

$$\widehat{\alpha}_i\widehat{\beta}_i = \alpha_i\beta_i.$$

The above ingredients allow us to generalize our result from (E.17) to

$$M(f) = \left(\sum_{i=0}^{k} \theta_i \cdot w_i/\pi_i\right) \cdot (1/f - 1), \quad 0 \leq f \leq 1, \tag{E.20}$$

where $\pi_i := \sigma/r_i$ and $\theta_i := \alpha_i\beta_i$, and $0 < \theta_i < 1$ with $\sum_{i=0}^{k} \theta_i = 1$. If mutations occur at continuous rate $\nu_i > 0$ for type-i throughout the lifetime of the cell, the result becomes

$$M(f) = \left(\sum_{i=0}^{k} \theta_i \cdot \nu_i/\sigma\right) \cdot (1/f - 1), \quad 0 \leq f \leq 1. \tag{E.21}$$

a restatement of (E.20) with $\nu_i = wr_i$.