Extreme value statistics of mutation accumulation in renewing cell populations

Philip Greulich[1, 2] and Benjamin D. Simons[3, 4, 5]

[1] *Mathematical Sciences, University of Southampton, Southampton, UK*
[2] *Insititute for Life Sciences, University of Southampton, Southampton, UK*
[3] *Cavendish Laboratory, University of Cambridge, Cambridge, UK*
[4] *Wellcome Trust-CRUK Gurdon Institute, University of Cambridge, Cambridge, UK*
[5] *Cambridge Stem Cell Institute, University of Cambridge, Cambridge, UK*
(Dated: October 12, 2017)

The emergence of a predominant phenotype within a cell population is often triggered by a rare accumulation of DNA mutations in a single cell. For example, tumors may be initiated by a single cell in which multiple mutations cooperate to bypass a cell's defense mechanisms. The risk of such an event is thus determined by the extremal accumulation of mutations across tissue cells. To address this risk, we study the statistics of the maximum mutation numbers in a generic, but tested, model of a renewing cell population. By drawing an analogy between the genealogy of a cell population and the theory of branching random walks, we obtain analytical estimates for the probability of exceeding a threshold number of mutations and determine how the statistical distribution of maximum mutation numbers scales with age and cell population size.

PACS numbers: x

Over the lifetime of an organism, its constituent cells continuously accumulate DNA mutations, which can affect the pathways that control cell proliferation and survival. Yet, due to gene multiplicity or functional redundancy [1–5], disruptions of such pathways may often be tolerated within a homeostatic tissue cell population. Evidence from studies of the cancer genome [2, 6, 7] suggest that the accumulation of a critical number of individually "neutral" or "near-neutral" mutations may, in many cases, be necessary to trigger a selective survival advantage on cycling cells – a process called "genetic" or "epistatic buffering" [2–4, 8–10]. The resulting proliferative advantage of mutated cells confers clonal dominance [11, 12] which, if sustained long-term [13], constitutes a potential tumour-initiating event. Crucially, since one cell within a tissue cell population is sufficient to trigger such an event [14], the risk of this occurring is naturally dominated by the statistics of rare events – in this case the extreme accumulation of a multiplicity of mutations within a cell, rather than by the cell population averages reaching some level of mutational burden. The statistics of extreme mutation accumulation represents, therefore, a question of both academic and practical interest.

The normal maintenance of adult renewing tissue, such as the skin epidermis or the gut epithelium, relies on the activity of stem cells, which divide to replenish functional differentiated cells lost through exhaustion or cell death [15, 16]. In addition to asymmetric divisions, which leave the stem cell population unchanged [17], in most of these tissues the frequent, stochastic loss of stem cells is compensated be replacement via neighbors that divide symmetrically [18, 19]. It is on this background, that these long-lived cells acquire mutations that may lead, in turn, to a selective growth advantage.

Historically, efforts to model how the serial acquisition of mutations can drive tumor progression have focused predominantly on population means or have neglected the potential for epistatic buffering [20–23]. The impact of stochastic cell fate dynamics on the statistics of rare mutational signatures has remained under-explored. However, recently, numerical studies have shown how maintenance mechanisms reliant on stochastic stem cell self-renewal can protect cell populations from extreme mutational acquisition events [24]. These findings have been reinforced by analytical studies based on a specific model of tumor-initiation involving a "double-hit" [25]. However, the statistical basis of cancer risk on rare event phenomena in renewing tissues remains poorly defined. Here, we present a generic theory for how properties of the extreme mutation number distribution scale with age and cell number, and how this determines the risk of accumulating a critical number of mutations. In particular, we elucidate how drift dynamics of the renewing cell population moderates the strength of fluctuations, diminishing the frequency of rare events. Besides its relevance for assessing the risk of tumor initiation, this theory also generically elucidates how a predominant phenotype can emerge in a cell population (e.g. bacteria) via the epistatic cooperation of individually (near-)neutral mutations.

To model the long-term accumulation of mutations in a renewing cell population we consider a stochastic model closely related to the Moran process in population genetics [26]. In this model, cells replicate through division, acquire mutations and are lost stochastically while the total number of cells N is maintained constant (the condition of homeostasis). Therefore, we assign a fixed number of 'sites' $i = 1, ..., N$ to the cells, where a cell at site i is char-

acterized by mutation number, m_i. When a cell at site i is lost, at rate λ, another cell at a random site j simultaneously divides symmetrically, producing a copy with the same mutational signature, to replace the lost cell on site i. In addition, any cell i with m_i mutations can acquire stochastically an additional mutation at a constant rate μ. Note that in stem cell populations, asymmetric cell divisions, where one of the daughter cells commits to terminal differentiation and loss, leave the configuration of mutations across cells invariant, and so need not to be considered explicitly. Their potential to effect additional mutations through division is incorporated as the mutation rates are decoupled from the loss/replacement rate. For simplicity, we do not distinguish between the loci of mutations in the genome, an approximation that is valid for low net mutational burden. Furthermore, we consider the scenario before transformation into a hyperproliferative state, in which the mutations' effect on proliferation is neutral. The model dynamics can be written as the process

$$m_i \xrightarrow{\lambda} m_j \ , \tag{1}$$

$$m_i \xrightarrow{\mu} m_i + 1 \ , \tag{2}$$

where sites i and j are chosen randomly.

In the following, we will address the risk $\bar{P}_N(m_c, T)$ that at least one cell in a population of N cells acquires a critical number of mutations m_c after time $t = T$. This corresponds to the probability that the maximal mutation number across the population, $m^* := \max(m_1, ..., m_N)$, reaches m_c, which is related to the cumulative distribution function (CDF) of m^*, $P_N^*(m_c, T) := \mathrm{Prob}(m^*(T) \leq m_c) = 1 - \bar{P}_N(m_c, T)$. In particular, we will study the dependence of the CDF's median and mean on N and T.

Before addressing the dynamics of mutation accumulation on the background of stochastic cell loss/replacement, as a benchmark, we first consider the case $\lambda = 0$ in which cells accumulate mutations independently. In this case, the model describes N independently distributed Poisson processes and an expression for the extreme mutation number distribution can be determined straightforwardly. Although this expression does not admit a simple scaling form [27], the dependence on N can be well-approximated for large μT by normally distributed random variables with a mean and variance μT (see Supplemental Material). From this correspondence it follows that, at large N and μT, the difference between maximum mutation number m^* and the population mean, $\Delta m^* := m^* - \langle m_i \rangle$ has a CDF $P_N^*(\Delta m_c) = \mathrm{Prob}(\Delta m^* < \Delta m_c)$ which follows a Gumbel distribution,

$$P_N^*(\Delta m_c) = e^{-e^{-X}} \text{with } X = \frac{\Delta m_c - \tilde{m}}{\sigma_N} - \ln \ln 2 \ , \tag{3}$$
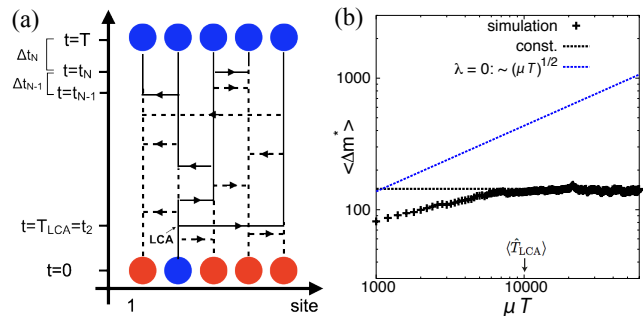

Figure 1. The genealogy and its implications. (a) Illustration of the history of mutation accumulation. Vertical lines represent mutational paths. Horizontal arrows mark loss/replacement events by which one cell copies its mutational configuration, in the direction of the arrow. Dashed lines are mutational paths that are lost, while bold lines are paths which survive until time $t = T$, constituting the genealogy. If T is large enough, the genealogy possesses a single root, the last common ancestor (LCA). (b) Mean difference between maximum mutation number and the population mean, $\langle \Delta m^* \rangle = \langle m^* \rangle - \langle m_i \rangle$, as a function of μT, for fixed $N = 10,000$, with $\lambda = \mu$. Points are results from Monte Carlo simulations; the dashed black line illustrates the saturation constant. The blue dashed line marks the scaling prediction for $\lambda = 0$, $\langle \Delta m^* \rangle \sim (\mu T)^{1/2}$.

with median $\tilde{m}$ and scaling width σ_N given by [28]

$$\tilde{m} \approx \sqrt{2\mu T \ln N}, \quad \sigma_N \approx \sqrt{\frac{\mu T}{2 \ln N}} \ . \tag{4}$$

The scaling estimate for the mean value $\langle \Delta m^* \rangle$ coincides with that of $\tilde{m}$ (see Supplemental Material). Thus, the CDF becomes narrowly peaked for large T and N around $\Delta m^* \sim (2\mu T \ln N)^{1/2}$.

In the case of a non-zero cell loss/replacement rate, $\lambda > 0$, any two cells may have a common ancestor and thus do not accumulate mutations independently. It is then instructive to consider the *genealogy* of the cell population, as illustrated in Fig. 1a. The genealogy describes the mutational history of all ancestors of cells at time $t = T$ and has the form of a binary tree, where branches connect daughter cells with their mothers [29]. It contains all mutational paths that start at $t = 0$ and reach the present. In considering the mutational statistics at time $t = T$, it is therefore sufficient to consider only mutations that occur on the genealogy [29, 30].

The tree structure of the genealogy is characterized by its branching times t_k, at which the branch number changes from $k - 1$ to k (see Fig. 1a), i.e. during the period $t_k < t < t_{k+1}$, the genealogy consists of k branches. The branching times can be determined by following the genealogy backwards in time $\hat{t} = T - t$; a coalescent process [30–32]. This results in branching times whose intervals $\Delta t_k := t_{k+1} - t_k$ are exponentially distributed, with $\mathrm{Prob}(\Delta t_k) = \langle \Delta t_k \rangle^{-1} e^{-\Delta t_k / \langle \Delta t_k \rangle}$ and mean branching

times (see Supplemental Material)

$$\langle \Delta t_k \rangle = \frac{N}{k(k-1)}\frac{1}{\lambda} \ .$$ (5)

Importantly, the accumulation of mutations along a single branch follows a simple, independent Poisson process [30] and thus the mean mutation number is simply $\langle m_i \rangle(T) = \mu T$.

For times T large enough, if one follows the coalescent process far backwards in time $\hat{t}$, the number of branches reduces until $\hat{t} = \hat{T}_{LCA} := T - t_2$, beyond which there is only a single one left, the *last common ancestor (LCA)* (see Fig. 1a). Thus, an LCA exists whenever $T > \hat{T}_{\mathrm{LCA}}$, which is on average $\langle \hat{T}_{\mathrm{LCA}} \rangle = \sum_{k=2}^{N}\langle \Delta t_k \rangle \approx N/\lambda$ (for $N \gg 1$). In that case, before the time $T_{\mathrm{LCA}} = T - \hat{T}_{\mathrm{LCA}} = t_2$, the genealogy corresponds to the mutational path of a single cell for which the maximum m^* equals the mutation number m. Hence, it follows that $\Delta m^* = m^* - \langle m_i \rangle > 0$ only for times larger than T_{LCA}, such that the statistics of Δm^* does not explicitly depend on the total time T if $T > \hat{T}_{\mathrm{LCA}}$. Indeed, Monte Carlo simulations of the model confirm this conjecture for $\langle \Delta m^* \rangle(T)$, as is illustrated in Fig. 1b for a high mutation rate $\mu = \lambda$, where a plateau is reached around $T \sim \langle \hat{T}_{\mathrm{LCA}} \rangle$. Note that this is in contrast to the case of purely asymmetric divisions, $\lambda = 0$, for which $\langle \Delta m^* \rangle \sim (\mu T)^{1/2}$.

To support this finding for $T > \hat{T}_{\mathrm{LCA}}$ quantitatively, we note that the branching times are random and exponentially distributed. Thus, the branching of the genealogy corresponds to a Markov process, a branching process with initially two branches at time T_{LCA} with branching rate *per branch* $\nu_k := 1/\langle \Delta t_k \rangle k$. By approximating the random accumulation of mutations along each branch (Poisson process) by diffusive random walks in the variables $m_i - \mu T$ (valid for $\mu T \gg 1$), the mutation accumulation of the genealogy becomes an unbiased *branching random walk (BRW)*. For unit branching rate, it has been shown [33] that the CDF of the maximum Δm^* of the BRW, $P_N^*(\Delta m_c, \tau) = \mathrm{Prob}(\Delta m^*(\tau) \le \Delta m_c)$, follows a Fisher-KPP-type equation [34, 35]

$$\partial_\tau P_N^* = D\frac{\partial^2 P_N^*}{\partial \Delta m_c^2} - P_N^*[1 - P_N^*] \ ,$$ (6)

with the dimensionless time $\tau = \nu t$ measured in units of the constant branching time ν^{-1}, and D, the diffusion constant of the random walk. The solution of this equation has the form

$$P_N^*(\Delta m_c, \tau) = f(\Delta m_c - \tilde{m}(\tau))$$ (7)

with the median of P_N^*, $\tilde{m}(\tau) = 2\sqrt{D}\,\tau + O(\ln \tau)$ [33].

Here, the branching rate ν_k is not constant. By a step-wise rescaling of time in units of branching times as $\tau(t) := \nu_k(t - t_k) + \tau_k$, with $\tau_k = \sum_{k'=2}^{k-1}\nu_{k'}\Delta t_{k'}$ for the largest $t_k < t$, $k > 2$, the corresponding BRW in the
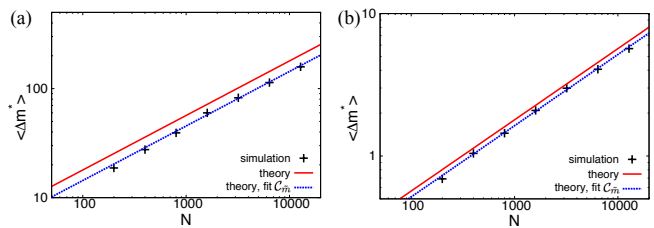

Figure 2. Mean maximum mutation number ahead of the mean, $\langle \Delta m^* \rangle$ as a function of N, for $T = 10\,N/\lambda$ such that $T > \hat{T}_{\mathrm{LCA}}$. Shown are the results of Monte Carlo simulations (pluses), and theoretical predictions from the BRW approximation, Eq. (8), for fitted numerical constant $\mathcal{C}_{\tilde{m}}^{\mathrm{fit}}$ (blue dashed line) and theoretically estimated value (solid red line), $\mathcal{C}_{\tilde{m}}^{\mathrm{th}} = 1.79$ for (a) $\mu = \lambda$ ($\mathcal{C}_{\tilde{m}}^{\mathrm{fit}} = 1.43$) and (b) $\mu = 0.001\lambda$ ($\mathcal{C}_{\tilde{m}}^{\mathrm{fit}} = 1.63$) .

time scale τ has unit branching rate and effective diffusion constant $D_k := \mu/2\nu_k$. While D_k does not explicitly depend on time, we can take the ensemble average over the branch numbers k, $D(\tau) := \langle D_k \rangle_k |_\tau = \mu N/(2\lambda) \times \langle (k-1)^{-1} \rangle_k |_\tau$, to get an effective time-dependent diffusion constant. According to Ref. [36] (see also Supplemental Material), for a diffusion constant $D(\tau')$ that decreases over time $\tau' < \tau = \tau(T)$, the CDF of the maximum of a BRW has the form of a Fisher-KPP wave, according to (7), but with

$$\tilde{m}(\tau) = \left[\int_0^\tau 2\sqrt{D(\tau')}\,d\tau'\right](1 - O(\tau^{-\frac{2}{3}})) \approx \mathcal{C}_{\tilde{m}}\sqrt{\frac{\mu N}{\lambda}} \ ,$$ (8)

where $\mathcal{C}_{\tilde{m}} = \int_0^\infty \sqrt{2\langle (k-1)^{-1} \rangle_k |_{\tau'}}\,d\tau'$ is independent of the model parameters N, T, λ and μ. Here, we took the limit $\tau(T) \to \infty$, which is valid for large N, as this is a unit rate branching process with branch number $k(T) = N \approx e^\tau$. Thereby, terms of $O(\tau^{-\frac{2}{3}})$ are omitted, and the integral becomes independent of N. A numerical evaluation (see Supplemental Material) yields $\mathcal{C}_{\tilde{m}} \approx 1.79$.

As expected, $\tilde{m}$ becomes independent of T if an LCA exists. The mean, $\langle \Delta m^* \rangle$, follows the same scaling in N and T, since due to the wave property of the CDF it only differs by a constant. This confirms our previous conjecture and the simulation results for $\langle \Delta m^* \rangle(T)$ (Fig. 1b). In Fig. 2, we compare the theoretical results from Eq. (8) with results for $\langle \Delta m^* \rangle$ from Monte Carlo simulations, as a function of N. Here, T was scaled with N ($T = 10\,N/\lambda$) to assure that $T > \hat{T}_{\mathrm{LCA}}$. The theory with fitted numerical constant $\mathcal{C}_{\tilde{m}}$ ($\mathcal{C}_{\tilde{m}}^{\mathrm{fit}}$, blue dashed line) shows excellent agreement with simulations, while the calculated value ($\mathcal{C}_{\tilde{m}}^{\mathrm{th}} = 1.79$, red line) shows some deviation. These deviations are expected due to contributions with small τ' in $\mathcal{C}_{\tilde{m}}$ (whose approximation is valid for large τ'). Remarkably, the predictions of the approximation are also valid for $\mu T \sim 1$ as shown in Fig. 2b for $\mu = 0.001\,\lambda$.

While the nonlinear form of the Fisher-KPP equation

does not admit an exact solution, the CDF's tail with $\bar{P}_N = 1 - P_N^*(\Delta m_c, \tau) \ll 1$ can be mapped onto a simple diffusion equation with time-varying diffusion constant (see Supplemental Material). Since variances add linearly in this case, for $T > \hat{T}_{LCA}$ and $\Delta m_c \gg \tilde{m}$, the CDFs tail is that of a non-normalized Gaussian function,

$$P_N^*(\Delta m_c) \sim 1 - \frac{N \sigma_{\text{eff}} \, e^{-\frac{\Delta m_c^2}{2\sigma_{\text{eff}}^2}}}{\sqrt{2\pi} \, \Delta m_c} \ , \tag{9}$$

with

$$\sigma_{\text{eff}}(\tau) = 2\sqrt{\int_0^\tau D(\tau') \, d\tau'} \approx \mathcal{C}_\sigma \sqrt{\frac{\mu N}{\lambda}} \ , \tag{10}$$

where $\mathcal{C}_\sigma \approx \sqrt{\int_0^\infty 2\langle(k-1)\rangle_k|_{\tau'} \, d\tau'} \approx 0.76$ (see Supplemental Material).

For $T < \hat{T}_{\text{LCA}}$ the population may not possess an LCA. In this case the genealogy fragments into k sub-genealogies $l = 1, ..., k$, each representing a subpopulation of n_l cells, while $N = \sum_{l=1}^k n_l$. Each sub-genealogy, however, can again be approximated by a branching random walk, with a CDF $P_l(\Delta m_c)$ having a Gaussian tail according to Eq. (9). Therefore, and since the subpopulations accumulate mutations independently from each other, the CDF of the whole population, $P_N^*(\Delta m_c)$, is approximated for large N by a Gumbel distribution according to Eq. (3) [28], scaled by σ_{eff}, and with effective number of independently distributed random variables, $k \approx N/\lambda T$ (see Supplemental Material). This CDF has then median and scaling width,

$$\tilde{m} \approx \mathcal{C}_\sigma \sqrt{2 \ln\left(\frac{N}{\lambda T}\right) \mu T}, \ \ \sigma_N \approx \mathcal{C}_\sigma \sqrt{\frac{\mu T}{2 \ln\left(\frac{N}{\lambda T}\right)}}, \tag{11}$$

where $\mathcal{C}_\sigma$ is according to Eq. (10). The same scaling applies to $\langle \Delta m^* \rangle$.

Thus, we expect the scaling $\Delta m^* \sim \sqrt{\ln N}$, as for $\lambda = 0$, but with a negative offset $-2\mathcal{C}_\sigma^2 \ln(\lambda T)\mu T$ under the square root. Figure 3 shows Monte Carlo simulations of $\langle \Delta m^* \rangle$ together with the theory, with fitted $\mathcal{C}_\sigma^{\text{fit}}$, which shows a good agreement in the shown range of N, for both large and small mutation rates, $\mu = \lambda$ and $\mu = 0.001\lambda$. Deviations from the theoretically approximated value $\mathcal{C}_\sigma^{\text{th}} = 0.76$ are expected, for the same reasons as for $\mathcal{C}_{\tilde{m}}$ before, and furthermore for very large N, since the extreme value distribution of Poisson variables differs from the Gaussian approximation in this limit.

Finally, we consider the risk of accumulating a critical mutation number m_c, $\bar{P}_N(m_c, T) = 1 - P_N^*(m_c, T)$. In Figure 4, results from stochastic simulations for $\bar{P}_N(T)$ are shown for $m_c = 6$ [20, 37] with parameters chosen to match physiological conditions of human epidermal stem cells [6, 7] (see figure caption), comparing a model of stochastic stem cell loss and replacement, $\lambda > 0$, with
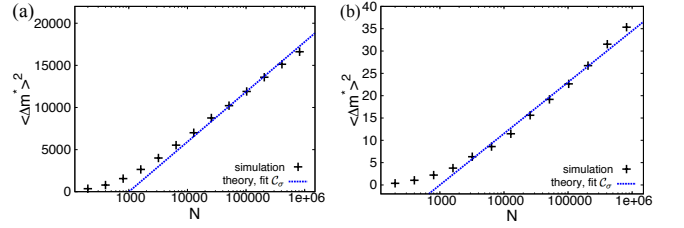

Figure 3. Squared mean maximum mutation number ahead of the population mean, $\langle \Delta m^* \rangle^2$ as a function of N, for fixed $T = 1000/\lambda$. Shown are the results of Monte Carlo simulations (pluses) and the theory, Eq. (11), with fit parameter $\mathcal{C}_\sigma^{\text{fit}}$ (dashed line) for (a) $\mu = \lambda$ ($\mathcal{C}_\sigma^{\text{fit}} = 1.13$) and (b) $\mu = 0.001\lambda$ ($\mathcal{C}_\sigma^{\text{fit}} = 1.59$).

the hypothetical case of asymmetric stem cell divisions, $\lambda = 0$. As the front of $P_N^*(m_c)$ moves with a speed $\mu + \partial_t \tilde{m}$, which is larger for $\lambda = 0$ (see Figure 1b), it reaches the critical value $m_c = 6$ earlier than for $\lambda > 0$, resulting in an earlier increase of $\bar{P}_N(m_c, T)$ in Figure 4a. Remarkably, the threshold of $m_c = 6$ mutations is already frequently exceeded for mean mutation numbers $\mu T < 1$, which demonstrates that the acquisition of a critical number of mutations is indeed dominated by extreme values. In Figure 4b, the *risk ratio* between the risk $\bar{P}_N|_{\lambda=0}$ for asymmetric divisions and that for symmetric divisions, $\bar{P}_N|_{\lambda>0}$, is shown. One observes a non-monotonic behavior: for intermediate times the ratio is large, while it approaches one for large times. The initial growth of the ratio is due to the smaller slope in the tail for $\lambda > 0$, $1/\sigma_N^{\lambda>0} \sim \sqrt{\ln(N/(\lambda T))}$ compared to the slope $1/\sigma_N^{\lambda=0} \sim \sqrt{\ln(N)}$ for $\lambda = 0$. Once $\bar{P}_N|_{\lambda=0}$ has saturated, however, $\bar{P}_N|_{\lambda>0}$ catches up with the latter until both approach $\bar{P}_N = 1$. This provides a theoretical foundation for claims that the risk of tumor initiation is decreased by stochastic stem cell loss and replacement [24, 25], yet this advantage is limited to intermediate time scales.

Until now, we have considered an "infinite-dimensional" process in which any cell may replace another. However, the majority of biological tissues are low-dimensional, where stem cell loss and replacement occur between neighboring cells in tubular (one-dimensional), epithelial (two-dimensional) or volumnar (three-dimensional) settings. Such situations can be modeled by embedding cells on a d-dimensional regular lattice, allowing replacement only between neighboring cells [38, 39]. In this case, our general theory remains valid, based on the mapping of the genealogy on a BRW; only the distribution of branching times, Δt_k, differs (see Supplemental Material). Nonetheless, only for $d = 1$ and $T > \hat{T}_{\text{LCA}}$, a significantly different scaling with $\tilde{m} \sim N\sqrt{\mu/\lambda}$ is observed, compared to the infinite-dimensional case.

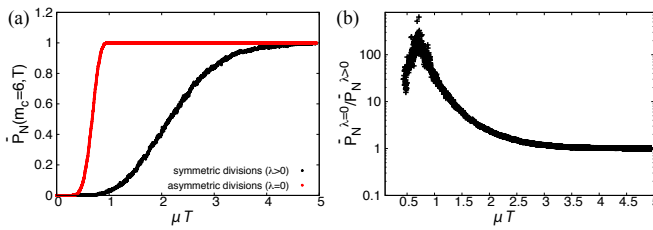In summary, we have studied the asymptotic behav-

Figure 4. Probability $\bar{P}_N(T)$ of accumulating a critical number of $m_c = 6$ mutations in an area of $0.1\,\mathrm{cm}^2$, corresponding to 10^5 cells. (a) $\bar{P}_N$ as function of time for physiological parameters of human eyelid [6, 7], with $\mu = 0.27/(63\,\mathrm{years})$ and $\lambda = 0.5/\mathrm{week}$, corresponding to $\lambda \approx 6000\mu$ (black points), and for the hypothetical case $\lambda = 0$ (red points). (b) The risk ratio $(\bar{P}_N)^{\lambda=0}/(\bar{P}_N)^{\lambda>0}$ for the two scenarios from (a).

ior, with time T and cell number N, of the maximum mutation number statistics in a renewing cell population, in which cells may be stochastically lost and replaced (Moran process). This is of importance if multiple neutral or near-neutral mutations can cooperate through epistasis to trigger hyper-proliferation, a potential tumor initiating event. We showed that, for a non-zero cell replacement rate, $\lambda > 0$, the difference between the average maximum mutation number and the population mean, $\langle \Delta m^* \rangle$, saturates to a constant value when scaled with T. Using an analogy to branching random walks, we showed that the value of this constant scales as $(\mu N/\lambda)^{1/2}$ for $\lambda > 0$, while in the absence of symmetric loss/replacement, $\lambda = 0$ (e.g. for asymmetric stem cell divisions only), $\langle \Delta m^* \rangle$ scales with time and cell number as $(\mu T \ln(N))^{1/2}$. If T is fixed, $\langle \Delta m^* \rangle$ scales as $(\mu T \ln(N/(\lambda T)))^{1/2}$ for $\lambda > 0$ and large $N \gg \lambda T$. From this result, it follows that at intermediate time scales, the risk of triggering hyper-proliferation is higher for asymmetric than for symmetric stem cell divisions, yet these probabilities converge for large times. Crucially, our theory also applies in a low-dimensional setting, where cell loss and replacement occurs between neighbors, albeit with a different scaling dependence for one-dimensional structures.

[1] A. G. Knudson, Proc. Natl. Acad. Sci. **68**, 820 (1971).
[2] A. Ashworth, C. J. Lord, and J. S. Reis-Filho, Cell **145**, 30 (2011).
[3] J. L. Hartman, B. Garvik, and L. Hartwell, Science **291**, 1001 (2001).
[4] J. A. G. M. de Visser, et al., Evolution **57**, 1959 (2003).
[5] Z. Gu, et al., Nature **421**, 63 (2003).
[6] B. D. Simons, Proc. Natl. Acad. Sci. **113**, 128 (2016).
[7] I. Martincorena, et al., Science **348**, 880 (2015).
[8] J. H. Moore, Nature Genetics **37**, 13 (2005).
[9] L. Jasnos and R. Korona, Nature Genetics **39**, 550 (2007).
[10] N. L. Komarova, Proc. Natl. Acad. Sci. **111**, 10789 (2014).
[11] L. Wagstaff, G. Kolahgar, and E. Piddini, Trends in Cell Biology **23**, 160 (2013).
[12] M. P. Alcolea, et al., Nature Cell Biology **16**, 615 (2014).
[13] S. Brown, et al., Nature (2017).
[14] G. M. Cooper, in *The Cell: A Molecular Approach. 2nd edition.* (Sinauer Associates;, 2000).
[15] L. Alonso and E. Fuchs, Proc. Natl. Acad. Sci. **100**, 11830 (2003).
[16] N. Barker, et al., Nature **449**, 1003 (2007).
[17] C. S. Potten, Cell Tissue Kinet. **7**, 77 (1974).
[18] E. Clayton, et al., Nature **446**, 185 (2007).
[19] C. Lopez-Garcia, A. M. Klein, B. D. Simons, and D. J. Winton, Science **330**, 822 (2010).
[20] P. Armitage and R. Doll, British Journal of Cancer **8**, 1 (1954).
[21] P. Armitage and R. Doll, British Journal of Cancer **11**, 161 (1957), 209.
[22] C. J. Portier, A. Kopp-Schneider, and C. D. Sherman, Mathematical Biosciences **135**, 129 (1996).
[23] I. Bozic, et al., Proc. Natl. Acad. Sci. **107**, 18545 (2010).
[24] P. T. McHale and A. D. Lander, PLoS computational biology **10**, e1003802 (2014).
[25] L. Shahriyari and N. L. Komarova, PLoS ONE **8**, e76195 (2013).
[26] P. A. P. Moran, Mathematical Proceedings of the Cambridge Philosophical Society **54**, 60 (1958).
[27] M. R. Leadbetter, G. Lindgren, and H. Rootzén, *Extremes and related properties of random sequences and processes* (Springer, New York, 1983).
[28] J.-P. Bouchaud and M. Marc, J. Phys. A: Math. Gen. **30**, 7997 (1997).
[29] J. Kingman, Journal of Applied Probability **19**, 27 (1982).
[30] R. R. Hudson, in *Oxford Surveys in Evolutionary Biology*, edited by D. Futuyama and J. Antonovics (1991), p. 1, 7th ed.
[31] J. Kingman, Stochastic Processes and their Applications **13**, 235 (1982).
[32] J. Wakely, in *Coalescent Theory: An Introduction* (W. H. Freeman, 2008).
[33] M. D. Bramson, Communications on Pure and Applied Mathematics **31**, 531 (1978).
[34] R. A. Fisher, Annals of Eugenics **7**, 355 (1937).
[35] A. Kolmogorov, I. Petrovskii, and N. Piscunov, Byul. Moskovskogo Gos. Univ. **1**, 1 (1937).
[36] M. Fang and O. Zeitouni, J. Stat. Phys. **149**, 1 (2012).
[37] I. P. M. Tomlinson, M. R. Novelli, and W. F. Bodmer, Proc. Natl. Acad. Sci. **93**, 14800 (1996).
[38] A. M. Klein, D. P. Doupé, P. H. Jones, and B. D. Simons, Phys. Rev. E **77**, 031907 (2008).
[39] A. M. Klein and B. D. Simons, Development **138**, 3103 (2011).