Punctuated mutagenesis promotes multi-step evolutionary adaptation in human cancers

Christopher Graser[1,2,3], Wenbo Wu[1], Cole Christini[1,4], Mia Petljak[5,6] and Franziska Michor[1,2,3,7,8,9*]

[1]Department of Data Science, Dana-Farber Cancer Institute, Boston, MA, USA

[2] Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA

[3]Department of Stem Cell and Regenerative Biology, Harvard University, Cambridge, MA, USA

[4]School of Computer Science, Carnegie Mellon University

[5]Department of Pathology, New York University School of Medicine, New York, NY, USA

[6]Laura and Isaac Perlmutter Cancer Center, New York University, New York, NY, USA

[7]Center for Cancer Evolution, Dana-Farber Cancer Institute, Boston, MA, USA

[8]The Eli and Edythe L. Broad Institute, Cambridge, MA, USA

[9]Ludwig Center at Harvard, Harvard Medical School, Boston, MA, USA

[*]Author for correspondence. Email: michor@jimmy.harvard.edu.

Abstract

The rate of acquisition of genomic changes in cancer has been the topic of much discussion, with several recent investigations finding evidence of punctuated evolution instead of gradual accumulation of such changes. Despite forays into the description and quantification of these punctuated events, the effects of such events on subsequent cancer evolution remain incompletely understood. Here we investigate how non-gradual mutagenesis affects the ability of tumor cells to acquire and retain fitness-enhancing adaptations. We find that punctuated mutagenesis significantly facilitates adaptation in scenarios where adaptation requires crossing a fitness valley, i.e. when multiple mutations are required which individually are maladaptive but jointly confer a fitness advantage. By increasing the probability that multiple mutations occur in close succession, punctuation increases the chance that mutants in a fitness valley mutate further to reach a fitness peak before going extinct. Analyzing data from The Cancer Genome Atlas, we find that tumors with signatures of APOBEC mutagenesis, which has been shown to proceed in episodic bursts, exhibit patterns consistent with higher rates of crossing fitness valleys. Lastly, we characterize how the interplay between this enhanced ability to cross fitness valleys and adaptation-limiting effects of clonal interference affects overall adaptability in complex fitness landscapes.

Introduction

Tumors evolve via the acquisition of randomly arising genetic and/or epigenetic alterations and their selection[1]. The rate at which tumor cells accumulate such genomic changes to produce potentially adaptive innovation is thus an important determinant of the capacity of tumor cells to adapt to diverse selection pressures[2,3], affecting their propensity to progress locally, to metastasize, or to evolve resistance to therapeutic intervention. Several methods to quantify mutation rates have been developed[4–6] and proxies for mutation rates such as tumor mutation burden often feature in prediction models of therapeutic outcomes[7,8].

Most existing methods deal with constant mutation rates[6,9], assuming gradual evolutionary change over time. However, accumulating evidence suggests that mutagenesis in cancer is fluctuating. Indeed, a recent study[10] demonstrated in vitro that mutations associated with DNA-editing activity of APOBEC cytidine deaminases[11] occur in episodic bursts, with more than 100-fold differences in the rate of APOBEC-associated mutations across otherwise identical cell culture replicates (Fig 1A). This episodicity observed in vitro aligned with patterns in previously published in vivo data investigating APOBEC mutagenesis in lung cancer[12,13], and data from intestinal crypts[14] later confirmed explicitly that this episodic pattern also occurs in patients. Other recent studies using genomic data to reconstruct evolutionary histories of tumors have found evidence of punctuated evolution across several cancer types[15–17]. These studies have demonstrated the existence of distinct phases of mutation bursts[17,18] (Fig 1B), challenging the prevailing paradigm of gradual emergence of mutant lineages (Fig 1C). Various mutagenic processes such as chromothripsis[19,20], chromoplexy[21] and other drivers of chromosomal instability[22] have been implicated in causing such punctuated patterns[23].

How punctuated mutagenesis affects the evolutionary dynamics of tumor cell adaptation across complex fitness landscapes remains incompletely understood. Addressing this question with existing experimental data is inconclusive, as only short time horizons are observed (single expansions[10]; time to a most recent common ancestor[17,18]), and as fitness differences between observed cell lineages are incompletely characterized. Mathematical modelling of human cancer genomics data, however, can elucidate the dynamics of adaptation during punctuated tumor evolution. Here, we set out to systematically investigate the evolutionary consequences of punctuated mutation acquisition during tumorigenesis using mathematical modeling and analysis of genomic data from The Cancer Genome Atlas (TCGA).

Results

Temporal clustering facilitates multi-step adaptation

We set out to investigate how *temporal clustering* of mutagenic events into distinct episodes affects the ability of a cell population to achieve multi-step evolutionary

adaptations. We considered two scenarios: tumor evolution via a single advantageous step, such as an oncogenic adaptation[24] ("1-step adaptations", Fig 1D), and accumulation of multiple mutations with synergistic fitness effects ("multi-step adaptations", Fig 1E). Indeed, the two hit-hypothesis[25] for tumor suppressor genes (TGSs) suggests that deactivation of one copy of a TSG can be inconsequential or maladaptive, while bi-allelic deactivation confers a selective advantage[26]. While some TSGs exhibit (context-dependent) haploinsufficiency[27–29], synergistic fitness effects for successive (epigenetic[27,28,30]) alterations in the same gene remain the prominent feature of TSGs. Analogously, there are ample examples of successive alterations across different genes acting synergistically[31–33]. High mutation rates are known to limit the rate of 1-step adaptations due to clonal interference[34–36] (Fig 1F), which has been highlighted as a potential consequence of punctuated cancer evolution[37]. The dynamics of multi-step adaptation have been studied extensively[26,38–44], elucidating two modes of evolution (Fig 1G): 'sequential fixation' refers to the scenario in which the second mutation only emerges after cells harboring the first mutant have taken over the population, while 'stochastic tunneling' refers to situations in which the second mutation arises sooner. In large populations, mutants with a selective disadvantage become vanishingly unlikely to reach fixation, so that stochastic tunneling becomes the dominant mode of evolution for multi-step adaptations.

To thus investigate how the rate of successful stochastic tunneling events – enabling multi-step adaptation – is affected by punctuated mutagenesis, we simulated selection dynamics of a population of cells which proliferate according to a Wright-Fisher process [45,46] and accumulate mutations according to a constant (Fig 1H) or temporally clustered (Fig 1I) rate. The Wright-Fisher process models evolution as successive, non-overlapping generations of constant size. Each new generation is populated by drawing with replacement from the cells in the previous generation, with probabilities proportional to their fitness (Fig 2A). Cells can accumulate mutations that change their fitness, i.e. increase their probability to survive to the next generation. We considered fitness landscapes in which two-step adaptations are the only way by which cells can increase their fitness, with fixed fitness ratios across subsequent two-step adaptations (Fig 2B,C).

Interestingly, for two-step adaptations for which the intermediate mutants that carry only one of the two required mutations have a strong selective disadvantage (Fig 2B), we found that the population undergoes substantially (3.55 times) more two-step adaptations when accumulating mutations in a temporally clustered rather than in a uniform way (Fig 2D). This effect arises because intermediate mutants in "fitness valleys" have a high chance of going extinct before acquiring the next mutation. During mutation bursts, acquiring the next mutation in time before the disadvantageous mutant goes extinct becomes more likely. Similarly, intermediate mutants are more

likely to emerge in a mutation burst. This temporal clustering of the likelihood of two succeeding mutation events increases the rate of successful two-step adaptations.

This effect also emerges for two-step adaptations that do not constitute proper fitness valleys, i.e. for which the intermediate mutant is not maladaptive (Fig 2E; 1.15-fold increase). Under stochastic selection dynamics, even mutants with a slight selective advantage (Fig 2C) have a high chance of going extinct due to random drift. Lineages acquiring sets of synergistic mutations, thus, often do so without prior fixation of each intermediate mutant – via stochastic tunneling[26]. In those regimes, temporal clustering of mutations therefore also increases the speed of adaptation, as confirmed with simulations (Fig S2A-D). Moreover, the proportion of such sets of synergistic mutations acquired via stochastic tunneling rather than sequential fixation also increases with temporal clustering (Fig S2E).

Our findings are not unique to models of constant population size, as confirmed with investigations of a branching process model (Methods, Fig S1) which yielded qualitatively similar results. Furthermore, our results also generalize to multi-step adaptations with arbitrarily many steps (SI1). To demonstrate this result, we compared a uniform mutation process with mutation rate μ to temporally clustered mutation processes in which mutations emerge exclusively during burst phases that constitute a fraction $\frac{1}{k}$ of the time, but at a k-fold increased rate, $k\mu$ (Fig S2A). We showed analytically that in the limit of the average mutation rate μ going to zero, the rate f_k of successful $(n+1)$ -step adaptations is k^n times higher in the temporally clustered mutation process than in the uniform mutation process. We validated this result with simulations of a Wright-Fisher process for two-step adaptations, $n = 1$, and a range of small values for μ (Fig S2F,G). As the mutation rate increases, this fold-increase $\frac{f_k}{f_1}$ becomes smaller. However, since the absolute rates f_k and f_1 are proportional to μ^{n+1} (SI1), the absolute effect of temporal clustering on the rate of fixations of mutants with multi-step adaptations increases in μ (Fig S2F). In light of the large fold-changes in mutation rates during a burst and before or after a burst[10,15,17,18], our findings suggest that the resulting effect on rates of multi-step adaptation is substantial.

Temporal clustering in an exploration-exploitation setting

The increased rate of valley crossing is driven by phases of high mutation rates, when the fold-increase in the chance of valley-crossing is larger than the fold-increase in the mutation rate. This effect improves adaptability while the fitness landscape offers scope for multi-step adaptation. However, in fitness landscapes which have global maxima, as the population adapts there is ever less scope for (further) multi-step adaptations, and excessive exploration of the landscape becomes costly.

To investigate the effect of temporal clustering in such an exploration-exploitation setting, we constructed a two-dimensional fitness landscape (Fig 3A) which is randomly re-drawn at regular intervals mimicking exposure to novel physical environments or drugs during cancer evolution and treatment (Methods). Between subsequent re-drawings, the population may move to a global fitness maximum where any further mutation decreases fitness. We considered mutation processes with different average mutation rates μ and different values of the clustering parameter k (Fig 3B) and investigated the average fitness of the cell population over long time horizons. As before, the clustering parameter k denotes the fold-increase of the mutation rate in burst phases relative to the average mutation rate μ. However, rather than fixing the out-of-burst mutation rate at zero, the burst duration is now held constant (Fig 3B) to control for differences in waiting time until a burst occurs.

We found that, in this scenario, relative to the uniform mutation process ($k = 1$) increasing k initially increases the average fitness of the population for any fixed μ (Fig 3C). Similarly, for a fixed k, increasing μ when starting at low μ-values increases the average fitness. However, moving further in either of these directions in the parameter space – towards increasing μ or towards increasing k – eventually brings about a decrease in average fitness (see SI2 and Fig 3D-F for detailed descriptions and plots of the dynamics). For any fixed k, average fitness peaks at intermediate μ, and the highest average fitness on the k-μ-plane is achieved by a $k > 1$. Maximizing average fitness, thus, involves temporal clustering even when there is a trade-off between exploration and exploitation.

Proxies for valley crossing and for temporal clustering found in patient data

We then sought evidence of this effect in patient data, leveraging whole exome sequencing (WES) data from TCGA. We reasoned that among tumors in which similar numbers of mutations had emerged, those tumors with more temporally clustered mutation processes in their evolutionary history would be more likely to have undergone a larger number of multi-step adaptations. We would, thus, expect the ratio of detectable indicators of multi-step adaptations, such as two deactivating mutations in a TSGs relative to the total number of mutations in a tumor sample, to correlate with indicators of fluctuations in the mutation rate history of a tumor such as those found for APOBEC mutagenesis.

Before considering the TCGA data, to confirm in silico that such a correlation would indeed emerge under biologically realistic parameters, we performed large-scale simulations of tumor growth from a single cell to up to realistically detectable cell numbers of 10^6–10^7 cells[47] (Methods). We recorded the fraction of mutations that emerged during burst phases, constructed a TSG deactivation score (Fig 4A,B) by counting the number of TSGs (modelled as mutations with synergistic fitness effects) with at least two mutations across the population and divided by the total number of mutations across the genome (Methods). Our simulation results confirmed that these

two readouts are correlated (Pearson correlation 0.42, $p<0.01$; Fig 4C), showing that in silico predictions about the relative ease of acquiring 2-step adaptations under mutation processes with more vs. less temporal clustering are well reflected in our score for TSG deactivation.

Next, to construct similar readouts from the TCGA data, we first performed SBS signature decomposition on each sample and computed the relative contribution of APOBEC associated mutation signatures (SBS2 and SBS13) vs all other mutation signatures[48] (Methods). Given the evidence for episodic APOBEC mutagenesis [10,14], we used this relative contribution as a score for temporal clustering. As a score for multi-step adaptations, a list of known TSGs (the TSGene 2.0 database[49]) was considered. For each TCGA sample, we counted the number of TSGs in this list which harbored at least two single nucleotide substitutions classified as a missense or nonsense mutation (Methods) normalized by the total number of single nucleotide substitutions in that sample. The TCGA WES data did not enable us to determine whether such mutations indeed deactivated different copies of a TSG, rather than appearing on the same copy. If the likelihood of appearing on the same copy varied with the contribution of the APOBEC signature, this association could confound our results. Since mutations that are interdependently generated on the same copy of a gene would likely appear in close spatial proximity[11,50,51], we thus verified in a robustness check that our results remain unchanged when filtering out mutations with close-by genomic coordinates (Methods).

We then ranked the different tumor types in TCGA by the average contribution of SBS2 and SBS13 relative to all detected mutation signatures, and for each tumor type, investigated the correlation between our proxies for multi-step adaptations and for temporal clustering. For the top four categories with the highest mean contribution of APOBEC-driven mutagenesis to all mutations, we observed a positive correlation between the two proxies. This correlation is significant at the 5% level (t-test on the Pearson correlation coefficient) for three out of four of these top four categories (Fig 4D, further results in Fig S3).

Consistent with the hypothesis that this correlation arises because TSG deactivation occurs more readily with APOBEC mutagenesis, we found that mutations which contribute to the TSG-deactivation score on average have a strongly increased likelihood of being caused by APOBEC (Methods, Fig 4E). This finding is consistent across cancer types and independent of whether probabilities are averaged over samples or over individual mutations, suggesting that this result is not driven by outlier samples with many TSG deactivations and many APOBEC associated mutations (Methods, Fig 4F).

An alternative score for multi-step adaptations, in which the list of TSGs is replaced with a list of pairs of genes with synergistic fitness effects[31](Methods) again exhibits a positive correlation with the relative contribution of APOBEC associated mutation signatures in all categories; this correlation is significant for two of these

categories (Fig S4). Taken together, these results show that tumors with a larger share of APOBEC-associated mutations are enriched for pairs of mutations with synergistic fitness effects.

Stochastic Tunneling vs Clonal Interference during Punctuated Evolution

In tumors evolving at high mutation rates, clonal interference may limit the rate at which a cell population manages to acquire and retain 1-step adaptations (Fig 1F). Previous literature has quantified this limiting effect as a function of the mutation rate and the distribution of fitness values of emerging mutants[52–54]. However, the effects of non-constant mutation rates in this setting have not been elucidated. We thus set out to investigate clonal interference in the setting of temporally clustered mutation rates (Fig 5A).

We first considered a setting in which cells can only acquire 1-step adaptations and measured the rate at which novel 1-step adaptations reach fixation in the population across simulations with different clustering parameters k (Fig 5B). We found that fixation rates quickly decrease as k increases (Fig 5C). This pattern emerges because the extent to which clonal interference reduces fixation rates disproportionately increases with the mutation rate; for a fixed average mutation rate temporal clustering thus increases the effects of clonal interference.

These findings stand in contrast to our results for 2-step adaptations whose rate of fixation increases when temporal clustering is introduced. However, if in-burst mutation rates are sufficiently large such that multiple clones in the population may acquire a 2-step adaptation independently before one of these clones has fixated, effects of clonal interference also play a role for 2-step adaptations. Performing simulations for 2-step adaptations, we found that fixation rates are non-monotone in k. While at low k increasing k leads to a steep increase in the fixation rate, this trend eventually levels off and becomes negative, with further increases in k leading to a decrease in the fixation rate.

Having observed that temporal clustering increases effects of clonal interference but also facilitates stochastic tunneling, we next set out to investigate the relative contribution of these two effects on overall adaptation rates. We considered a setting in which cells can increase their fitness through both 1-step and 2-step adaptations (Fig 5D).

We found that 1-step adaptations are acquired more quickly under the uniform mutation process (k=1), while 2-step adaptations are acquired more quickly under a temporally clustered mutation process (k=5; Fig 5E). The relative magnitude of both effects varies with time. As 1-step adaptations spread more readily in the population, differences in the speed at which they fixate manifest early in the dynamics. Over time, cells gradually exhaust the possibilities for 1-step adaptations and the difference between the adaptation rate in both processes becomes dominated by differences in

the speed of acquiring two-step adaptations. In this latter phase we observe large differences between the two mutation processes in the expected time until any given proportion of the possible two-step adaptations have spread in the population. Analogous time differences for reaching fixed proportions of one-step adaptations in the initial phase of the dynamics are much smaller (Fig 5E).

We also investigated how this pattern depends on the ruggedness of the fitness landscape (Fig S5A). To this end, we varied the relative proportions of loci with and without fitness valleys and quantified adaptability differences. We found that the duration of the initial phase in which the uniform process yields higher adaptability varies with the ruggedness of the fitness landscape: this phase is roughly three times as long when the proportion of loci with fitness valleys is 5% compared to when it is 95% (Fig S5B-D). However, over this range of proportions, we consistently observed that this initial phase only accounts for a small fraction of the time it takes to acquire all adaptations.

Discussion

Accumulating evidence suggests that mutagenesis in tumor cell populations proceeds in punctuated bursts rather than gradually; however, the effects of such punctuation on the ability of a tumor cell population to acquire and retain fitness-enhancing adaptations remains incompletely understood. Here we set out to investigate the evolutionary dynamics of these processes using mathematical modeling and genomics data analysis of human tumors.

We found that when a population acquires multi-step adaptations via stochastic tunneling, the rate of adaptation is substantially enhanced by punctuation. Stochastic tunneling is the predominant mode of evolution in large populations that traverse fitness valleys. However, stochastic tunneling also occurs when the necessary mutation steps to reach a substantial fitness advantage are not individually maladaptive, i.e. if steps do not constitute a fitness valley but confer a slight fitness advantage. Punctuated evolution therefore facilitates the accumulation of sets of mutations that jointly and synergistically confer a fitness advantage, while the fitness effect of carrying only a subset of these mutations can range from being strongly maladaptive to being moderately adaptive.

For the limit of low average mutation rates we show analytically that the rate of stochastic tunneling under a temporally clustered processes is k^n times larger than under the time-invariant process, where k is the fold-increase in the mutation rate relative to the average mutation rate during mutation bursts in the temporally clustered process, and n is the number of mutation steps that is required to exit the fitness valley. Moving towards higher average mutation rates, this fold-change in the

stochastic tunneling rates decreases, but the absolute difference between the rates increases.

While temporal clustering facilitates multi-step adaptations, it also impedes 1-step adaptations due to clonal interference. We examined the interplay between these contrasting effects in a setting in which cells can acquire both types of adaptations. We showed that the relative importance of both effects varies over time. Since 1-step adaptations tend to be acquired more readily, clonal interference matters most in the early phase of the process. As the share of possible 2-step adaptations relative to 1-step adaptations increases, differences in the tunneling rate become the more relevant determinant for the speed of adaptation. A uniform mutation rate, thus, makes a population more efficient at finding local fitness maxima. However, temporal clustering allows the population to more quickly move between local maxima, thus speeding up the search for a global maximum.

While our theoretical results are easily verified in simulation settings where we can choose a fitness landscape, applying these results to real data remains challenging because of the inherent complexities in determining fitness effects in biological systems. Estimating fitness effects of individual mutations requires either intricate experimental approaches[55] or large patient cohorts to control for differences in the (epi-)genetic background against which these mutations emerge; and such estimations become considerably harder when considering joint fitness effects of sets of multiple mutations.

To circumvent these challenges, we focused on TSGs as a representative set of gene pairs with synergistic effects and on one mutation process known to tend to fluctuate over time – APOBEC-driven mutagenesis. Consequently, we observed only a subset of groups of mutations with synergistic fitness effects in the data and possibly only a small fraction of the variability in mutagenic punctuation between different samples. Nevertheless, we find that these scores significantly correlate, which aligns with our model predictions.

Our results have implications for both mechanistic and statistical modeling of tumor evolution. A key quantity in mechanistic models such as those used to optimize treatment schedules[56] is the likelihood that resistance-conferring adaptations arise during therapy; our results show that this likelihood depends on the temporal dynamics of the mutation process. Analogously, in statistical models, incorporating measures of punctuated mutagenesis may improve the ability to predict treatment outcomes by accounting for how these temporal dynamics shape adaptability.

Methods

Simulations of a Wright-Fisher and a branching Process in an unbounded fitness landscape

In the simulations presented in Fig 2 and Fig S1, mutations occur between selection steps and are independent of divisions. In the Wright-Fisher process simulations (Fig 2 C) with a temporally clustered mutation rate, the out-of-burst mutation probability per cell and per update step (between two divisions) was set to 0.01 and was increased to 0.1 in bursts. For the corresponding uniform mutation rate trajectory, this probability was simply set equal to the total amount of mutations that occurred in the simulations under the temporally clustered mutation rate, divided by the total number of update steps.

Mutation rates in the branching model simulations (Fig S1) were chosen analogously. However, to achieve temporally equidistant mutation bursts in the branching process simulations, we scaled the burst duration and the time between bursts by the population size.

Simulations Wright-Fisher Process in exploration/exploitation setting

To produce the results shown in Fig 3, we ran agent-based stochastic simulations of a Wright-Fisher Process with 20 cells. Cells were characterized by their location on a discrete fitness landscape (a 30 by 30 two-dimensional lattice). The fitness values f associated with the positions on the lattice were drawn independently at random as $f = 0.5 + x^4$ where $x \sim U[0,1]$ is distributed uniformly between zero and one. The lower bound of 0.5 ensures that cells at all positions have a non-negligible probability of dividing.

Raising x to the fourth power creates a landscape with few peak-positions surrounded by many valley-positions with little variation in fitness.

In each selection step, one cell was randomly chosen to divide with probability proportional to its fitness, and replaced a cell which was drawn uniformly at random. Mutations occurred between consecutive selection steps, and the mutating cell was chosen uniformly at random amongst all cells in the population (independently of the preceding selection step). If a cell mutated, it would move to a location in the fitness landscape chosen uniformly at random from the set of (at most 8) locations in the Moore neighborhood of its current location.

The rate at which mutation events occurred was governed by the parameters μ and k. Mutation bursts lasted 10^2 division events and started every 10^3 division events. Fitness landscapes were redrawn every 50714 ($\approx 50000 + \frac{10^3}{7}$) division events, to periodically vary the relative timing of mutation bursts and re-drawings of the fitness landscape. In this manner, we avoid artifacts in our results caused by phase-alignment of bursts and re-drawings.

Larger-scale simulations

We simulated tumor evolution as a branching process, starting from a single unmutated cell, up to a randomly drawn target cell number between 10^6 and 10^7 cells. We assume a constant death rate, so that over the course of a simulation the probability that the next event is a death event rather than a division event remains fixed (at a value of 0.35). In case of a death event, one cell picked uniformly at random is removed from the population. In case of a division event, one cell is chosen to divide with probabilities proportional to its fitness.

We assume that the number of mutations per division follows a binomial distribution $Bin(10^5,\mu)$, and for each simulation we randomly draw a baseline mutation probability μ uniformly between 0 and $3{\cdot}10^{-5}$. In mutation bursts, this probability gets multiplied by a factor of 20 or 50, again chosen randomly for each simulation. The population enters a burst phase with probability $\frac{n}{10}$ per division, where n is the current population size, and exits a burst phase with probability$\frac{n}{3}$, so that in expectation bursts phases last 3 generations, and start every 10 generations. The first cell in a simulation is initialized to be in a burst with probability $\frac{3}{13}$ consistent with the expected time spent in bursts given those parameters.

The fitness effect of mutations is modeled as follows. The starting cell has a fitness of one. Each new mutation has an additive effect on the cells fitness. We assume that mutations occur uniformly at random across the genome ant that multi-step adaptations can occur in a fraction θ = 0.01 of the genome, roughly aligning with the ratio of the number of TSGs vs the total number of genes in the human genome. We subdivide this part of the genome into 200 TSGs which are all hit by a mutation with equal probability. For each individual TSG, the first mutation reduces a cell's fitness by 0.05. The second mutation increases fitness by 0.15, and all further mutations are fitness neutral. For the remaining $(1 - \theta)$ fraction of the genome, we assume that mutations are fitness-neutral with a probability of 0.9, and that fitness effects are otherwise drawn from a Gaussian with mean equal to −0.005 and standard deviation equal to 0.005.

To produce the results in Fig 4A-B, we keep track of all mutations that arise in a population, and remove all mutations with less than 1% variant allele frequency from the output, as those would be unlikely to be detected in WES. Moreover, we keep track of the fraction of mutations that a sample acquired during a mutation burst. For each simulation, we then count the number of TSGs for which there are at least two mutations found in the population and divide this by the total number of mutations in the sample.

Analysis of TCGA WES data

Whole-exome sequencing data was acquired from The Cancer Genome Atlas. We performed mutation signature decomposition using the *cosmic fit()* function in python from the package SigProfilerAssignment[57] version 0.1.8 with cosmic version 3.4.

Moreover, we used the mutation-level signature probabilities generated via SigProfilerAssignment to compute mean probabilities of signatures SBS2 and SBS13 for individual SBSs (Fig 4 D,E).

To compute our TSG deactivation score, we filtered the SNV data for missense and nonsense mutations in genes belonging to the TSGene 2.0 database[49]. For each sample we calculated the number of TSGs with at least two mutations and divided by the total number of SNVs.

Analogously, to construct our synergistic mutations score, we filtered for missense and nonsense mutations in frequently co-mutated gene pairs identified by Gu et al.[31] [1], and for each sample divided the number of pairs in this list with mutations in both genes by the number of SNVs. [29]This enrichment for co-occurrences of mutations in both genes in a pair suggests that mutations in the two genes tend to have a synergistic fitness effect. For each sample in the TCGA WES data we thus computed the number of gene-pairs from this list for which there is at least one non-synonymous single nucleotide substitution mutation in each of the two genes, and divided this number by the total number of single nucleotide substitutions in the sample.

As a robustness-check, we investigated whether the results change if we require a minimum distance between the genomic locations of the mutations in TSGs that we count to our TSG deactivation score. APOBEC has been linked to clustered mutagenesis through processes of kataegis and omikli[11,50,51]. If samples with higher APOBEC activity have higher rates of having multiple close-by mutations on the same allele, and if such a sample has multiple deactivating mutations in a TSG, these deactivations might be more likely to lie on the same allele compared to samples with lower APOBEC activity. Since we use appearances of more than one deactivating mutation as a proxy for bi-allelic deactivation of TSGs, such a pattern would bias our results. To account for that, we explored filtering TSG mutations in the data based different minimum-distance-thresholds ranging from 1 to 100 bp and found no impact on our results.

[1] Enrichment for co-occurrence of mutations in pairs of genes, relative to the product of the frequencies of each individual gene being mutated suggests that the two mutations have a synergistic fitness effect.

References

1. Merlo LMF, Pepper JW, Reid BJ, Maley CC. Cancer as an evolutionary and ecological process. *Nature Reviews Cancer 2006 6:12*. 2006;6(12):924-935. doi:10.1038/nrc2013

2. Yates LR, Campbell PJ. Evolution of the cancer genome. *Nature Reviews Genetics 2012 13:11*. 2012;13(11):795-806. doi:10.1038/nrg3317

3. Neinavaie F, Ibrahim-Hashim A, Kramer AM, Brown JS, Richards CL. The Genomic Processes of Biological Invasions: From Invasive Species to Cancer Metastases and Back Again. *Front Ecol Evol*. 2021;9:681100. doi:10.3389/FEVO.2021.681100

4. Araten DJ, Golde DW, Zhang RH, et al. A Quantitative Measurement of the Human Somatic Mutation Rate. *Cancer Res*. 2005;65(18):8111-8117. doi:10.1158/0008-5472.CAN-04-1198

5. Williams MJ, Werner B, Heide T, et al. Quantification of subclonal selection in cancer from bulk sequencing data. *Nature Genetics 2018 50:6*. 2018;50(6):895-903. doi:10.1038/s41588-018-0128-6

6. Werner B, Case J, Williams MJ, et al. Measuring single cell divisions in human tissues from multi-region sequencing data. *Nature Communications 2020 11:1*. 2020;11(1):1-9. doi:10.1038/s41467-020-14844-6

7. Galuppini F, Dal Pozzo CA, Deckert J, Loupakis F, Fassan M, Baffa R. Tumor mutation burden: From comprehensive mutational screening to the clinic. *Cancer Cell Int*. 2019;19(1):1-10. doi:10.1186/S12935-019-0929-4

8. Sha D, Jin Z, Budczies J, Kluck K, Stenzinger A, Sinicrope FA. Tumor mutational burden as a predictive biomarker in solid tumors. *Cancer Discov*. 2020;10(12):1808-1825. doi:10.1158/2159-8290.CD-20-0522

9. Altrock PM, Liu LL, Michor F. The mathematics of cancer: integrating quantitative models. *Nature Reviews Cancer 2015 15:12*. 2015;15(12):730-745. doi:10.1038/nrc4029

10. Petljak M, Alexandrov LB, Brammeld JS, et al. Characterizing Mutational Signatures in Human Cancer Cell Lines Reveals Episodic APOBEC Mutagenesis. *Cell*. 2019;176(6):1282-1294.e20. doi:10.1016/J.CELL.2019.02.012

11. Petljak M, Dananberg A, Chu K, et al. Mechanisms of APOBEC3 mutagenesis in human cancer cells. *Nature 2022 607:7920*. 2022;607(7920):799-807. doi:10.1038/s41586-022-04972-y

12. Jamal-Hanjani M, Wilson GA, McGranahan N, et al. Tracking the Evolution of Non–Small-Cell Lung Cancer. *New England Journal of Medicine*. 2017;376(22):2109-2121. doi:10.1056/NEJMOA1616288/SUPPL_FILE/NEJMOA1616288_DISCLOSURES.PDF

13. Lee JK, Lee J, Kim S, et al. ClonalHistory & genetic predictors of transformation into small-cell carcinomas from lung adenocarcinomas. *Journal of Clinical Oncology*. 2017;35(26):3065-3074. doi:10.1200/JCO.2016.71.9096/SUPPL_FILE/DS_2016.719096_2.XLSX

14. Wang Y, Robinson PS, Coorens THH, et al. APOBEC mutagenesis is a common process in normal human small intestine. *Nature Genetics 2023 55:2*. 2023;55(2):246-254. doi:10.1038/s41588-022-01296-5

15. Navin N, Kendall J, Troge J, et al. Tumor Evolution Inferred by Single Cell Sequencing. *Nature*. 2011;472(7341):90. doi:10.1038/NATURE09807

16. Sottoriva A, Kang H, Ma Z, et al. A Big Bang model of human colorectal tumor growth. *Nat Genet*. 2015;47(3):209. doi:10.1038/NG.3214

17. Gao R, Davis A, McDonald TO, et al. Punctuated copy number evolution and clonal stasis in triple-negative breast cancer. *Nature Genetics 2016 48:10*. 2016;48(10):1119-1130. doi:10.1038/ng.3641

18. Minussi DC, Nicholson MD, Ye H, et al. Breast Tumors Maintain a Reservoir of Subclonal Diversity During Expansion. *Nature*. 2021;592(7853):302. doi:10.1038/S41586-021-03357-X

19. Stephens PJ, Greenman CD, Fu B, et al. Massive genomic rearrangement acquired in a single catastrophic event during cancer development. *Cell*. 2011;144(1):27-40. doi:10.1016/J.CELL.2010.11.055

20. Zhang CZ, Spektor A, Cornils H, et al. Chromothripsis from DNA damage in micronuclei. *Nature 2015 522:7555*. 2015;522(7555):179-184. doi:10.1038/nature14493

21. Baca SC, Prandi D, Lawrence MS, et al. Punctuated evolution of prostate cancer genomes. *Cell*. 2013;153(3):666-677. doi:10.1016/J.CELL.2013.03.021

22. Drews RM, Hernando B, Tarabichi M, et al. A pan-cancer compendium of chromosomal instability. *Nature 2022 606:7916*. 2022;606(7916):976-983. doi:10.1038/s41586-022-04789-9

23. Davis A, Gao R, Navin N. Tumor evolution: Linear, branching, neutral or punctuated? *Biochimica et Biophysica Acta (BBA) - Reviews on Cancer*. 2017;1867(2):151-161. doi:10.1016/J.BBCAN.2017.01.003

24. Chakravarty D, Gao J, Phillips S, et al. OncoKB: A Precision Oncology Knowledge Base. *JCO Precis Oncol*. 2017;2017(1):1-16. doi:10.1200/PO.17.00011

25. Knudson AG. Mutation and Cancer: Statistical Study of Retinoblastoma. *Proceedings of the National Academy of Sciences*. 1971;68(4):820-823. doi:10.1073/PNAS.68.4.820

26. Iwasa Y, Michor F, Nowak MA. Stochastic tunnels in evolutionary dynamics. *Genetics*. 2004;166(3):1571-1579. doi:10.1534/GENETICS.166.3.1571

27. Paige AJW. Redefining tumour suppressor genes: exceptions to the two-hit hypothesis. *Cell Mol Life Sci*. 2003;60(10):2147. doi:10.1007/S00018-003-3027-6

28. Wang LH, Wu CF, Rajasekaran N, Shin YK. Loss of Tumor Suppressor Gene Function in Human Cancer: An Overview. *Cellular Physiology and Biochemistry*. 2019;51(6):2647-2693. doi:10.1159/000495956

29. Park S, Supek F, Lehner B. Higher order genetic interactions switch cancer genes from two-hit to one-hit drivers. *Nature Communications 2021 12:1*. 2021;12(1):1-10. doi:10.1038/s41467-021-27242-3

30. Issa JP. The Two-Hit Hypothesis Meets Epigenetics. *Cancer Res*. 2022;82(7):1167-1169. doi:10.1158/0008-5472.CAN-22-0405

31. Gu Y, Yang D, Zou J, et al. Systematic interpretation of comutated genes in large-scale cancer mutation profiles. *Mol Cancer Ther*. 2010;9(8):2186-2195. doi:10.1158/1535-7163.MCT-10-0022

32. Hsu MS, Yu JC, Wang HW, et al. Synergistic effects of polymorphisms in DNA repair genes and endogenous estrogen exposure on female breast cancer risk. *Ann Surg Oncol*. 2010;17(3):760-771. doi:10.1245/S10434-009-0802-0/METRICS

33. Deming DA, Leystra AA, Nettekoven L, et al. PIK3CA and APC mutations are synergistic in the development of intestinal cancers. *Oncogene 2014 33:17*. 2013;33(17):2245-2254. doi:10.1038/onc.2013.167

34. Gerrish PJ, Lenski RE. The fate of competing beneficial mutations in an asexual population. *Genetica 1998 102:0*. 1998;102(0):127-144. doi:10.1023/A:1017067816551

35. Wilke CO. The Speed of Adaptation in Large Asexual Populations. *Genetics*. 2004;167(4):2045-2053. doi:10.1534/GENETICS.104.027136

36. Fogle CA, Nagle JL, Desai MM. Clonal Interference, Multiple Mutations and Adaptation in Large Asexual Populations. *Genetics*. 2008;180(4):2163. doi:10.1534/GENETICS.108.090019

37. Sun R, Hu Z, Curtis C. Big Bang Tumor Growth and Clonal Evolution. *Cold Spring Harb Perspect Med*. 2018;8(5). doi:10.1101/CSHPERSPECT.A028381

38. Komarova NL, Sengupta A, Nowak MA. Mutation–selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. *J Theor Biol*. 2003;223(4):433-450. doi:10.1016/S0022-5193(03)00120-6

39. Nowak MA, Michor F, Komarova NL, Iwasa Y. Evolutionary dynamics of tumor suppressor gene inactivation. *Proceedings of the National Academy of Sciences*. 2004;101(29):10635-10638. doi:10.1073/PNAS.0400747101

40. Weinreich DM, Chao L. Rapid evolutionary escape by large populations from local fitness peaks is likely in nature. *Evolution (N Y)*. 2005;59(6):1175-1182. doi:10.1111/J.0014-3820.2005.TB01769.X

41. Weissman DB, Desai MM, Fisher DS, Feldman MW. The rate at which asexual populations cross fitness valleys. *Theor Popul Biol*. 2009;75(4):286-300. doi:10.1016/J.TPB.2009.02.006

42. Proulx SR. The rate of multi-step evolution in Moran and Wright–Fisher populations. *Theor Popul Biol*. 2011;80(3):197-207. doi:10.1016/J.TPB.2011.07.003

43. Ashcroft P, Michor F, Galla T. Stochastic Tunneling and Metastable States During the Somatic Evolution of Cancer. *Genetics*. 2015;199(4):1213-1228. doi:10.1534/GENETICS.114.171553

44. Van Egeren D, Madsen T, Michor F. Fitness variation in isogenic populations leads to a novel evolutionary mechanism for crossing fitness valleys. *Communications Biology 2018 1:1*. 2018;1(1):1-9. doi:10.1038/s42003-018-0160-1

45. Wright S. Evolution in Mendelian Populations. *Genetics*. 1931;16(2):97. doi:10.1093/GENETICS/16.2.97

46. Fisher R. On the dominance ratio. *Proc R Soc Edinb*. 1922;42:321-341.

47. Fischer BM, Olsen MWB, Ley CD, et al. How few cancer cells can be detected by positron emission tomography? A frequent question addressed by an in vitro study. *Eur J Nucl Med Mol Imaging*. 2006;33(6):697-702. doi:10.1007/S00259-005-0038-6/METRICS

48. Alexandrov LB, Nik-Zainal S, Wedge DC, et al. Signatures of mutational processes in human cancer. *Nature 2013 500:7463*. 2013;500(7463):415-421. doi:10.1038/nature12477

49. Zhao M, Kim P, Mitra R, Zhao J, Zhao Z. TSGene 2.0: an updated literature-based knowledgebase for tumor suppressor genes. *Nucleic Acids Res*. 2016;44(D1):D1023-D1031. doi:10.1093/NAR/GKV1268

50. Nik-Zainal S, Alexandrov LB, Wedge DC, et al. Mutational processes molding the genomes of 21 breast cancers. *Cell*. 2012;149(5):979-993. doi:10.1016/J.CELL.2012.04.024

51. Mas-Ponte D, Supek F. DNA mismatch repair promotes APOBEC3-mediated diffuse hypermutation in human cancers. *Nat Genet*. 2020;52(9):958-968. doi:10.1038/S41588-020-0674-6

52. Campos PR, Adami C, Wilke CO. Modelling Stochastic Clonal Interference. Published online 2004:21-38. doi:10.1007/978-3-642-18734-6_2

53. Campos PRA, De Oliveira VM. Mutational effects on the clonal interference phenomenon. *Evolution*. 2004;58(5):932-937. doi:10.1111/J.0014-3820.2004.TB00427.X

54. Park SC, Krug J. Clonal interference in large populations. *Proc Natl Acad Sci U S A*. 2007;104(46):18135-18140. doi:10.1073/PNAS.0705778104

55. Salehi S, Kabeer F, Ceglia N, et al. Clonal fitness inferred from time-series modelling of single-cell cancer genomes. *Nature*. 2021;595(7868):585-590. doi:10.1038/S41586-021-03648-3

56. McDonald TO, Cheng YC, Graser C, Nicol PB, Temko D, Michor F. Computational approaches to modelling and optimizing cancer treatment. *Nature Reviews Bioengineering 2023 1:10*. 2023;1(10):695-711. doi:10.1038/s44222-023-00089-7

57. Díaz-Gay M, Vangara R, Barnes M, et al. Assigning mutational signatures to individual samples and individual somatic mutations with SigProfilerAssignment. *Bioinformatics*. 2023;39(12). doi:10.1093/BIOINFORMATICS/BTAD756

58. Watson HW, Galton F. On the Probability of the Extinction of Families. *The Journal of the Anthropological Institute of Great Britain and Ireland*. 1875;4:138. doi:10.2307/2841222

Conflicts of Interest

M.P. is a shareholder in Vertex Pharmaceuticals and a compensated consultant for the GLG Network.
F.M. is a co-founder of and has equity in Harbinger Health, has equity in Zephyr AI, and serves as a consultant for both companies. She is also on the board of directors of Recursion Pharmaceuticals. She declares that none of these relationships are directly or indirectly related to the content of this manuscript.
The other authors declare no conflicts of interest.

Supplementary Information

Derivations of results on multi-step adaptation rates in the limit of infrequent mutations

We discussed differences in the rate of multi-step adaptations f_k as a function of a temporal clustering parameter k. Here we formalize this discussion, and derive the result presented in the main text.

We consider processes in which multi-step adaptations occur via stochastic tunneling rather than sequential fixation, i.e. processes in which mutants with a selective disadvantage are negligibly unlikely to reach fixation. We assume that mutations happen sufficiently infrequently, so that we can neglect valley crossing scenarios in which the same mutation in a multi-step adaptation sequence occurs multiple times. Additionally, we require that valley crossing happens solely via stochastic tunneling, i.e. that the probability that a maladaptive mutant fixates is vanishingly low.

We denote the average rate at which cells acquire mutations by μ, and compare dynamics under different mutation processes which we index by a temporal clustering parameter $k \geq 1$. Starting every d units of time, a process with parameter k undergoes a

mutation burst lasting $\frac{d}{k}$ units of time, during which mutations occur at a rate $k\mu$. Outside of bursts, the mutation rate is 0. We assume that the time $\frac{(k-1)d}{k}$ between subsequent burst phases is sufficiently long such that we can restrict attention to multi-step adaptations that occur during a single burst, and that the duration of each burst $\frac{d}{k}$ is sufficiently long relative to the time that it takes to cross a fitness valley, so that valley crossings, which fail because the burst phase ended but which would have been successful, have a negligible effect on f_k.

As stated in the main text, with this setup we can show that for any population dynamics process for which the above limits can be motivated, the fold-increase in the rate f_k at which $(n+1)$-step adaptations occur in a process with clustering parameter k relative to the uniform process approaches k^n.

$$\lim_{d\to\infty}\lim_{\mu\to 0}\frac{f_k}{f_1}=\lim_{\mu\to 0}\lim_{d\to\infty}\frac{f_k}{f_1}=k^n$$

Deriving this result for any selection process indeed becomes straightforward, if it can be motivated that the mutation rate does not affect the fate of any mutant lineage once the first mutant in this lineage has emerged.

Somewhat more formally, we index the steps in an n-step adaptation sequence by $i \in \{1,...,n\}$, and denote the expected size of the lineage descending from mutant i by $Y_i(t)$. This quantity $Y_i(t)$ of course depends on the specificities of the selection process, but for our derivations it suffices to only require that $Y_i(t)$ does not depend on the mutation rate.

Since all intermediate mutants have a selective disadvantage, and since we assume that we are in a parameter regime in which the likelihood that a disadvantageous mutant fixates is negligibly low, it follows that the integral $\int_0^\infty Y_i(t)dt$ converges.

$$\int_0^\infty Y_i(t)dt<\infty$$

In the limit of a low (and for now time-invariant) mutation rate μ, the probability that the i'th mutant produces a further mutant $P(i \to i+1)$ is simply the product of the mutation rate and this integral.

$$P(i\to i+1)=\mu\int_0^\infty Y_i(t)dt$$

The probability that the first mutant spawns a sequence of n+1 mutants can, thus, be written as follows.

$$P(1 \to n+1) = \prod_{i=1}^{n} \left(\mu \int_0^\infty Y_i(t)dt) \right)$$
$$= \mu^n \prod_{i=1}^{n} \left(\int_0^\infty Y_i(t)dt) \right)$$

Lastly, we use $\tilde{\mu}$ to denote the rate at which new mutants with only one mutation emerge. For a given selection process, this rate may vary with the population size. We arrive at the following formulation for the rate at which new advantageous mutants emerge f_1.

$$f_1 = \tilde{\mu} \cdot P(1 \to n+1)$$
$$= \tilde{\mu} \cdot \mu^n \prod_{i=1}^{n} \left(\int_0^\infty Y_i(t)dt) \right)$$

Finally, we can introduce our clustering parameter k to scale both u and $\tilde{u}$, and multiply by a factor of $\frac{1}{k}$ to arrive at the average rate of valley-crossing in the temporally clustered process with parameter k.

$$f_k = \frac{1}{k} \cdot (k \cdot \tilde{\mu}) \cdot (k \cdot \mu)^n \prod_{i=1}^{n} \left(\int_0^\infty Y_i(t)dt) \right)$$
$$= k^n \cdot \tilde{\mu} \cdot \mu^n \prod_{i=1}^{n} \left(\int_0^\infty Y_i(t)dt) \right)$$
$$= k^n \cdot f_1$$

One consequence of our assumption that none of the maladaptive intermediate mutants reaches fixation is that the composition of the population once the final mutant emerges in the limit of $\mu \to 0$ does not depend on the mutation rate. For population dynamics models with a constant population size in which the probability that an emerging mutant (absent further mutation events) reaches fixation only depends on this composition, such as the Wright-Fisher process, we can interpret $\frac{f_k}{f_1}$ therefore also as the ratio of the rates at which mutants with n+1 mutations fixate.

Analogously, in certain models of branching evolution in which the prospects of one branch do not depend on other co-evolving branches, such as the Galton-Watson process[58], whether the lineage of an emerging mutant with (n + 1) mutations survives is independent of the mutation rate in the limit of $\mu \to 0$. The ratio $\frac{f_k}{f_1}$ therefore also reflects the ratio of the rates at which surviving lineages with (n+1) adaptations arise in such models.

The above result suggests that the effect of temporal clustering on relative rates of valley crossing is substantial, and gets exponentially stronger the wider the fitness valley is (n). We validated this result with simulations of a Wright-Fisher process for a range of small values for μ (Fig S2 B-E). As we move away from the limit of rare

mutations, the relative effect of temporal clustering $\frac{f_k}{f_1}$ gets smaller, as will be discussed below. However, as the absolute rates f_k and f_1 are proportional to μ^{n+1} (Appendix), the absolute effect of temporal clustering on the rate of fixations of mutants with multi-step adaptations initially increases in μ (Fig S2 D).

Dynamics in exploration-exploitation setting

Our simulations show that for given μ the effect of increasing the clustering parameter k on the average fitness becomes negative at some k (Fig 3C). In the simulations in this section, this effect is driven by two main factors. Firstly, after a re-drawing, the population often is not in a local maximum of the fitness landscape, and hence may be able to increase its fitness by single mutations (1-step adaptations), without tunneling. Waiting for a burst to occur and thereby delaying such local explorations decreases average fitness. Second, once the population has found a peak in the fitness landscape, having very high mutation rates in bursts temporarily scatters cells to lower points in the landscape, and may even cause drift to points of lower fitness. Both of these effects become apparent when considering a representative snapshot of the simulations (Fig 3D). The misalignment of re-drawing and burst causes the population to remain in a valley for several divisions until the first burst occurs. Moreover, relative to similar snapshots of simulations with lower k (Fig 3E,F), once the population has reached a peak, the scattering during bursts at higher k leads to sharper decreases in fitness, and the population may even remain in a lower fitness point after a burst (Fig 3D).

This increased scattering can also be seen when going from $k = 1$ to $k = 5$ (Fig 3E,F). However, at $k = 1$, the population spends much time in local maxima, as large jumps in fitness only occur right after re-drawings, whereas at $k = 5$ jumps occur also long after the landscape was re-drawn and the population found a local maximum. These later jumps to higher points in the landscape correspond to tunneling events.

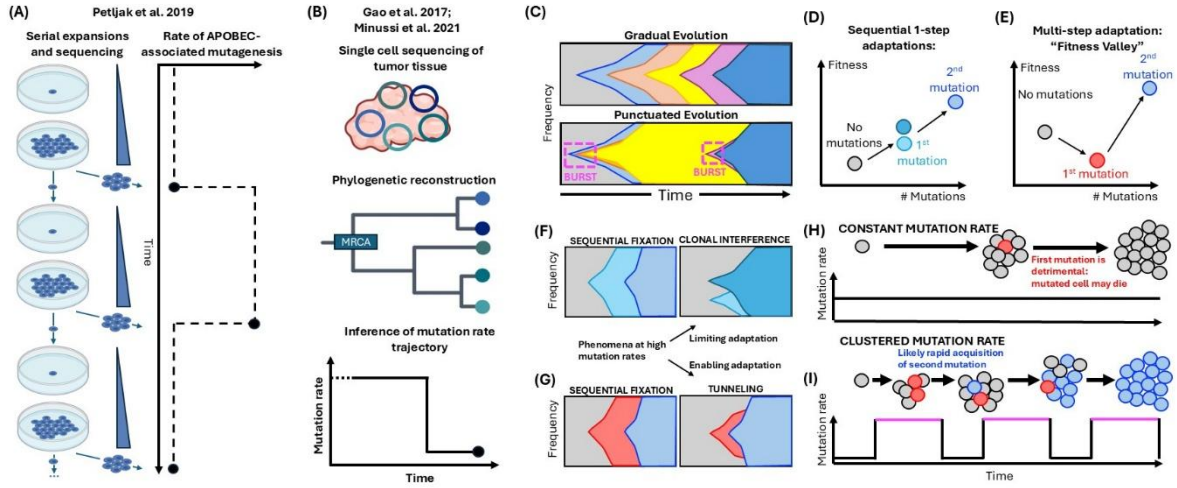

Figure 1 **Population dynamics under gradual vs punctuated mutagenesis**. **(A)** Evidence for punctuated APOBEC mutagenesis (Petljak et al. 2019) from successive cell culture expansions, seeded with single progenitors from the preceding expansion. Sequencing of the expanded populations revealed large fluctuations in APOBEC-associated mutagenic signatures SBS2 and SBS13. **(B)** Evidence for punctuated copy-number evolution (Gao et al. 2017, Minussi et al. 2021). Patterns in branch lengths of reconstructed phylogenetic trees reveal fluctuations in the rates of copy number alterations. Dynamics prior to most recent common ancestor (MRCA) are unidentifiable. **(C)** Gradual vs. punctuated evolution. In gradual evolution novelty emerges and spreads at a constant rate over time. In punctuated evolution novelty emerges and spreads during distinct burst phases. **(D)** Fitness schematic of two 1-step adaptations which each independently confer a fitness advantage. **(E)** Fitness schematic for a 2-step adaptation, in which carrying one mutation confers a fitness disadvantage but carrying two mutations confers a fitness advantage. **(F,G)** Schematic of possible evolutionary dynamics for two 1-step adaptations (F) and of modes of valley-crossing with or without prior fixation of the first mutant (G). Sequential fixation occurs at low mutation rates where emerging mutant lineages are likely to have fixated or gone extinct before the next mutation occurs. At higher mutation rates, the second mutation can occur in a multi-clonal population. **(H)** Sketch of failing two-step adaptation under a uniform (time-invariant) mutation rate. **(I)** Sketch of a successful two-step adaptation via stochastic tunneling under a punctuated mutation process with distinct clusters of high mutation rates.

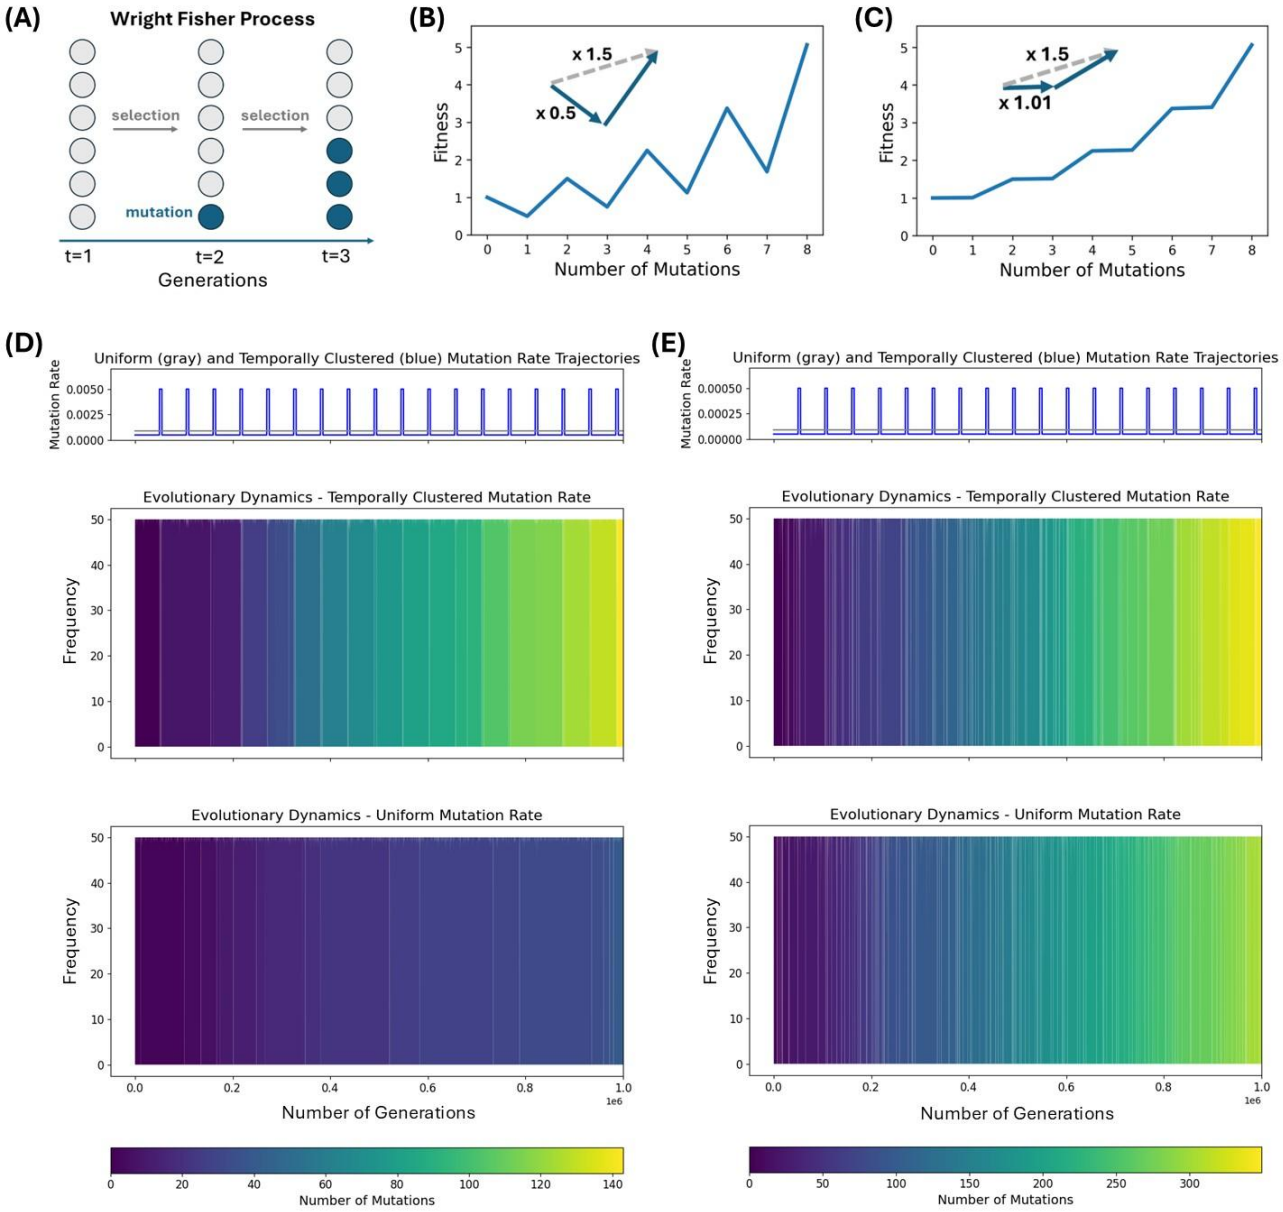

Figure 2 **Simulation results: valley-crossing under uniform vs. temporally clustered mutation rates. (A)** Schematic of a Wright Fisher Process. **(B)-(C)** Fitness landscapes used for the simulations in Panel (D) and (E) respectively. Mutations move a cell from left to right through the landscape. Each 2-step adaptation corresponds to a fitness increase by a factor of 1.5. Having an odd number of mutations comes at a multiplicative fitness disadvantage of 0.5 in (B) and an advantage of 1.01 in panel (D). **(D)-(E)** Simulation results for a Wright-Fisher process with 50 cells. The mutation rate trajectories in each panel are chosen such that the total expected number of mutations under the uniform trajectory is identical to that under the temporally clustered trajectory.

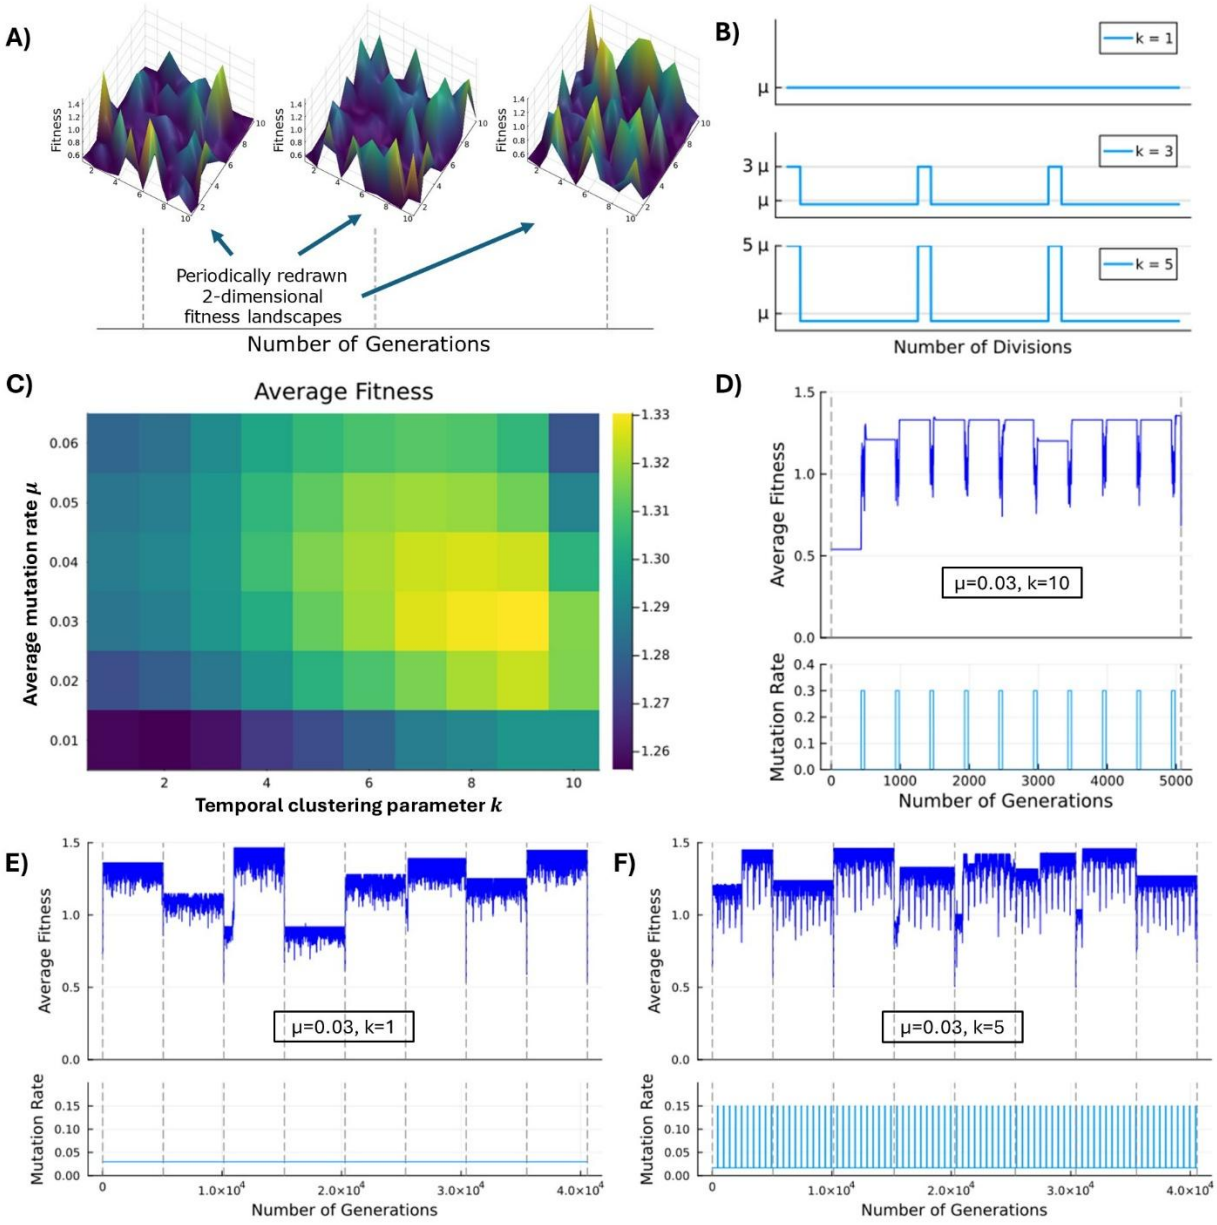

Figure 3 **Exploration and exploitation with temporally clustered mutation rates. (A)** Cells move through two-dimensional fitness landscapes. These fitness landscapes are randomly redrawn every 50714 divisions. **(B)** Mutation rate trajectories are parameterized with a mean mutation rate μ, and a clustering parameter k. **(C)** Average fitness in simulations of a population of 20 cells in a Wright Fisher Process. Simulations were run for 10^8 division events. **(D)-(F)** Average fitness trajectories for representative snippets of the simulations in **(C)**.

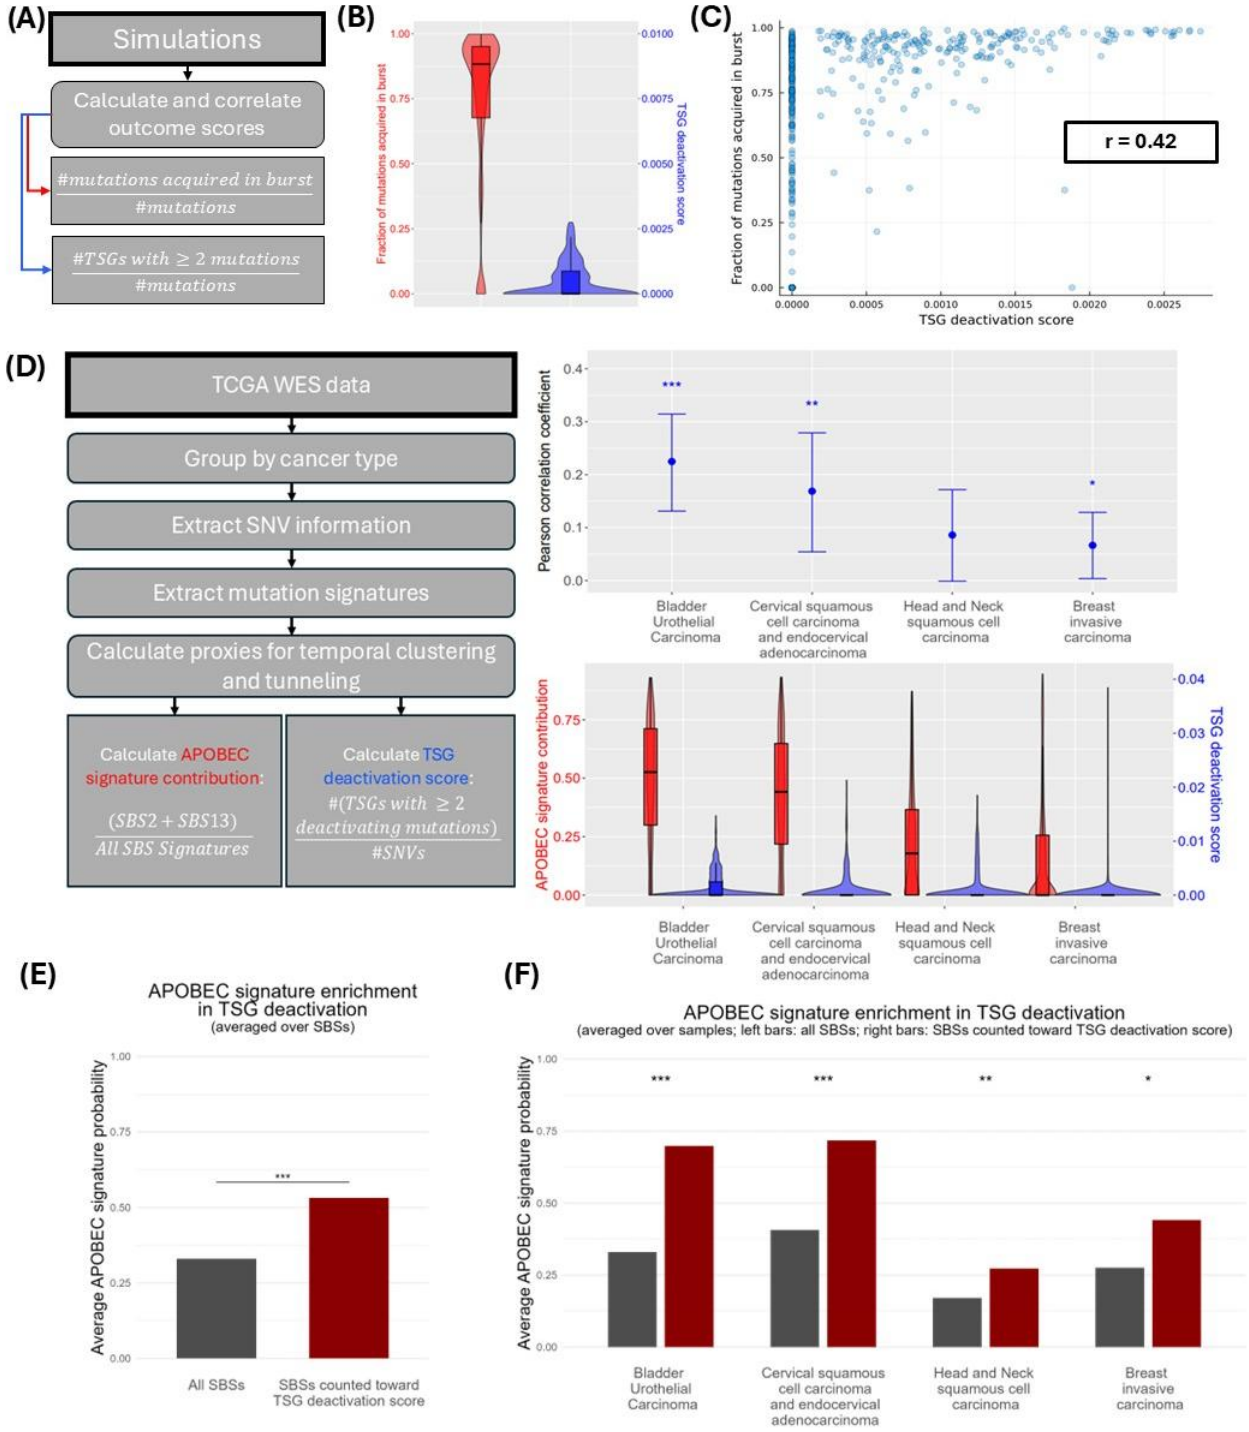

Figure 4 **Proxies for valley crossing and temporal clustering in simulations and in TCGA data. (A)** Sketch of the simulation analysis workflow. **(B)** Distributions of the fraction of mutations acquired during mutation bursts and the TSG deactivation score in simulations**.** **(C)** Joint distribution of the quantities in (B) indicates strong correlation. **(D)** Schematic of

analysis workflow for TCGA data (methods) and results for the four cancer types with highest APOBEC signature contribution. **(E)** Probabilities that single base substitutions (SBSs) in samples of the cancer types in (C) were caused by APOBEC associated signatures SBS2 and SBS13. On the left, these probabilities are averaged over all SBSs. On the right, only those SBSs are considered which are classified as “nonsense” or “missense”, and which appear in a TSG with at least 2 such deactivating SBSs. The average probabilities between both groups are significantly different (t-test, $p<0.001$). **(F)** Average probabilities analogous to those in panel (D) were constructed for each individual sample and then averaged across samples. Samples without deactivated TSGs were excluded. The average probabilities between groups are significantly different for the four cancer types (*, ** and *** respectively indicate that the p-value in a t-test lies below 0.05, 0.01 and 0.001).

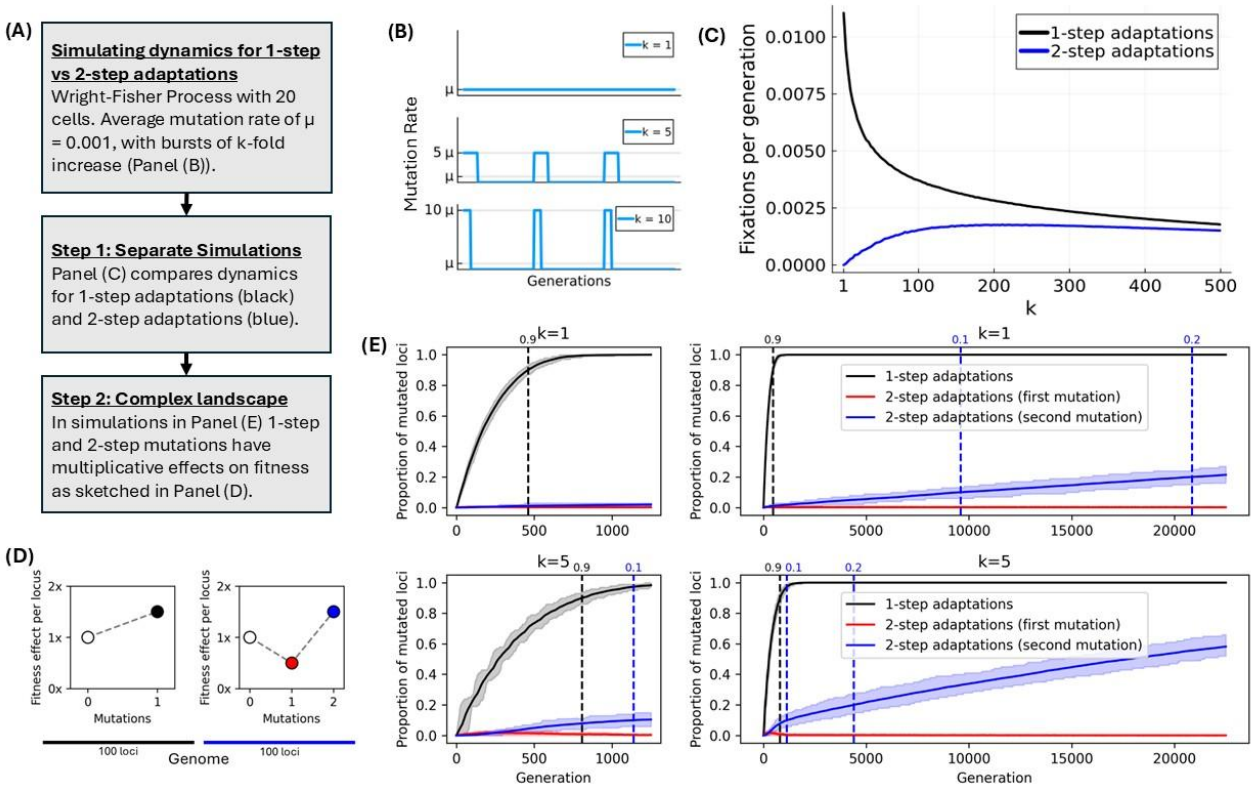

Figure 5 **Stochastic tunneling vs clonal interference under temporally clustered mutation rates: (A)** Schematic of the simulation approach. **(B)** Sketch of the mutation process. **(C)** Simulation results: measuring fixations of higher-fitness mutations per generation as a function of the clustering parameter and of the fitness landscape. In the 1-step adaptation setting fitness is defined as $1.5^{\#mutations}$. Fitness in the 2-step adaptation setting is defined as in Fig 2B. Results are averaged over simulation runs with 10^7 generations. **(D)** Sketch of fitness landscape for simulations shown in (E). Cells start with an unmutated genome of 200 loci. Mutations on each locus have independent multiplicative effects on cell fitness. In the first 100 loci a single mutation confers a multiplicative fitness change of 1.5. In the remaining 100, 2 mutations are required to reach this multiplicative fitness change of 1.5, with the first mutation conferring a multiplicative fitness change of 0.5. **(E)** Simulation results for adaptation with genome sketched in (D). Results are averaged over 100 simulation runs per value of k. Shaded regions indicate 10^{th} and 90^{th} percentiles. Smaller plots (left) are zoomed in versions of the first 1200 generations of the larger plots (right), with identical color-coding.

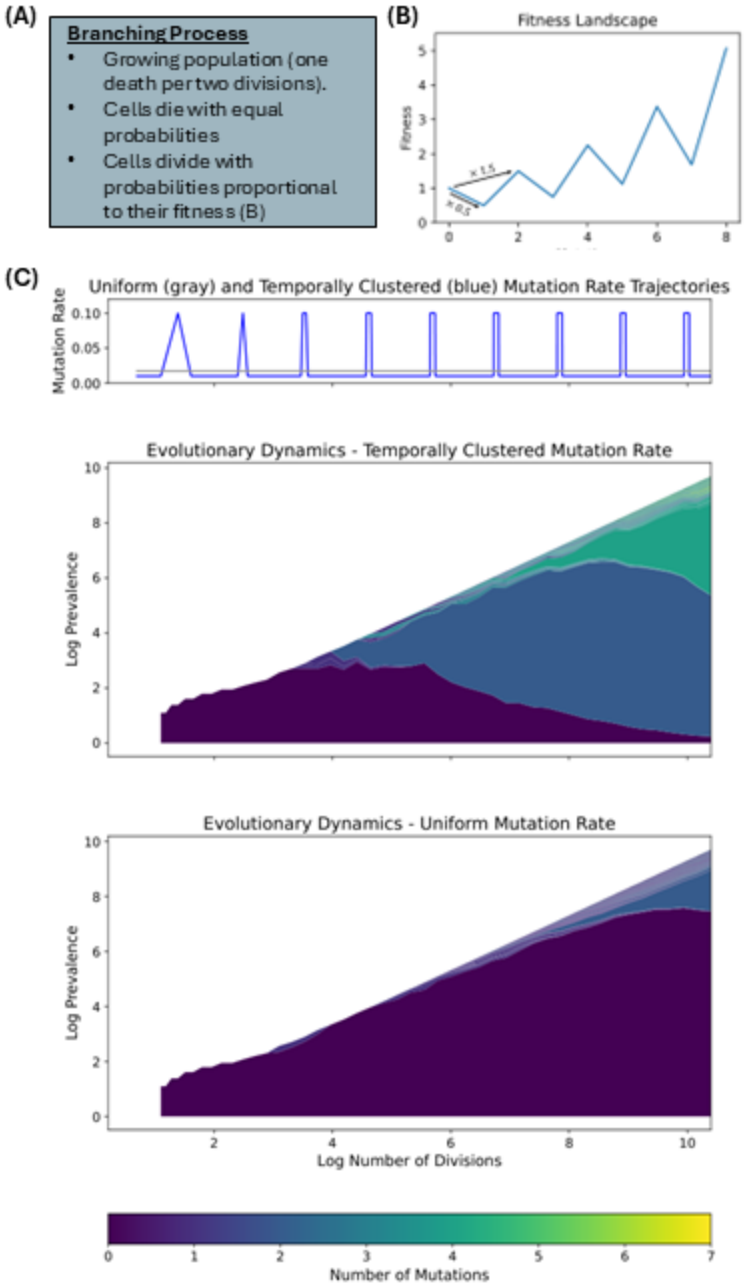

Supplementary Figure1: **Simulation results: valley-crossing under uniform vs. temporally clustered mutation rates in branching process. (A)** Schematic of fitness landscape as in Fig 2. **(B)** Schematic of Branching Process model. **(C)** Simulation results for the Branching Process model. The mutation rate trajectories in (C) are chosen such that the total expected number of mutations under the uniform trajectory is identical to that under the temporally clustered trajectory.

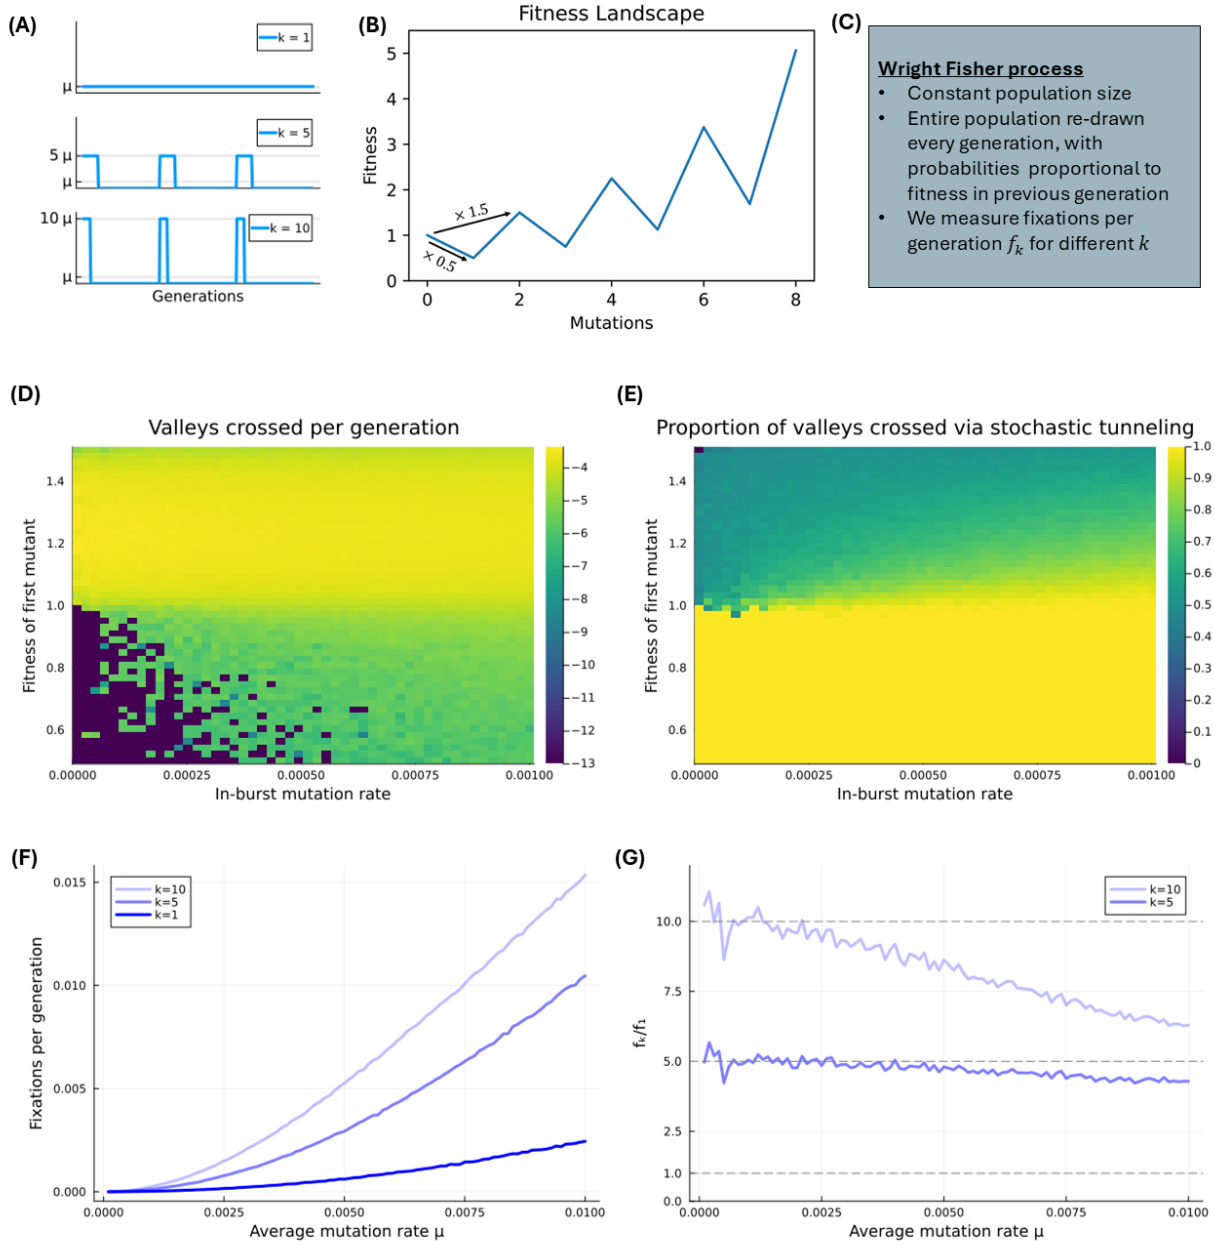

Supplementary Figure 2: **Effect of temporal clustering on stochastic tunneling rates as a function of the mutation rate. (A)** Sketch of the mutation process with clustering parameter k. **(B)** Sketch of the fitness landscape as in Fig 2. **(C)** Description of the simulations. **(D)-(E)** Simulation results of a Wright-Fisher process with 100 cells over 10^6 generations with a fixed average mutation rate of10^{-5} per cell per generation and varying temporal clustering k. The fitness is as in (B), with subsequent peaks differing in fitness by a factor of 1.5. However, on the y-axis in (F) and (G) we vary the fitness of the valleys relative to the preceding peak from 0.5 (as in (B)) to 1.5. Panel (D) shows the number of valleys crossed per generation (color bar in log10 values), where contrary to the previous panels it is no longer true that every valley crossing leads to a fixation – the next mutant might emerge before that. Panel (E) the share of these valley crossings that occur without fixation of the first mutant.**(F)** Simulation results of the process described in (A)-(C). Each simulation ran for$10^5/(k\mu)$ generations. The y-axis measures fixations of adaptive mutations per generation (see Fig 5D). **(G)** Fixations per generation relative to the uniform mutation process (k=1) for k=5 and k=10.

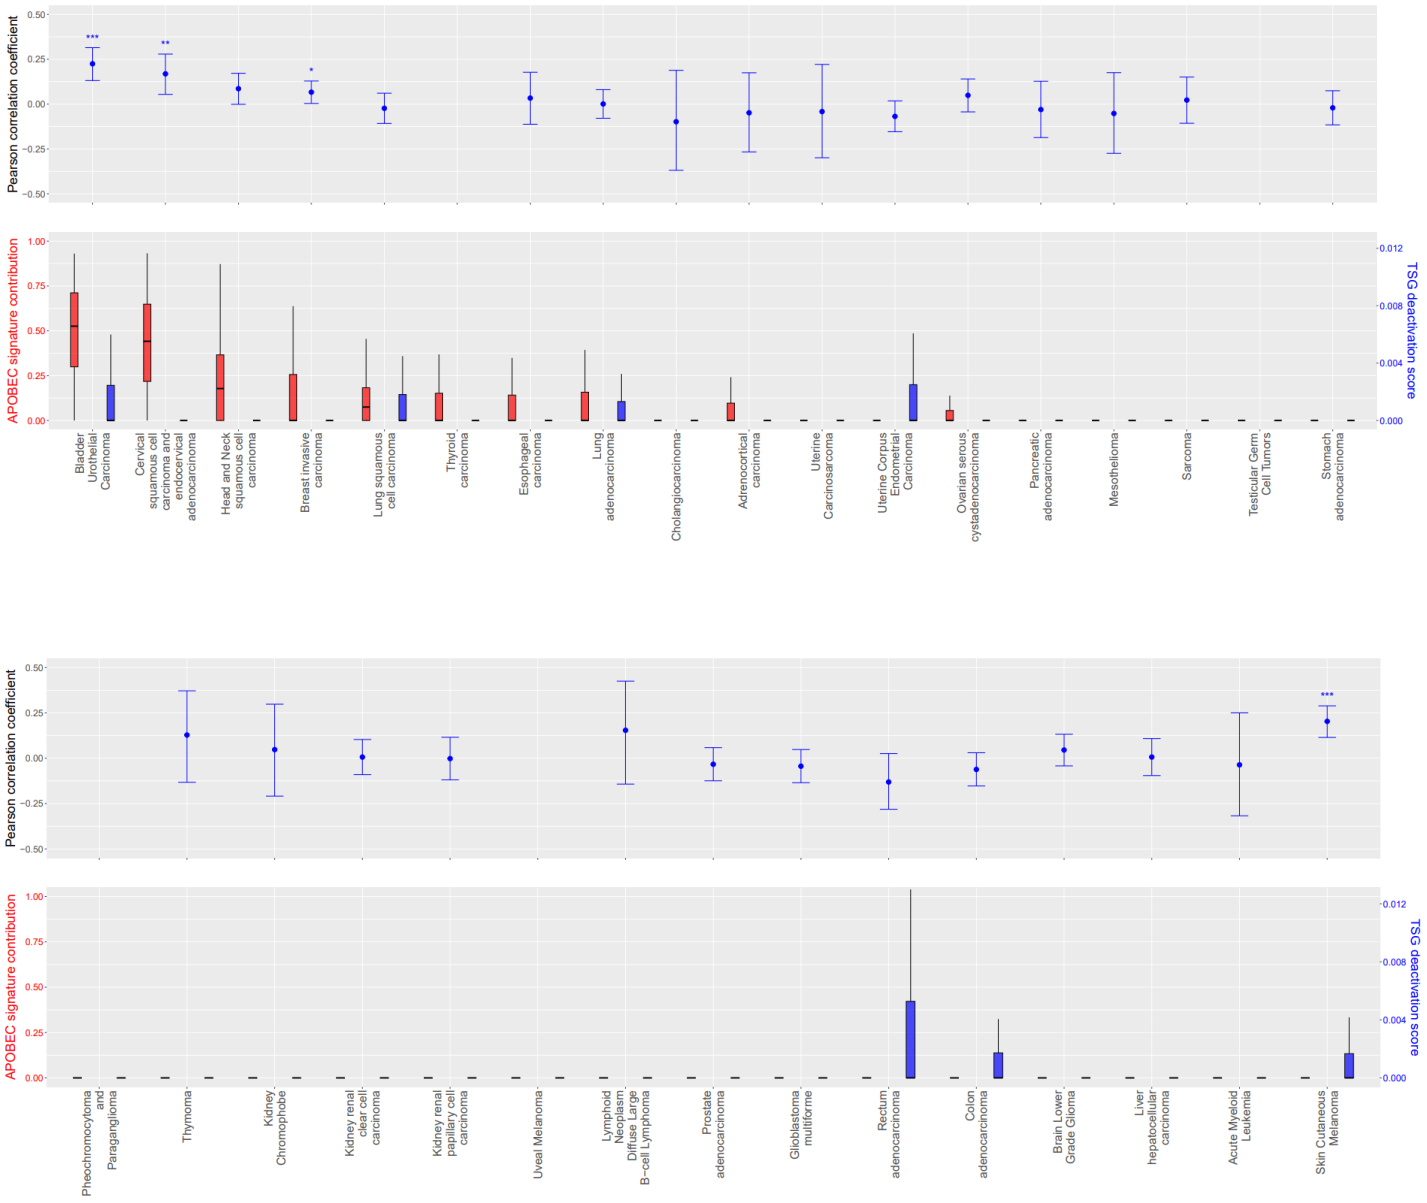

Supplementary Figure 3: **TSG deactivation scores in TCGA.** Distribution and correlation of APOBEC signature contribution and TSG deactivation scores as in Fig 4, sorted by cancer type.

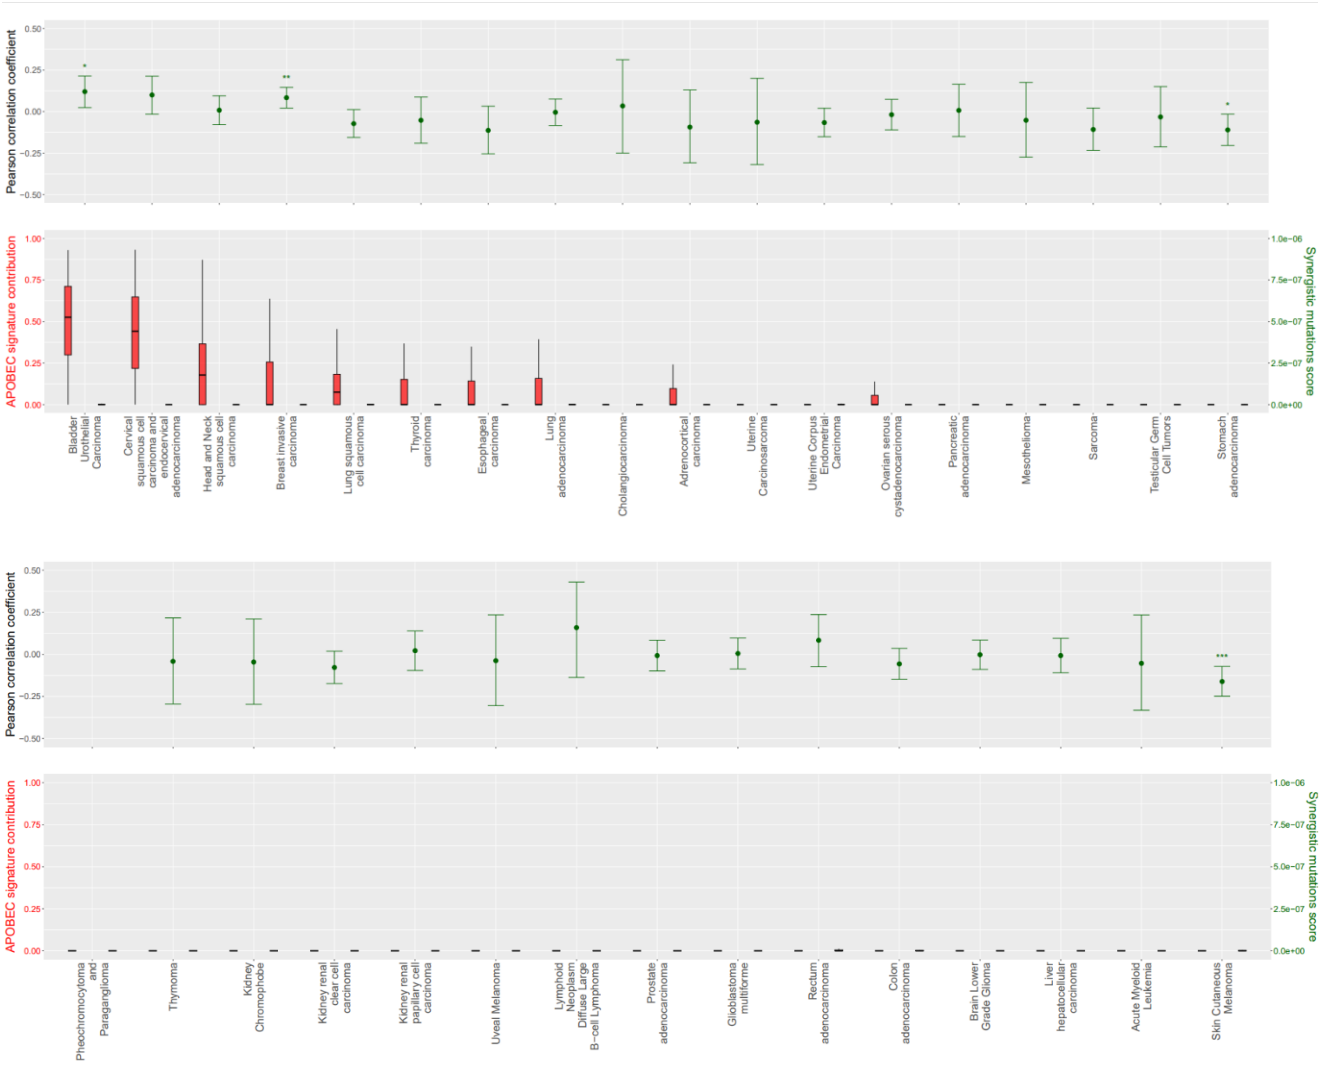

Supplementary Figure 4: **Synergistic mutation scores in TCGA.** Distribution and correlation of APOBEC signature contribution and synergistic mutation scores, sorted by cancer type.

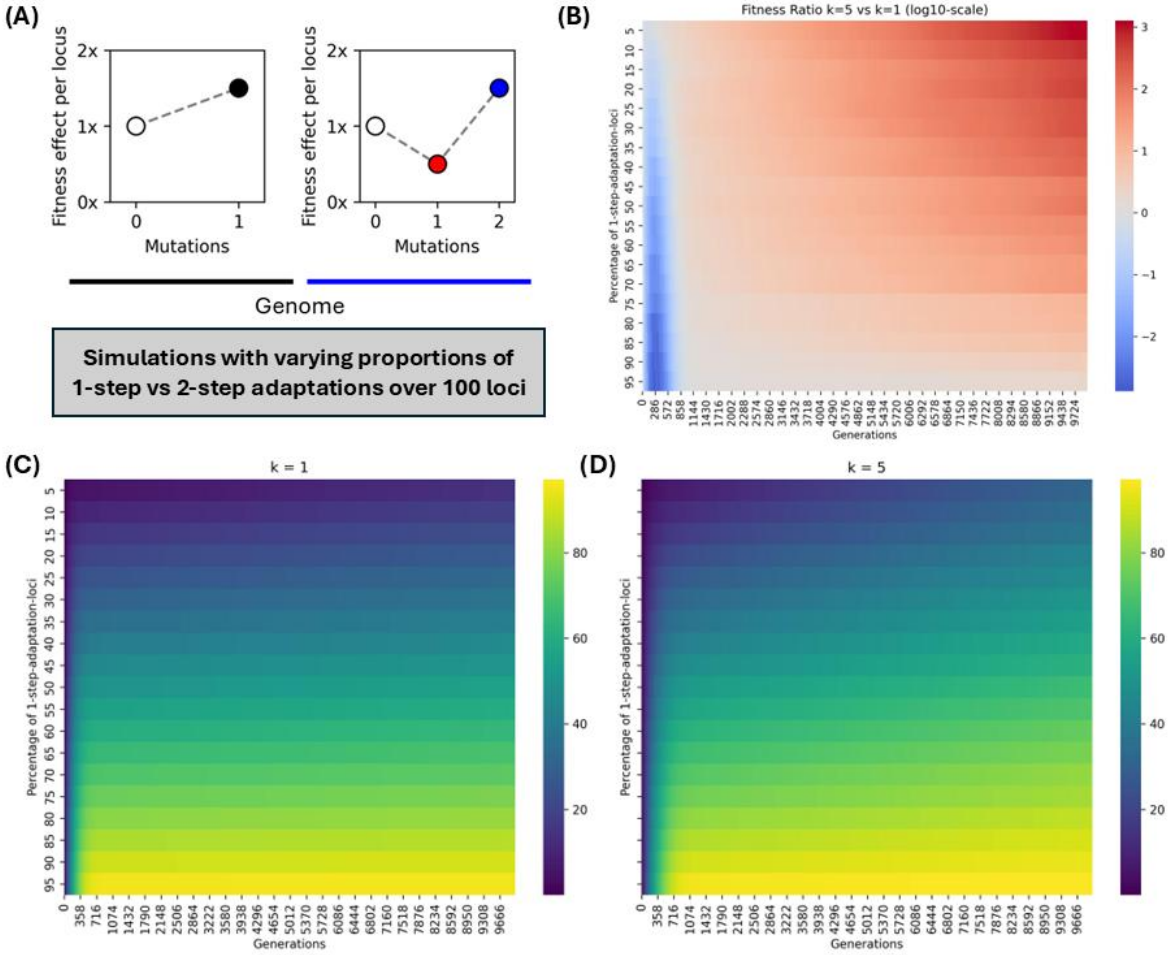

Supplementary Figure 5: **Temporal clustering and adaptation in fitness landscapes with varying ruggedness. (A)** Sketch of the setup. Like in Fig 5, we assume that cells have two types of loci on which one-step or two-step adaptations are possible respectively. Fitness effects per locus are as in Fig 5. Here, we run simulations in which we vary the relative shares of the two types of loci. **(B)** Simulation results over time, averaged for 100 simulation replicates per row and clustering parameter k. Colors indicate the fitness ratio between simulations with k=5 and k=1 on a log10 scale. **(C)-(D)** show the corresponding absolute numbers of loci with adaptative mutations (i.e. loci in the black and blue states in panel (A)).