


Robustness and predictability of evolution in bottlenecked populationsOsmar Freitas ¹, Lindi M. Wahl ², and Paulo R. A. Campos ^{1,*}¹*Evolutionary Dynamics Lab, Physics Department, Federal University of Pernambuco, Recife-PE, 50670-901, Brazil*²*Applied Mathematics, Western University, London, Ontario N6A 5B7, Canada* (Received 30 November 2020; revised 3 March 2021; accepted 2 April 2021; published 19 April 2021)

Deterministic and stochastic evolutionary processes drive adaptation in natural populations. The strength of each component process is determined by the population size: deterministic components prevail in very large populations, while stochastic components are the driving mechanisms in small ones. Many natural populations, however, experience intermittent periods of growth, moving through states in which either stochastic or deterministic processes prevail. This growth is often countered by population bottlenecks, which abound in both natural and laboratory populations. Here we investigate how population bottlenecks shape the process of adaptation. We demonstrate that adaptive trajectories in populations experiencing regular bottlenecks can be naturally scaled in time units of generations; with this scaling the time courses of adaptation, fitness variance, and genetic diversity all become relatively insensitive to the timing of population bottlenecks, provided the bottleneck size exceeds a few thousand individuals. We also include analyses at the genotype level to investigate the impact of population bottlenecks on the predictability and distribution of evolutionary pathways. Irrespective of the timing of population bottlenecks, we find that predictability increases with population size. We also find that predictability of the adaptive pathways increases in increasingly rugged fitness landscapes. Overall, our work reveals that both the adaptation rate and the predictability of evolutionary trajectories are relatively robust to population bottlenecks.

DOI: [10.1103/PhysRevE.103.042415](https://doi.org/10.1103/PhysRevE.103.042415)**I. INTRODUCTION**

As we face the evolutionary emergence of novel pathogens on a global scale, the need for predictive models of microbial evolution has never been more apparent. In microbial adaptation, both deterministic processes such as selection, and stochastic processes such as mutation and drift, drive the adaptive process. Nonetheless, predictability may be an attainable goal, at least partially, because deterministic evolutionary processes prevail in the large population sizes that characterize many microbial communities [1–3].

These large population sizes, however, may not be constant across time. Microbial populations can experience periods of rapid growth, when resources are plentiful, and population bottlenecks are often the necessary consequence of these rapid expansions. Thus, population bottlenecks are repeatedly observed in seasonality, predator-prey dynamics, and pathogen-host dynamics, as well as being a key feature of many experimental protocols for the study of adaptation [4]. While deterministic processes dominate the period of population expansion before the bottleneck, bottlenecks themselves can increase the stochasticity of the evolutionary trajectory. How do these conflicting forces play out in the adaptive trajectory? Do bottlenecks reduce or enhance adaptation, and do they reduce or enhance predictability?

We will address these questions using a computational model of an adapting population, in which we can manipulate

the population size, the length of the growth phase, and the severity of the bottleneck. Central to any such model is the fitness landscape, that is, the mapping between genotype (or phenotype) and the expected growth rate of that type; fitness landscapes have been the focus of both theoretical and experimental investigations stretching back over decades (for review, see Refs. [1,2,5–8]). In the absence of epistasis, the genotype-to-fitness landscape can be described as “smooth” (equivalently: single-peaked, correlated). In this case, each beneficial allele makes a fixed contribution to fitness—fixed in the sense that the benefit does not change depending on the genetic background in which the allele occurs. Maximum fitness is realized by a genotype that carries all possible beneficial alleles, and thus the landscape has a single fitness peak. In contrast, a “rugged,” multi-peaked or uncorrelated landscape may have multiple fitness peaks, and the fitness contribution of particular alleles depends on the genetic background, that is, on alleles at other loci in the genome. For an adapting population, the end-points of adaptation may be less predictable on rugged landscapes, however, the evolutionary trajectories through sequence space may also be more constrained by sign epistasis [3,9].

Multiple lines of experimental evidence, largely from microbial populations, suggest that fitness effects in fitness landscapes are not fully correlated, that is, epistasis is prevalent ([10–12], see Ref. [13] for review). However, measures of predictability and parallel evolution in several experiments strongly suggest that fitness landscapes are not fully uncorrelated either; many mutations retain their beneficial effect across a range of genetic backgrounds [14–16]. Thus,

*paulo.acampos@ufpe.br

microbial adaptation seems to occur on fitness landscapes that lie somewhere between completely smooth and maximally rugged landscapes [16].

In the modeling literature, single-peaked fitness landscapes are far more tractable mathematically, and have been widely used [17–21]. Both the Block model [22,23] and the NK model [24,25] offer fitness landscapes that are tunably rugged, that is, they are partially correlated and parameter values can be used to manipulate the extent of this correlation. Several examples of empirically derived fitness landscapes are also now available, demonstrating varying degrees of epistasis ([2,7,8], for review see Ref. [1]). Recent statistical analyses, however, have demonstrated that any single realization of a fitness landscape is likely to incorporate idiosyncratic features or misleading correlations that disappear on examination of an ensemble of landscapes [26]. Thus, while empirical fitness landscapes are clearly the gold standard for examining any particular evolutionary trajectory, the use of any single empirical landscape to draw general conclusions about adaptation may be misleading.

Here, we investigate a well-established model of a partially correlated fitness landscape, the NK model [24], and simulate adapting asexual populations with changing population sizes. Using the NK model allows us to generate a large number of replicate fitness landscapes, and draw general conclusions over this ensemble. This approach also allows us to independently manipulate the overall degree of epistasis, population size, and mutation rate.

Population bottlenecks are a standard protocol in many microbial evolution experiments, and their effects on evolutionary trajectories have been a topic of renewed interest recently [27,28]. Previous work has suggested that the adaptation rate of large asexual populations will be maximized at an intermediate bottleneck ratio [29]. Because of the stochasticity inherent in bottlenecks, another natural hypothesis is that bottlenecks should reduce the predictability of evolution. In the sections to follow, we examine each of these hypotheses, demonstrating that both adaptation and predictability are surprisingly robust to population bottlenecks.

II. METHODS

The simulation tracks a population of asexual individuals whose fitness values are determined by their genotype \vec{S} . Each individual is represented as a genome of L loci, $\vec{S} = (s_1, s_2, \dots, s_L)$, where $s_\alpha = \{0, 1\}$. The L -dimensional genotype space consists of all $M = 2^L$ possible sequences.

A fitness landscape is a mapping from a set of genotypes, equipped with some notion of adjacency, nearness, distance or accessibility, into fitness values [30]. In the NK fitness landscape model considered here, there exist interdependencies (epistatic interactions) between the loci comprising the genome: If an allele at a given locus changes, then it affects both the fitness contribution of the focal locus and the contributions of any loci that are epistatically coupled to the focal locus. In particular, the degree of epistasis is tuned by the parameter K . Thus, the contribution of locus j to the overall fitness, ω_j , is a function given by $\omega_j = g[s_j, \Pi(j)]$, that depends on state of locus j , s_j , and on the state of a set of K neighbors, $\Pi(j)$. In our formulation, the K neighbors are

randomly chosen among the remaining $L - 1$ loci. Therefore, for each locus j , a lookup table consisting of 2^{K+1} entries is built, whose values, ω_j , are drawn from an uniform distribution (0,1]. Finally, the overall fitness of the genotype is taken as the arithmetic mean of the contributions of all loci, i.e.,

$$f = \frac{1}{L} \sum_j \omega_j. \quad (1)$$

In the limit $K = 0$, the NK-fitness landscape is said to be additive, as the fitness contribution of each element of \vec{S} is independent of all others. This means that the allele at each locus can be optimized independently. The resulting fitness landscape is smooth and single-peaked. At the other extreme, $K = L - 1$, the change of a single locus adjusts the contribution of all elements to the overall fitness. In this case, the fitness landscape is maximally uncorrelated, and typically contains a large number of fitness peaks.

A. Simulation protocol

In our simulation protocol, initially the population is isogenic and all sequences are set to the antipode of S_{\max} , the global optimum of the fitness landscape. Accordingly, the Hamming distance from the initial genome to the global optimum of the fitness landscape is always equal to L . The population size is variable: we simulate a population of initial size N_0 that experiences a growth phase for τ discrete generations (doublings), reaching size N_f , where $N_f = 2^\tau N_0$. At the end of the growth phase, the population of size N_f is subject to a bottleneck protocol; in particular, N_0 individuals are chosen uniformly at random to form the next founding population.

During the growth phase, the population doubles its size at every time step. Following the standard Wright-Fisher model with nonoverlapping generations, the individuals replicate proportionally to their fitness values. In particular: for each of the N_g newborn individuals in generation g , a uniform random number is generated and compared to the normalized sum of fitnesses in generation $g - 1$ to determine the parent of that newborn. Newborn individuals inherit the parental genotype, but mutations occur randomly with a fixed probability U per generation. During evolution, we keep track of the evolutionary pathway, and unless stated otherwise, the process is only stopped when the population reaches the global optimum, S_{\max} , of the fitness landscape.

B. Characterization of the evolutionary pathways

We note that in many previous theoretical and especially in computational studies of adaptive walks, the population is assumed to remain isogenic during the adaptive process [31,32]. Departures from this assumption are negligible in the strong-selection weak-mutation regime, i.e. when $U \ll (N \log N)^{-1}$. For the mutation rates and population sizes studied here, however, the population is more accurately described as a cloud of mutants around a given sequence. If mutation rates are not too large the cloud of mutants will cluster around the sequence of maximum fitness value.

As mentioned, we track the entire array of sequences visited ($\vec{S}_0, \vec{S}_1, \dots, \vec{S}_n$) in a given evolutionary pathway. Particularly, we are concerned with the most fit genotype along

the evolutionary trajectory. When the most fit sequence in the population changes, it is added to the list. We note that alternative methods to build the evolutionary pathway can be used, such as following lines of descent, as implemented by Szendro *et al.* [3]. While our approach characterizes aspects of the evolutionary trajectory of the *population*, lines of descent trace the evolutionary pathway of a single genotype. Whether the qualitative results we present here are robust to such differences in pathway definition remains an interesting question for future work.

After many independent trials, an ensemble of trajectories is then considered to produce the statistical analysis. The starting and ending points of these evolutionary trajectories are always the same (unless stated otherwise): $S_n = S_{\max}$, and S_0 is the antipode of S_{\max} . However, the number of steps, n , in the evolutionary pathway is variable and dictated by the dynamics themselves. Before analysis, the trajectories are purged of loops (if a sequence appears more than once in an evolutionary pathway, which characterizes a loop structure, the pathway is redefined with the loop removed). This is to avoid situations in which a fitter genotype is generated and then lost (does not increase in frequency), not truly characterizing a displacement of the population to a new fitter sequence.

Among the quantities of interest are the predictability and the mean path divergence [1,33–35]. Predictability with respect to the paths is a measure of the repeatability of the evolutionary pathway, and its inverse can be seen as the number of effective pathways exploited by the dynamics. Predictability is defined as

$$P_2 = \sum_{q_\alpha} \phi^2(q_\alpha), \quad (2)$$

where $\phi(q_\alpha)$ is the observed frequency of trajectory q_α in the ensemble, and the sum is taken over all trajectories in that ensemble. Alternatively, the mean path divergence is used to assess the similarity between evolutionary pathways, and relies on the estimate of the pairwise divergence $d(q_\alpha, q_\beta)$ between paths q_α and q_β . For each sequence S_α comprising pathway q_α , we estimate the Hamming distance from that sequence to every sequence in q_β , $h(S_\alpha, q_\beta)$. The lowest distance is then stored, and the process is repeated until all sequences in q_α are rated. The process is then repeated in reverse, from q_β to q_α . The divergence $d(q_\alpha, q_\beta)$ is taken as the mean value of those shortest Hamming distances [33,36]:

$$d(q_\alpha, q_\beta) = \frac{1}{n_{q_\alpha} + n_{q_\beta}} \left(\sum_{S_\alpha \in q_\alpha} h(S_\alpha, q_\beta) + \sum_{S_\beta \in q_\beta} h(S_\beta, q_\alpha) \right), \quad (3)$$

and the mean path divergence is then estimated as

$$\bar{d} = \sum_{q_\alpha} \phi(q_\alpha) \sum_{q_\beta} \phi(q_\beta) d(q_\alpha, q_\beta), \quad (4)$$

where the sum is carried out over the ensemble of evolutionary pathways. A previous study has shown that the two quantities, P_2 and \bar{d} , are negatively correlated, although such a claim may not be generalizable, as such correlations may strongly depend on the topological properties of the underlying fitness landscape [26].

III. ANALYTICAL RESULTS

1. Expected fitness values

The expected fitness value of the global optimum (GO) of the NK fitness landscape can be predicted analytically for any sequence length L , for the two asymptotic cases $K = 0$ (a smooth landscape) and $K = L - 1$ (a maximally rugged landscape). These expectations are not only of theoretical interest, but, importantly, allow us to independently validate the computational implementation of the fitness landscape described above.

We use an additive fitness function in which the contribution of locus i , ω_i , is drawn from a uniform distribution, $U(0, 1]$. When $K = 0$, the contribution of locus i to the GO fitness, f_{\max} , is the maximum of the two possible fitness values at locus i . It is straightforward to demonstrate that the expected value of the maximum of two draws from $U(0, 1]$ is $2/3$. Thus, the GO fitness f_{\max} is given by $1/L$ times the sum of L independent, identically distributed random variables, each of which has expected value $2/3$. The expected value of the GO, $E[f_{\max}]$, is thus $2/3$.

When $K = L - 1$, each of 2^L possible sequences is independently assigned a fitness value. Each of these fitness values is computed as $1/L$ times the sum of L draws from $U(0, 1]$. The cumulative density function (cdf) for the sum of L draws from $U(0, 1]$ is given by the Irwin-Hall distribution [37,38]:

$$H_L(x) = \frac{1}{L!} \sum_{i=0}^{\lfloor x \rfloor} (-1)^i \binom{L}{i} (x - i)^L.$$

Let F denote the random variable for the fitness of a sequence. To determine the expected value of the GO, we first compute the cdf of the GO, $M(x)$. This gives the probability that the maximum of 2^L independently drawn values of F is less than or equal to x :

$$M(x) = \text{Prob}(\text{maximum of } 2^L \text{ values of } F \leq x) \quad (5)$$

$$= \text{Prob}(\text{maximum of } 2^L \text{ draws from } H_L \leq Lx) \quad (6)$$

$$= \text{Prob}(\text{each of } 2^L \text{ draws from } H_L \leq Lx) \quad (7)$$

$$= \prod_{k=1}^{2^L} H_L(Lx) \quad (8)$$

$$= \prod_{k=1}^{2^L} \frac{1}{L!} \sum_{i=0}^{\lfloor Lx \rfloor} (-1)^i \binom{L}{i} (Lx - i)^L. \quad (9)$$

The expected value of the GO is then given by integrating the product of x with the probability density function associated with $M(x)$:

$$E[f_{\max}] = \int_0^\infty x \frac{d}{dx} M(x) dx. \quad (10)$$

In Fig. 1, we illustrate these analytical predictions of $E[f_{\max}]$, for $K = 0$ and $K = L - 1$, along with results for the observed value of f_{\max} , averaged over 100 000 simulated landscapes. As expected, the analytical predictions agree with simulation results when $K = 0$ or $K = L - 1$, and in all other cases give upper and lower bounds on $E[f_{\max}]$. We further observe that for the relatively short sequences we investigate,

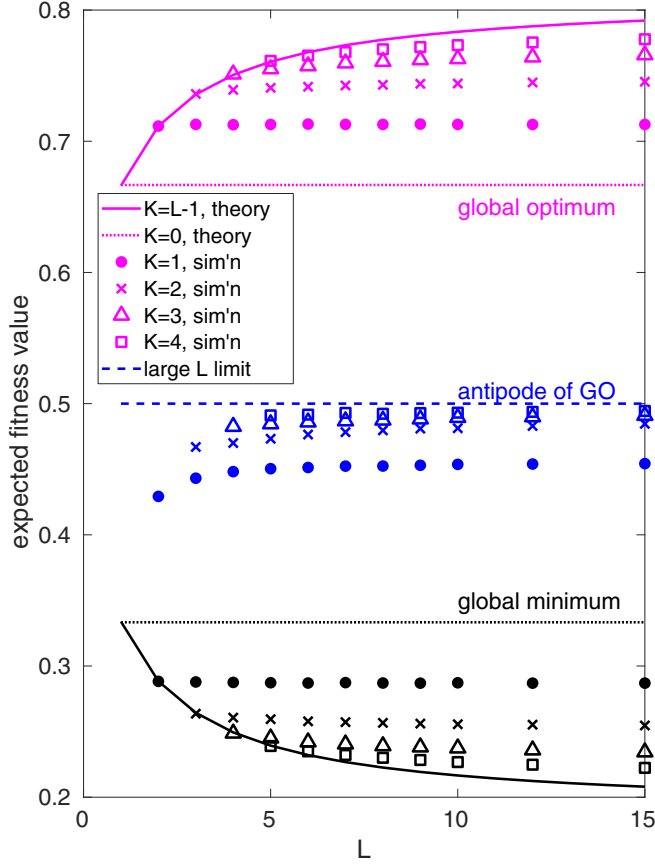


FIG. 1. Global optimum, antipode of global optimum and global minimum fitness values, versus sequence length, L . The analytical prediction [Eq. (10), evaluated numerically] for the global optimum fitness, $E[f_{\max}]$, is shown for $K = L - 1$ (magenta solid line), along with the analytical prediction of $E[f_{\max}] = 2/3$ for $K = 0$ (magenta dotted line). Analogous analytical predictions for the global minimum fitness are shown for comparison (black lines). Simulation results are shown for comparison for $K = 1, 2, 3$ and 4 (magenta and black symbols as indicated). Simulation results for the expected fitness of the antipode of the global optimum are shown in blue, along with the asymptotic expectation (0.5 for large L and $K = L - 1$, dashed line). Simulation results show the mean across 100,000 randomly generated fitness landscapes in each case. Error bars for simulation results are similar to symbol heights and omitted for clarity.

the GO fitness depends strongly on the ruggedness of the landscape (K) and weakly on sequence length (L).

For comparison, we also plot the global minimum fitness (black) and the fitness of the antipode of the GO (blue) in Fig. 1. By analogous arguments, it is straightforward to demonstrate that the expected value of the global minimum is simply $1 - E[f_{\max}]$. When $K = 0$, the global minimum corresponds to the antipode of the GO, and thus the expected value of the global minimum and the antipode of the GO is $1/3$ (black dotted line). When $K = L - 1$, the fitness of the GO antipode is given by a randomly chosen fitness from the landscape, conditioned by the fact that the chosen fitness is not the GO. Thus, the GO antipode fitness approaches the mean landscape fitness, $1/2$, as L increases. For smaller values of L , the ex-

pected fitness of the GO antipode is less than 0.5 , because this conditioning has a more pronounced effect when L is small.

2. Wright-Fisher comparison

To assess the impact of population bottlenecks on predictability, we would like to compare results with constant-sized populations. In this subsection we demonstrate that our simulation protocol, when $\tau = 1$, is formally equivalent to a Wright-Fisher population maintained at a constant size.

Consider a discrete time Wright-Fisher model with constant population size N , in which the i th individual in the population has absolute fitness W_i . This parent individual i has a Poisson-distributed number of offspring with expected value W_i , and these offspring are sampled to form the next generation. It is standard to assume that the number of offspring is large, such that offspring can be sampled with replacement, that is, each offspring is selected independently with a fixed probability. To maintain a constant population size, the sampling probability must be $1/\bar{W}$, where \bar{W} is the mean population fitness. The probability generating function (pgf) for the descendants of individual i in the next generation, $F_i(x)$, is then given by the composition of the Poisson offspring pgf, $\exp[W_i(x - 1)]$ and the binomial sampling pgf, $(1 - p) + px$ (where $p = 1/\bar{W}$ is the sampling probability):

$$F_i(x) = \exp \left[W_i \left(\left(1 - \frac{1}{\bar{W}} \right) + \frac{1}{\bar{W}} x - 1 \right) \right] \quad (11)$$

$$= \exp \left[\frac{W_i}{\bar{W}} (x - 1) \right]. \quad (12)$$

Thus, the net effect of this process—a large, Poisson-distributed number of offspring, followed by independent sampling with a constant probability—is the same as a Poisson distribution of descendants with mean $\frac{W_i}{\bar{W}}$.

In the simulations to follow, the population size doubles for τ generations, with each individual contributing offspring to the next generation in proportion to their relative fitness. The population is then sampled with sampling probability $2^{-\tau}$. When $\tau = 1$, the contribution of the i th member of the initial population (of size N_0) to the next population of size N_0 (after one cycle of growth and one bottleneck) is therefore given by

$$f_i(x) = \exp \left\{ \frac{2W_i}{\bar{W}} \left[\left(1 - \frac{1}{2} \right) + \frac{1}{2} x - 1 \right] \right\} \quad (13)$$

$$= \exp \left[\frac{W_i}{\bar{W}} (x - 1) \right] = F_i(x). \quad (14)$$

Thus, in the results to follow, cases illustrated for $\tau = 1$ (often the extreme or asymptotic cases) are equivalent to a standard discrete time Wright-Fisher model at fixed population size N_0 . In other words, the case $\tau = 1$ reveals the behavior of a Wright-Fisher population in the absence of population bottlenecks.

IV. SIMULATION RESULTS

We will first investigate the role of population bottlenecks in adaptation in Sec. IV A. In particular, we are interested in the conditions that optimize the rate of adaptation. In Sec. IV B, we turn our attention to the characterization of




FIG. 2. Fitness trajectories, that is, mean population fitness versus time for different bottleneck ratios (upper panels) and fitness versus bottleneck size at different times (lower panels). Time is measured in units of bottlenecks (left panels) and doublings (right panels). The parameter values are mutation rate $U = 1 \times 10^{-4}$, sequence size $L = 8$, epistasis parameter $K = 2$, and N_f is set at $N_f = 2^{15} = 32\,768$. The bottleneck sizes are indicated in the legends. In the bottom panels the curves correspond to fixed numbers of bottlenecks (or doublings) as indicated in the legends.

evolutionary pathways and how their statistical properties are affected by the bottlenecks.

A. Optimal adaptation rate

When investigating which bottleneck protocol optimizes the adaptation rate, we must consider both the underlying demographics that will be deemed “comparable,” and the “rate” we seek to optimize [4]. Here, we consider the rate at which the mean population fitness increases, that is, the rate of at which *de novo* beneficial mutations occur, survive, and increase mean fitness. It is also necessary to define the time unit in this rate: as will become apparent in the figures to follow, the optimal adaptation rate differs if we consider the adaptation rate per “birth” event in the population, per generation (doubling), or per bottleneck. Since τ doublings occur per bottleneck, it is clear that $t_{\text{doubling}} = \tau t_{\text{bottleneck}}$. Similarly, the population begins at size N_0 and doubles τ times, such that the total number of births between bottlenecks is $2N_0(2^\tau - 1)$, and thus the total number of births after several bottlenecks is $t_{\text{birth}} = 2N_0(2^\tau - 1)t_{\text{bottleneck}}$. The factor $2N_0(2^\tau - 1)$ corresponds to the sum of a finite geometric series of first term $2N_0$ and common ratio 2. Once more, we highlight that nonoverlapping generations are considered. Each of these rates may have practical relevance: If the population is resource-limited, then a limited total number of new births may be possible, and thus the adaptation rate per birth might be the critical factor in an evolutionary rescue scenario; if we are concerned with environmental change that occurs at a pace set by calendar

time, then the adaptation rate per generation time may be relevant; if we consider an experimental population for which the bottleneck process itself is labor-intensive, then adaptation per bottleneck will be the appropriate rate to compare across cases. Demographically, bottleneck ratios will be compared across populations that have either the same initial population size, or the same final population size.

B. Fitness trajectories

Figure 2 shows average fitness trajectories for different bottleneck ratios. The data correspond to an average of 1000 replicates each for over 50 distinct randomly drawn fitness landscapes. Here, the population size after the growth phase is fixed, $N_f = 2^{15} = 32\,768$. The fitness trajectories are shown in time units of bottlenecks (left panels) and in time units of doublings (right panels).

Figure S11 [39] displays how the mean selective effect of the beneficial and the proportion of beneficial mutations evolve over time. The rapid drop of both quantities is a clear signature of the pattern of diminishing returns [40], and explains why the rate of increase of fitness slows down with time.

We do not expect the bottleneck ratio to produce any considerable discrepancy in the long-term fitness attained by the populations; however bottlenecks do play a clear role at earlier stages of adaptation. These findings are better summarized in the lower panels of Fig. 2. Here, curves correspond to distinct times at which fitness was reported. The existence of an

optimal bottleneck ratio seems to be inherent to the dynamics when time is given in units of bottlenecks, being observed at different times and across landscape structures; Fig. S2 [39] shows analogous results for varying degrees of epistasis in the landscape. Moreover, these results predict that sampling about 10–20% of the population will maximize the speed of adaptation per bottleneck; this agrees with previous studies addressing the impact of bottleneck ratios on the fixation probability of single mutations, predicting that the optimum ratio occurs around $1/e^2$ [29,41].

In contrast with this result, when time is measured in units of doublings we observe that fitness becomes monotonic in bottleneck size (right panels in Figs. 2 and S2 [39]). In other words, larger effective population sizes lead to a higher rate of adaptation, although this increase begins to saturate for very large population sizes. Finally, when time is measured in units of births (see Fig. S1 [39]), one recovers the scenario shown on the left in Figs. 2 and S2 [39], and once again the highest adaptation rates are found at intermediate bottleneck sizes. These qualitative results hold regardless of the mutation rate, sequence size and the topography of the fitness landscape, as demonstrated in Figs. S3–S7 [39].

In the results above we changed the bottleneck ratio by varying N_0 , while holding N_f fixed. Analogous results, with N_0 constant while N_f varies, are shown in Fig. S8 [39]. Whether time is measured in units of bottlenecks and doublings, the mean population fitness displays a monotonic increase with N_f in this case. In contrast, the mean population fitness is reduced as N_f grows when time is expressed in units of births (see Fig. S9 [39]). Overall, we conclude that when time is measured in population doublings, the adaptation rate increases monotonically in larger populations, irrespective of population bottlenecks.

C. Understanding the optimal adaptation rate at intermediate bottleneck sizes: Genetic diversity and fitness variance

As we know, the fundamental theorem of natural selection states that the rate of increase in fitness is proportional to the genetic variance in fitness [42]. Therefore, it is crucial to understand how bottleneck sizes affect the lineage genetic diversity and fitness variation. Several measures of genetic variation can be used to assess the amount of genetic variation among individuals within as well as between populations. The simplest and most commonly used are the genotypic richness (number of different genotypes in the population) n , heterozygosity H , and various related measures of entropy, such as the Shannon entropy [43]. These measures can be unified through the use of Hill numbers, which capture essential properties of genetic diversity in a population [44].

The Hill diversity number of order a is defined as

$$D_a = \left(\sum_i p_i^a \right)^{\frac{1}{1-a}}, \tag{15}$$

where p_i is the frequency of genotype i in the population. These diversity numbers can be understood as the weighted sum of each p_i to the power $a - 1$, where the weights are themselves the p_i . We then take the $(a - 1)$ th root of that sum, thus D_a is simply a *weighted $(a - 1)$ -Norm of the vector of*




FIG. 3. Mean population fitness, fitness variance and change in fitness Δf as a function of time. Time is expressed in units of doublings. The parameter values are $N_f = 2^{15} = 32768$, mutation rate $U = 1 \times 10^{-4}$, sequence size $L = 8$ and epistasis $K = 2$. The bottleneck sizes are $N_0 = 32$ (blue dashed-lines), $N_0 = 4096$ (orange dashed-lines) and $N_0 = 16384$ (green dashed-lines). Δf is simply the mean population fitness at time $t + 1$ minus the mean population fitness at time t , for t in the units indicated.

genotype frequencies. The Hill diversity numbers can be easily related to the three commonly used measures of diversity:

- (1) $D_0 =$ number of different genotypes $= n$;
- (2) $D_1 = \exp(S)$ [44], where S is the Shannon entropy;
- (3) $D_2 = \frac{1}{1-H}$, where $H = 1 - \sum_i p_i^2$ is the heterozygosity.

Thus, while diversity is commonly measured by several measures with differing units (species richness, Shannon or related entropies, heterozygosity), the Hill numbers provide a unified understanding and consistent units across which these measures can be compared. In particular, we note that D_a of any order measures diversity in units of genotypes, and can be interpreted as an effective number of genotypes or number of lineages in the population. While D_0 simply counts the number of distinct genotypes in the population, as the order a increases the contribution of rare types to the corresponding diversity D_a is reduced.

Figure 3 shows the time evolution of the mean population fitness, the fitness variance as well the fitness increase at each time step. Three different bottleneck sizes are compared: $N_0 = 2^5 = 32$, $N_0 = 2^{12} = 4096$ and $N_0 = 2^{14} = 16384$. In all cases the population size at the end of the growth phase is $N_f = 2^{15}$, and mean fitness and fitness variance are compared (in this figure) just after each population bottleneck, for a population of size N_0 .

First and as expected, in the earlier stages of adaptation, while adaptation occurs at a faster pace, the fitness variance is considerably enhanced. As expected, in each population the

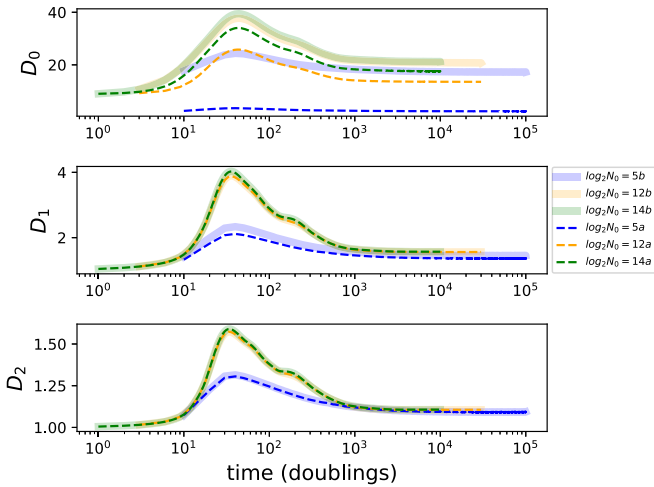


FIG. 4. Hill diversity numbers D_0 , D_1 , and D_2 versus time. Time is expressed in units of bottlenecks (doublings) on the left (right) panels. The measures of diversity are presented at the end of the growth phase (solid lines) and just after the bottleneck protocol (dashed lines). The parameter values are $N_f = 2^{15} = 32\,768$, mutation rate $U = 1 \times 10^{-4}$, sequence size $L = 8$ and epistasis parameter $K = 2$. The bottleneck sizes are $N_0 = 32$ (blue lines), $N_0 = 4096$ (orange lines) and $N_0 = 16\,384$ (green lines) as indicated in the legends. The symbol a in the legend means just after bottlenecks, whereas b means just before bottlenecks.

adaptation rate (slope in top panel) is greatest when the fitness variance is maximized; we can confirm this by comparing the fitness variance with the change in fitness per timestep (bottom panel). We also note in the center panels that the peak fitness variance is reduced for the smallest population $N_0 = 32$, but is not sensitive to the bottleneck for less severe bottleneck ratios. As the fitness variance is measured just after a bottleneck, this loss in fitness variance is attributed to the loss of lineages through the bottleneck; severe bottlenecks retard the adaptation rate through the loss of fitness variance. We will return to this effect in Fig. 4 to follow.

As can be corroborated from Fig. 3, measuring time in doublings reconciles the time courses of mean fitness and fitness variance, and isolates the effects of the bottlenecks. We see clearly here that only the most severe bottleneck ratio in this example (2^{-10}) reduces fitness variance and thus retards the adaptation rate. Importantly, we note that when $N_0 = 2^{14}$ and $N_f = 2^{15}$ (green lines), the simulated population dynamics are mathematically equivalent to a discrete time Wright-Fisher process, in a population of fixed size 2^{14} (see Analytical Results, above). Thus, the green lines plot the time course of adaptation *that would be obtained in the absence of population bottlenecks*. We see that when time is measured in units of population doublings, the rate of adaptation is almost insensitive to bottlenecks, once the bottleneck size exceeds several thousand individuals (also illustrated in top right panel of Fig. 2).

In Fig. 4, the Hill numbers D_0 , D_1 , and D_2 are plotted over time for the same populations ($N_0 = 2^5$, $N_0 = 2^{12}$, and $N_0 = 2^{14}$, with $N_f = 2^{15}$). Values computed just before the population bottleneck (solid lines), and immediately after the bottleneck protocol (dashed lines) are shown. Recall that the

zeroth order Hill number, corresponding to the number of distinct lineages, is very sensitive to low frequency components, whereas both D_1 and D_2 put less emphasis on rare genotypes. Taking time in units of doublings allows us to isolate the effect of the bottlenecks on genetic diversity. We see that the first and second order Hill numbers, D_1 and D_2 , change only negligibly over the course of a single bottleneck (solid versus dashed lines). In contrast, the zeroth order Hill number D_0 is greatly reduced, especially when the dilution ratio N_0/N_f is small (blue lines correspond to a ratio of $2^{-10} \approx 1 \times 10^{-3}$). This loss of rare lineages, revealed by D_0 , has long term consequences for the population; although a single bottleneck has little effect on D_1 or D_2 , we see that both of these diversity measures are greatly reduced in the $N_0 = 32$ population. Thus, the loss of rare lineages through the bottleneck feeds forward, resulting in overall reduced diversity at later times, even as measured by higher order metrics that are less sensitive to rare lineages.

D. The effect of bottlenecks on predictability

At this point, it is our aim to carry out a more detailed analysis of the evolutionary pathways at the genotypic level. As mentioned previously, two quantities, predictability and mean path divergence, may elucidate the dynamics. Computationally, simulations which quantify predictability and mean path divergence are costly, as they require that the starting and ending points of the trajectories are the same. We first investigate adaptive trajectories assuming that the population starts from the antipode of S_{\max} , the global optimum of the fitness landscape. In this case, the Hamming distance from the starting point to the global optimum takes its maximum value, being equal to the sequence size L , $d_{GO} = L$. While the route is smooth for an additive landscape ($K = 0$), it becomes increasingly tortuous as the landscape becomes more rugged. For low mutation rates the population is easily trapped by local optima of the landscape, and so the time needed to reach the global optimum rises substantially. Our estimates for the predictability and mean path divergence are taken over a sample of 50 randomly generated fitness landscapes, and for each landscape 1000 replicates are simulated.

Figure 5 displays the predictability, P_2 , and mean path divergence, \bar{d} , for two distinct scenarios: in the first case (upper panels), the bottleneck size N_0 is kept constant while the population size at the end of the growth phase, N_f , varies. Here, the predictability grows steadily with N_f , whereas the mean path divergence drops off. As previously argued, by augmenting N_f while N_0 is fixed, there is an increased supply of beneficial mutations between bottlenecks, as well as increased interference among those mutations, which in turn increases the selection coefficients of those mutations that fix [17,45,46]. As an overall result, the adaptation dynamics approach a more deterministic regime. These results also corroborate the existence of a negative correlation between predictability and mean path divergence [26]. In the second case (N_f fixed, N_0 varied), we observe a monotonic increase of predictability with N_0 , along with the expected decline in mean path divergence. We note further that predictability rises steeply with N_0 when N_0 is small, but saturates at larger N_0 values. In both panels, we therefore find that irrespective of

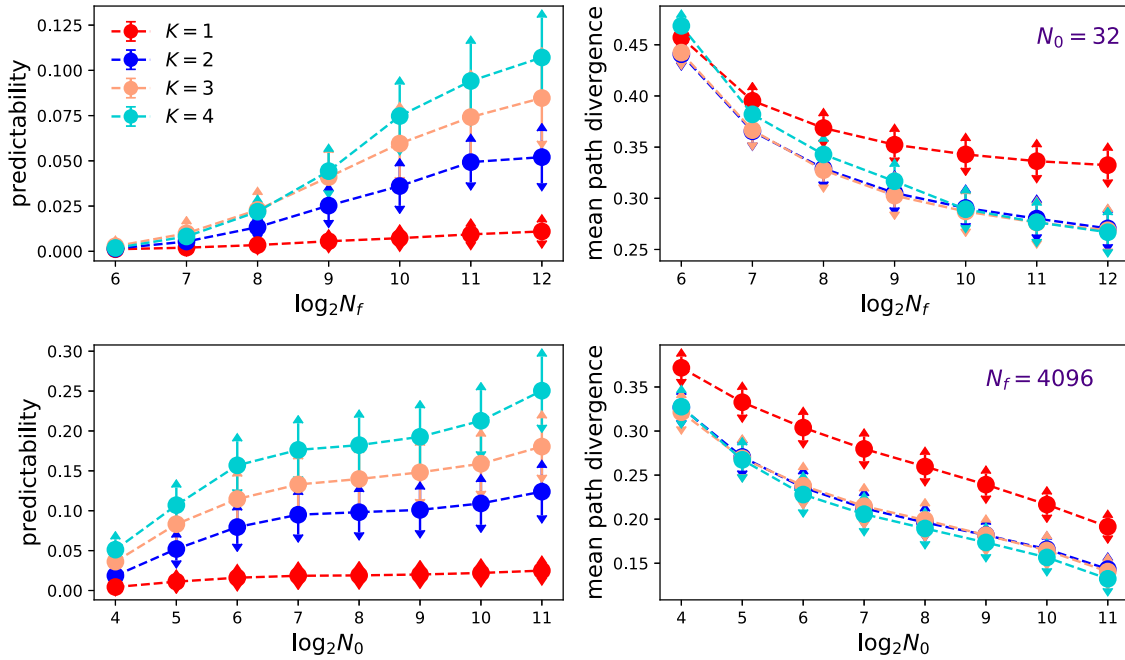


FIG. 5. Predictability and mean path divergence. In the upper panels both quantities are shown as a function of the population size at the end of the growth phase N_f . In these panels the population size after the bottleneck is set at $N_0 = 32$. In the lower panels both quantities are shown as a function of the population size after the bottleneck N_0 . In these panels the population size N_f is set at $N_f = 4096$, whereas the mutation rate is set at $U = 5 \times 10^{-2}$ and sequence size at $L = 8$. The epistasis parameter K is displayed in the legends.

the timing of population bottlenecks, predictability increases and saturates with an increasing mutational supply.

Figure 5 also demonstrates that predictability increases with the ruggedness of the fitness landscape (increasing K). Since these simulation results involve trajectories that are constrained to start at the antipode and end at the GO, pathways between this starting and ending point become ever more constrained as epistasis increases. We also note that predictability is more responsive to the ruggedness of the fitness landscape than the mean path divergence. While the predictability rises monotonically with K , the mean path divergence seems to be bounded already at intermediate ruggedness, and the curves for $K = 2, 3$, and 4 nearly collapse. Overall, when considering adaptive trajectories that are constrained to start at the antipode and end at the GO, any increase in epistasis (landscape ruggedness) very quickly reduces the divergence of pathways, and more gradually increases predictability.

To understand the role of sequence size L on the distribution of evolutionary pathways, a more general simulation protocol is necessary. The main limitation of the previous protocol is that the study cannot be generalized to larger sequence sizes or to less restrictive ranges of correlation in the fitness landscape, because the population tends to become trapped at local maxima, making the global optimum essentially inaccessible. To address this issue, instead of initiating the dynamics from the antipode of S_{\max} , the evolutionary process is now initiated from genotypes that are placed at a given Hamming distance, $d_{GO} < L$, from the global optimum [3]. Results are shown in Fig. S12 [39] for $d_{GO} = 5$, where the sequence size as well as the epistasis parameter K are varied such that the degree of correlation among fitness effects of mutations, $\gamma = 1 - (K + 1)/L$, remains unchanged [47,48]. In the plot, this correlation is set at $\gamma = 0.75$ (upper panels) and $\gamma = 0.5$

(lower panels). It is clear from these results that predictability decreases while the mean path divergence increases with the sequence size L . Note that here both quantities are measured with respect to the paths, i.e., the starting and ending points for each ensemble of trajectories are the same. The data refer to an average over 10 distinct starting points, all chosen at the same distance d_{GO} from the global optimum. Although we were unable to identify an explicit dependence of the predictability on the sequence size L , we observe empirically, at least for high level of correlation γ , that the mean path divergence scales roughly with L , $\bar{d} \propto L$ (see Fig. S10 [39]). As the landscape becomes more rugged, the accuracy of this scaling approximation is reduced.

In Fig. 6, the predictability with respect to the ending points, denoted by $P_{2,\text{ending}}$, is shown as a function of N_0 . The quantity $P_{2,\text{ending}}$ gives the probability that two randomly chosen paths terminate at the same genotype. Here, the evolutionary trajectories begin at a fixed distance from the global optimum, and are simulated up to a fixed time, $t_{\text{bottleneck}} = 2000$. In the left panel, the dependence on the sequence size L is shown as the correlation γ is fixed. We observe that in spite of the Hamming distance from the starting points to the global optimum being the same, $d_{GO} = 5$, the predictability decreases with L . Note that for $L = 8$ and $K = 3$ the predictability is very close to one, meaning that the global optimum is easily accessible. As L increases, the accessibility of the global optimum decreases and so does the predictability. In the right panel we investigate the role of the mutation rate. Higher mutation rates lead to higher predictability with respect to the ending points as they also allow the population to more easily escape from local maxima of the landscape and ultimately reach the global optimum. When the mutation rate U and bottleneck size are sufficiently




FIG. 6. Predictability with respect to the ending points, $P_{2,ending}$. In all plots, the population size at the end of the growth phase is $N_f = 32\,768$. In the left panel, the mutation rate is $U = 0.005$ whereas the sequence size L and epistasis parameter K are both varied such that the correlation is $\gamma = 0.5$. In the right panel, the sequence size is $L = 8$ and the epistasis parameter is set at $K = 3$. The Hamming distance from starting points to the global optimum is five, $d_{GO} = 5$. The simulation data refers to an average over 10 distinct fitness landscapes, and 10 random starting points for each landscape. The dashed-lines correspond to the predictability with respect to ending points for random adaptation walks (RAW) and for S-weighted walks (SWW).

small, predictability becomes independent of the mutation rate. To understand this effect, we additionally simulated two versions of adaptive walks: random adaptation walks (RAW) and the S-weighted walks (SWW) [26,49]. In the former, the walker randomly chooses one of its fitter neighbors, whereas in the SWW version the next step is chosen with probability proportional to the fitness advantage of its fitter neighbors. From these results, we find that in the limit $N_0/N_f \ll 1$ the predictability $P_{2,ending}$ lies between the two adaptive walk variants. We hypothesize that when the population size and mutation rate are sufficiently small, the simulated populations exist in the strong-selection-weak-mutation regime [50], and thus the dynamics are well-approximated by adaptive walk dynamics. In particular, the predictability with respect to the ending points, $P_{2,ending}$ at low mutation rates is sensitive to the earliest stages of adaptation, in which the selective effects of mutations are expected to be larger. This explains why the dynamics in those limits are well captured by adaptive walk dynamics.

In all of the scenarios illustrated here, predictability exhibits a monotonic dependence on the bottleneck size N_0 , meaning that the underlying dynamics become increasingly deterministic for larger populations, not only with respect to the paths but also with respect to the ending points. We return to this result in the Discussion.

E. Multidimensional scaling

Information visualization techniques allow for diverse displays of multidimensional data sets, and offer powerful approaches for visualizing patterns that are otherwise obscured in the data [51]. The visualization of multidimensional data requires a reduction of dimensionality and clustering meth-

ods, thus providing two- or three-dimensional representations [52,53]. Multidimensional scaling is one such technique in which multidimensional objects are positioned relative to each other in a lower dimensional space by interpreting the dissimilarities between objects as distances [54].

Recall that $d(q_\alpha, q_\beta)$ denotes the divergence (a measure of dissimilarity) between evolutionary paths q_α and q_β . We can therefore apply multidimensional scaling to the matrix of distances, $D = [d(q_\alpha, q_\beta)]$, over a bundle of evolutionary paths, where $\alpha, \beta = 1, \dots, N$.

Figure 7 provides two-dimensional representations of the ensemble of distinct evolutionary pathways collected in one thousand replicates, each starting with an isogenic population at the antipode of S_{max} , and ending at the global optimum S_{max} of the fitness landscape. It is important to emphasize that the x_1 and x_2 coordinates have no physical meaning as the positions in the plot are defined by the eigenvectors and eigenvalues of the reference matrix D (for more details about classical multidimensional scaling, please see Ref. [54]). In the upper panels of Fig. 7, N_f is set at $N_f = 4096$, whereas $N_0 = 64$ (left panel) and $N_0 = 2048$ (right panel). When the bottleneck size, N_0 , is increased, the bundle of evolutionary paths becomes less widespread over genotype space, and some paths are more frequently used (paths that are used with frequency of 5% or higher are highlighted). These two facts—that the paths are less scattered over genotype space, and that the frequency of the most-used paths increases—explain both the increase in predictability and decrease in the mean path divergence seen in the lower panels of Fig. 5.

The lower panels of Fig. 7 show analogous results if the bottleneck size $N_0 = 32$ is fixed while N_f is increased. On the left, $N_f = 128$, and the effective population size is correspondingly small. In this case the distribution of evolutionary




FIG. 7. Multidimensional scaling plot of the evolutionary pathways. In the upper panels N_f is set at $N_f = 4096$, whereas $N_0 = 64$ (left upper panel) and $N_0 = 2048$ (right upper panel). In the lower panels, N_0 is set at $N_0 = 32$, whereas $N_f = 128$ (left lower panel) and $N_f = 4096$ (right lower panel). The data correspond to a fixed fitness landscape with epistasis parameter $K = 2$. Those evolutionary pathways that achieve a frequency higher than 0.05 are highlighted in the plot (dark circles, with numbers indicating path frequency).

paths is quite diffuse. Moreover, none of the paths reach a frequency of 5%. In contrast, when N_f is large ($N_f = 4096$, right panel), the distribution of trajectories becomes substantially more compact. We note that three paths achieve high frequencies; in particular, the most frequent path is used in nearly 33% of the independent runs. This robust change is likewise captured by the measure of predictability in the upper panels of Fig. 5. Note that the predictability for $N_f = 4096$ is over ten-fold larger than that estimated for $N_f = 128$.

V. DISCUSSION

An understanding of population bottlenecks is critical to evolutionary biology, owing to their powerful influence on genetic diversity and the evolution of species [55,56]. There is also compelling evidence that bottleneck events shaped the distribution of modern human diversity [57–59]. Mathematical modeling of population bottlenecks, in particular, can be valuable in the development of optimal strategies to curb the adaptive evolution of pathogens [60,61]. In this contribution, we used an epistatic fitness landscape to explore two hypotheses: whether the adaptation rate is maximized by population bottlenecks at an intermediate bottleneck ratio, and whether population bottlenecks reduce the predictability of adaptation.

One of the most compelling results we obtain here is that adaptive trajectories in populations experiencing regular bottlenecks can be reconciled when time is scaled in units of generations (in our case, population doublings). We demonstrate that the time course of fitness increase, fitness variance and genetic diversity are all insensitive to population bottlenecks, provided the bottleneck size exceeds a few thousand

individuals, when time is expressed in population doublings (Figs. 3 and 4). Thus, in contrast with previous predictions that adaptation *per bottleneck* is fastest at intermediate bottleneck ratios, we demonstrate that the adaptation rate *per generation* is simply maximized by experimental protocols that maximize the supply of beneficial mutations, that is, by the largest effective population size (Fig. 2, right panels).

Overall these results imply that the “natural” time unit for adaptation is generations, irrespective of the number of generations that elapse between population bottlenecks, as long as the bottleneck size is not too small. We demonstrate that small bottleneck sizes can retard adaptation through the elimination of rare lineages, but this effect disappears when N_0 is several thousand, rather than tens of individuals. Thus, for most microbial populations, the adaptation rate *per generation* will be largely insensitive to the bottleneck ratio, and maximized by maximizing the supply rate of beneficial mutations.

As well as investigating the role of population bottlenecks at the macroscopic level by examining the rate of adaptation, we also investigated their role at a microscopic level by tracking the predictability and divergence of evolutionary trajectories. From this perspective, we observe a higher level of determinism when either the bottleneck size N_0 or the population size at the end of the growth phase N_f is increased. In other words, predictability is maximized in larger populations, irrespective of population bottlenecks. Similar to the adaptation process, we observe a pronounced increase in predictability, P_2 , with the bottleneck size N_0 when N_0 is small, while predictability saturates as N_0 increases.

More generally, the increases in predictability that we observe in larger populations can be ascribed to an increased

determinism in the underlying process due to the combined effects of three factors: more beneficial mutations are generated, these mutations have longer times to increase in frequency, and there is greater interference among mutant lineages, thus promoting those that confer larger selective advantage.

Although it is quite difficult to estimate the effective population size for complex fitness landscapes, previous studies relying on *stairway to heaven* landscapes provide some hints on the dependence of the effective population size on the parameters of the model [45]. Those previous results, however, are limited by the assumptions of a constant beneficial mutation rate and an invariant distribution of selective effects over the adaptation process—assumptions that are violated in the pattern of diminishing returns we observe here. Nonetheless, this previous work supports our natural expectation that the effective population size is directly proportional to both N_0 and τ , and thus increases with N_f . Therefore, increasing either N_0 or N_f , in fact, increases the effective population size, and thus makes the evolutionary process more deterministic [62].

In all the scenarios investigated here, we observed a monotonic dependence of predictability on both N_0 and N_f . In contrast, a nonmonotonic dependence of predictability measures on population size has been observed in both an empirical and simulated fitness landscape [3]. We note that the increase in predictability with population size observed

in Ref. [3] was due to increased clonal interference among beneficial mutations, while the reduction in predictability at higher population sizes was due to the appearance of second-step mutations, increasing when the value of NU^2 was in the range 10^{-6} – 10^{-7} . For computational efficiency in the very large fitness landscapes simulated here, we have studied comparatively large mutation rates, in which the value of NU^2 is typically several orders of magnitude larger than this threshold, and thus the appearance of second-step mutations is practically assured [with the mutation-limited dynamics described for some cases in Fig. 6(B) as an exception]. We hypothesize that in principle, as the mutation rate increases, alternating regimes in which predictability increases or decreases with population size may be possible, as higher-order mutational neighbourhoods become newly accessible. This would be a clear avenue for future work.

ACKNOWLEDGMENTS

O.F. is supported by Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). P.R.A.C. is supported by CNPq under Projects No. 302569/2018-9 and No. 406594/2018-0. L.M.W. is funded by the Natural Sciences and Engineering Research Council of Canada under Grant No. RGPIN-2019-06294.

-
- [1] J. A. G. De Visser and J. Krug, Empirical fitness landscapes and the predictability of evolution, *Nat. Rev. Genet.* **15**, 480 (2014).
 - [2] M. Lässig, V. Mustonen, and A. M. Walczak, Predicting evolution, *Nature Ecol. Evol.* **1**, 0077 (2017).
 - [3] I. G. Szendro, J. Franke, J. A. G. de Visser, and J. Krug, Predictability of evolution depends nonmonotonically on population size, *Proc. Natl. Acad. Sci. USA* **110**, 571 (2013).
 - [4] J. S. LeClair and L. M. Wahl, The impact of population bottlenecks on microbial adaptation, *J. Stat. Phys.* **172**, 114 (2018).
 - [5] H. A. Orr, The genetic theory of adaptation: A brief history, *Nat. Rev. Genet.* **6**, 119 (2005).
 - [6] I. Fragata, A. Blanckaert, M. A. D. Louro, D. A. Liberles, and C. Bank, Evolution in the light of fitness landscape theory, *Trends Ecol. Evol.* **34**, 69 (2019).
 - [7] N. C. Wu, L. Dai, C. A. Olson, J. O. Lloyd-Smith, and R. Sun, Adaptation in protein fitness landscapes is facilitated by indirect paths, *Elife* **5**, e16965 (2016).
 - [8] C. Bank, S. Matuszewski, R. T. Hietpas, and J. D. Jensen, On the (un) predictability of a large intragenic fitness landscape, *Proc. Natl. Acad. Sci. USA* **113**, 14085 (2016).
 - [9] D. L. Hartl, What can we learn from fitness landscapes? *Curr. Opin. Microbiol.* **21**, 51 (2014).
 - [10] D. M. Weinreich, N. F. Delaney, M. A. DePristo, and D. L. Hartl, Darwinian evolution can follow only very few mutational paths to fitter proteins, *Science* **312**, 111 (2006).
 - [11] Z. D. Blount, J. E. Barrick, C. J. Davidson, and R. E. Lenski, Genomic analysis of a key innovation in an experimental *Escherichia coli* population, *Nature* **489**, 513 (2012).
 - [12] L. Perfeito, A. Sousa, T. Bataillon, and I. Gordo, Rates of fitness decline and rebound suggest pervasive epistasis, *Evolution* **68**, 150 (2013).
 - [13] E. R. Jerison and M. M. Desai, Genomic investigations of evolutionary dynamics and epistasis in microbial evolution experiments, *Curr. Opin. Genet. Dev.* **35**, 33 (2015).
 - [14] H. A. Wichman, M. R. Badgett, L. A. Scott, C. M. Boulianne, and J. J. Bull, Different trajectories of parallel evolution during viral adaptation, *Science* **285**, 422 (1999).
 - [15] S. F. Elena and R. E. Lenski, Evolution experiments with microorganisms: The dynamics and genetic bases of adaptation, *Nat. Rev. Genet.* **4**, 457 (2003).
 - [16] C. R. Miller, P. Joyce, and H. A. Wichman, Mutational effects and population dynamics during viral adaptation challenge current models, *Genetics* **187**, 185 (2011).
 - [17] P. J. Gerrish and R. E. Lenski, The fate of competing beneficial mutations in an asexual population, *Genetics* **102**, 127 (1998).
 - [18] H. A. Orr, The rate of adaptation in asexuals, *Genetics* **155**, 961 (2000).
 - [19] P. R. A. Campos and V. M. d. Oliveira, Mutational effects on the clonal interference phenomenon, *Evolution* **58**, 932 (2007).
 - [20] M. M. Desai and D. S. Fisher, Beneficial mutation-selection balance and the effect of linkage on positive selection, *Genetics* **176**, 1759 (2007).
 - [21] S.-C. Park and J. Krug, Clonal interference in large populations, *Proc. Natl. Acad. Sci. USA* **104**, 18135 (2007).
 - [22] A. S. Perelson and C. A. Macken, Protein evolution on partially correlated landscapes, *Proc. Natl. Acad. Sci. USA* **92**, 9657 (1995).
 - [23] H. A. Orr, The population genetics of adaptation on correlated fitness landscapes: The block model, *Evolution* **60**, 1113 (2006).
 - [24] S. A. Kauffman, *The Origins of Order* (Oxford University Press, Oxford, 1993).

- [25] C. A. Macken and A. S. Perelson, Protein evolution on rugged landscapes, *Proc. Natl. Acad. Sci. USA* **86**, 6191 (1989).
- [26] S. M. Reia and P. R. Campos, Analysis of statistical correlations between properties of adaptive walks in fitness landscapes, *R. Soc. Open Sci.* **7**, 192118 (2020).
- [27] T. Wein and T. Dagan, The effect of population bottleneck size and selective regime on genetic diversity and evolvability in bacteria, *Genome Biol. Evol.* **11**, 3283 (2019).
- [28] L. Garoff, F. Pietsch, D. L. Huseby, T. Lilja, G. Brandis, and D. Hughes, Population bottlenecks strongly influence the evolutionary trajectory to fluoroquinolone resistance in *Escherichia coli*, *Mol. Biol. Evol.* **37**, 1637 (2020).
- [29] L. M. Wahl, P. J. Gerrish, and I. Saika-Voivod, Evaluating the impact of population bottlenecks in experimental evolution, *Genetics* **162**, 961 (2002).
- [30] S. Wright, Evolution in mendelian populations, *Genetics* **16**, 97 (1931).
- [31] D. R. Gifford, S. E. Schoustra, and R. Kassen, The length of adaptive walks is insensitive to starting fitness in *aspergillus nidulans*, *Evolution: Int. J. Organ. Evol.* **65**, 3070 (2011).
- [32] J. de Lima Filho, F. Moreira, P. Campos, and V. M. De Oliveira, Adaptive walks on correlated fitness landscapes with heterogeneous connectivities, *J. Stat. Mech.: Theory Exp.* (2012) P02014.
- [33] A. E. Lobkovsky and E. V. Koonin, Replaying the tape of life: quantification of the predictability of evolution, *Front. Gen.* **3**, 246 (2012).
- [34] A. E. Lobkovsky, Y. I. Wolf, and E. V. Koonin, Predictability of evolutionary trajectories in fitness landscapes, *PLoS Comput. Biol.* **7**, e1002302 (2011).
- [35] S. W. Roy, Probing evolutionary repeatability: neutral and double changes and the predictability of evolutionary adaptation, *PloS One* **4**, e4500 (2009).
- [36] M. Manhart and A. V. Morozov, Statistical physics of evolutionary trajectories on fitness landscapes, in *First-passage Phenomena and Their Applications* (World Scientific, Singapore, 2014), pp. 416–446.
- [37] J. Irwin, On the frequency distribution of the means of samples from a population having any law of frequency with finite moments, with special reference to Pearson's type II, *Biometrika* **19**, 225 (1927).
- [38] P. Hall, The distribution of means for samples of size n drawn from a population in which the variate takes values between 0 and 1, all such values being equally probable, *Biometrika* **19**, 240 (1927).
- [39] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.103.042415> for additional figures.
- [40] I. Gordo and P. R. Campos, Evolution of clonal populations approaching a fitness peak, *Biology Lett.* **9**, 20120239 (2013).
- [41] J. Hubbarde and L. Wahl, Estimating the optimal bottleneck ratio for experimental evolution: The burst-death model, *Math. Biosci.* **213**, 113 (2008).
- [42] R. A. Fisher, *The Genetical Theory of Natural Selection* (Clarendon Press, Oxford, UK, 1930).
- [43] A. Chao, L. Jost, T. Hsieh, K. Ma, W. B. Sherwin, and L. A. Rollins, Expected shannon entropy and shannon differentiation between subpopulations for neutral genes under the finite island model, *PloS One* **10**, e0125471 (2015).
- [44] M. O. Hill, Diversity and evenness: A unifying notation and its consequences, *Ecology* **54**, 427 (1973).
- [45] P. R. Campos and L. Wahl, The effects of population bottlenecks on clonal interference, and the adaptation effective population size, *Evol.: Int. J. Organ. Evol.* **63**, 950 (2009).
- [46] P. R. A. Campos and L. M. Wahl, The adaptation rate of asexuals: Deleterious mutations, clonal interference and population bottlenecks, *Evolution* **64**, 1973 (2010).
- [47] P. R. Campos, C. Adami, and C. O. Wilke, Optimal adaptive performance and delocalization in nk fitness landscapes, *Physica A* **304**, 495 (2002).
- [48] L. Ferretti, B. Schmiegelt, D. Weinreich, A. Yamauchi, Y. Kobayashi, F. Tajima, and G. Achaz, Measuring epistasis in fitness landscapes: The correlation of fitness effects of mutations, *J. Theor. Biol.* **396**, 132 (2016).
- [49] H. A. Orr, A minimum on the mean number of steps taken in adaptive walks, *J. Theor. Biol.* **220**, 241 (2003).
- [50] J. H. Gillespie, A simple stochastic gene substitution model, *Theore. Popul. Biol.* **23**, 202 (1983).
- [51] C. Ware, *Information Visualization: Perception for Design* (Morgan Kaufmann, Burlington, MA, 2019).
- [52] L. Van Der Maaten, E. Postma, and J. Van den Herik, Dimensionality reduction: A comparative review, *J. Mach. Learn. Res.* **10**, 13 (2009).
- [53] J. T. Machado and A. M. Lopes, Multidimensional scaling and visualization of patterns in prime numbers, *Commun. Nonlinear Sci. Numer. Simul.* **83**, 105128 (2020).
- [54] M. A. Cox and T. F. Cox, Multidimensional scaling, in *Handbook of Data Visualization* (Springer, Berlin, 2008), pp. 315–347.
- [55] E. Mayr *et al.*, *Animal Species and Evolution* (Harvard University Press, Cambridge, 1963).
- [56] M. Slatkin, In defense of founder-flush theories of speciation, *Am. Nat.* **147**, 493 (1996).
- [57] A. R. Rogers and H. Harpending, Population growth makes waves in the distribution of pairwise genetic differences, *Mol. Biol. Evol.* **9**, 552 (1992).
- [58] W. Amos and J. I. Hoffman, Evidence that two main bottleneck events shaped modern human genetic diversity, *Proc. R. Soc. B: Biol. Sci.* **277**, 131 (2010).
- [59] N. J. Fagundes, A. Tagliani-Ribeiro, R. Rubicz, L. Tarskaia, M. H. Crawford, F. M. Salzano, and S. L. Bonatto, How strong was the bottleneck associated to the peopling of the americas? New insights from multilocus sequence data, *Genet. Mol. Biol.* **41**, 206 (2018).
- [60] J. T. McCrone and A. S. Lauring, Genetic bottlenecks in intraspecies virus transmission, *Curr. Opin. Virol.* **28**, 20 (2018).
- [61] E. Rousseau, M. Bonneault, F. Fabre, B. Moury, L. Mailleret, and F. Grognaud, Virus epidemics, plant-controlled population bottlenecks and the durability of plant resistance, *Philos. Trans. R. Soc. B* **374**, 20180263 (2019).
- [62] D. E. Rozen, M. G. J. L. Habets, A. Handel, and J. A. G. M. de Visser, Heterogeneous adaptive trajectories of small populations on complex fitness landscapes, *PloS One* **3**, e1715 (2008).