

A dynamical limit to evolutionary adaptation

Matthew J. Melissa^a and Michael M. Desai^{1 a}

^aDepartment of Organismic and Evolutionary Biology, Department of Physics, Quantitative Biology Initiative, and NSF-Simons Center for Mathematical and Statistical Analysis of Biology, Harvard University, Cambridge MA 02138

This manuscript was compiled on July 28, 2024

Natural selection makes evolutionary adaptation possible even if the overwhelming majority of new mutations are deleterious. However, in rapidly evolving populations where numerous linked mutations occur and segregate simultaneously, clonal interference and genetic hitchhiking can limit the efficiency of selection, allowing deleterious mutations to accumulate over time. This can in principle overwhelm the fitness increases provided by beneficial mutations, leading to an overall fitness decline. Here, we analyze the conditions under which evolution will tend to drive populations to higher versus lower fitness. Our analysis focuses on quantifying the boundary between these two regimes, as a function of parameters such as population size, mutation rates, and selection pressures. This boundary represents a state in which adaptation is precisely balanced by Muller's ratchet, and we show that it can be characterized by rapid molecular evolution without any net fitness change. Finally, we consider the implications of global fitness-mediated epistasis, and find that under some circumstances this can drive populations towards the boundary state, which can thus represent a long-term evolutionary attractor.

rapid evolution | evolutionary attractor | Muller's ratchet | fitness-mediated epistasis

Evolution is often thought of as an optimization process, in which natural selection pushes populations inevitably uphill, towards a local optimum in the fitness landscape (1). However, much recent work has shown that in many populations, numerous linked mutations often arise and segregate simultaneously (2–9). In these *rapidly evolving* populations, natural selection is much less efficient: it cannot act on each mutation independently (10). As a result, deleterious mutations can often fix, which can slow down adaptation or even reverse its direction, leading to declining fitness over time.

Extensive previous work has studied the accumulation of deleterious mutations via Muller's ratchet (11, 12), particularly in models in which beneficial mutations are either negligible or can be treated as a rare perturbation (13–15). Similarly, numerous studies have considered the accumulation of beneficial mutations (i.e. adaptation) when deleterious mutations are absent (16–18) or can be treated as a perturbation (19, 20). However, we lack an understanding of the interplay between the accumulation of beneficial and deleterious mutations more generally. Except in special cases (e.g. when clonal interference is absent (21), or when all beneficial and deleterious mutations have the same fitness effect (22), or in a regime where fixations are dominated by single driver mutations (23)), this has made it impossible to answer a very basic question: given a particular set of population genetic parameters (population size, mutation rate, and fitness landscape), will a population tend to increase or decrease in fitness? In other words, under what circumstances can evolution act as an optimization process, and when do populations actually move towards less-optimal genotypes?

Here, we analyze this interplay between beneficial and

deleterious mutations in rapidly evolving populations, in the regime where both types of mutations can be important. Our analysis leverages recent work in traveling wave models of evolutionary dynamics, and in particular our recently introduced *moderate selection, strong-mutation* (MSSM) approximation (24). Using this approach, we predict the conditions under which populations will tend to increase or decrease in fitness (i.e. where the rate of change in mean fitness, v , is positive or negative). The boundary surface between these two regions of the parameter space, at which $v = 0$, corresponds to a state in which beneficial and deleterious mutations accumulate in a balanced way. While the fitness trajectory of a population in the $v = 0$ state appears neutral, the evolutionary dynamics of these populations can be strongly nonneutral (as has been suggested by several earlier studies; see e.g. (21–23, 25)). For example, a steady state accumulation of weakly deleterious mutations may be offset by the fixation of beneficial mutations under moderate or strong selection. We also consider additional surfaces of the parameter space on which patterns of molecular divergence and genetic diversity would suggest a population has evolved neutrally or nearly neutrally, but in fact mask a balance between the competing signatures of positive and negative selection.

We conclude by considering how our results and the structure of the fitness landscape determine the long-term outcomes of evolution. For example, it is natural to expect that beneficial mutations become less common (and deleterious mutations more common) as a population increases in fitness. This will tend to lead a population not towards a local optimum, but instead towards the $v = 0$ state (see e.g. Goyal et. al. (22)

Significance Statement

We often think of evolution as an optimization process that increases the fitness of a population over time through the action of natural selection on random mutations. However, this process is imperfect, and deleterious mutations can occasionally spread through a population. Depending on the mutation rates and selection pressures involved, this can overwhelm adaptation and lead to fitness decline. Here, we calculate the conditions under which natural selection can drive a population to higher fitness, and characterize the scenarios that lead to fitness declines. Our results quantify the limits on the ability of natural selection to drive populations towards a fitness optimum. We analyze how these constraints can lead populations to steady-state evolutionary “attractors” that are not fitness optima.

M.J.M. and M.M.D. designed the project; M.J.M. conducted the simulations and analysis; M.J.M. and M.M.D. wrote the paper.

The authors declare no conflicts of interest.

¹To whom correspondence should be addressed. E-mail: mdesai@oeb.harvard.edu

and Schiffels et al. (23)). More generally, recent empirical work has identified a consistent pattern of diminishing returns epistasis: beneficial mutations tend to have weaker effects as populations increase in fitness (26–28). An analogous pattern for epistasis on deleterious mutations is less clear, but recent work has identified a trend in which deleterious mutations are more costly in more-fit backgrounds (29, 30). We show here that, depending on details of the landscape and the starting point, these and other patterns of fitness-mediated epistasis can often (but not always) drive a population towards the $v = 0$ state. Thus the $v = 0$ state can in some circumstances represent a long-term evolutionary attractor, and define the extent to which evolution can act to optimize fitness. We refer to this as a *dynamical limit on evolutionary adaptation*, because it arises due to the stochastic nature of the evolutionary process, and can be reached far from any fitness peak.

Model

We model the evolution of a population of N haploid individuals, in which random mutations arise within a specific genomic region at a total rate U . We assume that recombination can be neglected within this region on the relevant timescales. We assume that each new mutation confers a fitness effect, s , drawn from some distribution of fitness effects (DFE), $\rho(s|\vec{g})$, that depends on the genotype $\vec{g}$ of an individual as well as the environment the population evolves in (its fitness landscape). The DFE includes both beneficial and deleterious mutations, with beneficial mutations corresponding to $s > 0$ and deleterious mutations corresponding to $s < 0$. To be more precise, a mutation with effect s increments an individual's (log) fitness X by an amount s , and we assume offspring numbers are drawn from a multinomial distribution each generation; the expected offspring number of an individual with fitness X is $e^{X-\bar{X}}$, where $\bar{X}$ denotes the mean fitness of the population.

The genotype-dependence of $\rho(s|\vec{g})$ has been termed *macroscopic epistasis* (31). This macroscopic epistasis can arise due to individual *microscopic epistatic* interactions among specific mutations, which collectively determine the overall DFE for a given genotype. We make the key assumption that similar genotypes share a similar $\rho(s|\vec{g})$, and in particular, that those genotypes simultaneously present in a population (which are similar because of their relatedness by common ancestry) share the same DFE, $\rho(s)$. This allows us to solve for the dynamics by treating $\rho(s)$ instantaneously as a constant parameter. We note that this assumption can be satisfied even in the presence of pervasive microscopic epistasis, as long as idiosyncratic interactions among mutations largely “average out” in contributing to the full distribution $\rho(s|\vec{g})$.

For simplicity, we focus on a few simplifying forms of $\rho(s)$ in our analysis. For instance, we consider the case where all beneficial mutations have effect s_b and all deleterious mutations have effect $-s_d$ (with $s_b, s_d > 0$ by convention). This example is useful for building general intuition, and is motivated by recent work showing that the evolutionary dynamics of rapidly evolving populations can in many cases be well-captured by a DFE consisting of a single appropriately-chosen “predominant” effect size (17, 18, 32). We also consider the cases of exponentially-distributed (and more generally gamma-distributed) effects of beneficial and of deleterious mutations, though our analysis can be extended to more general DFEs relatively straightforwardly. Importantly, we make no assumption

that the DFEs of beneficial mutations and of deleterious mutations are the same or similar in shape or in scale.

Results

The central goal of our analysis is to determine whether a population will tend to increase or decrease in fitness for a given set of parameters: the population size, N , the mutation rate, U , and the distribution of fitness effects, $\rho(s)$. Because our goal is to determine whether $v > 0$ or $v < 0$ for a given set of parameters, we focus on analyzing the boundary between these two regimes. This boundary is by definition a $v = 0$ surface where the mean fitness of the population does not on average either increase or decrease. We will find it useful to write the average rate of change in mean fitness, v , in terms of the fixation probability of a new mutation, $p_{\text{fix}}(s)$,

$$v = NU \int \rho(s) s p_{\text{fix}}(s) ds. \quad [1]$$

To find the $v = 0$ surface we then set Eq. (1) equal to 0, which gives a constraint on the parameters N , U , and $\rho(s)$ that defines the $v = 0$ surface in parameter space.

We can clearly have $v = 0$ if selection on deleterious mutations is sufficiently strong and beneficial mutations are sufficiently rare that no selected mutations fix at all, and the evolutionary dynamics are entirely neutral (i.e. if $\rho(s)p_{\text{fix}}(s)$ is negligible for all s). Apart from this trivial case, the $v = 0$ state by definition involves substantial accumulation of deleterious mutations (at least relative to the accumulation of beneficial mutations) which can be facilitated by the effects of linked selection and clonal interference. For instance, deleterious mutations may routinely hitchhike along with, or hinder the fixation of, a beneficial mutation (33–35). Interference among multiple beneficial mutations may also substantially reduce the rate at which they can fix in the population (10, 36). To obtain an accurate description of the $v = 0$ state, our expression for the fixation probabilities $p_{\text{fix}}(s)$ must therefore take these effects into account.

Frequent interference among mutations is a defining feature of *rapid evolution*, which has been the focus of much recent theoretical work (13, 18, 20, 37–40). Broadly speaking, this work uses traveling wave models, which first analyze the steady-state distribution of fitness within the population (the “traveling wave of fitness”), and then use this as the basis for computing the fixation probabilities of new mutations, $p_{\text{fix}}(s)$, and the average rate of fitness increase or decline, v . Most work on traveling wave models has been done by considering only beneficial mutations (16–18) or only deleterious mutations (14), or by focusing on one type and treating the other perturbatively (20). A key exception is the *moderate selection, strong-mutation* (MSSM) approximation we have recently introduced (24), which can be applied to analyze rapidly evolving populations for which both beneficial and deleterious mutations affect the dynamics in a substantial way. Here, we use this MSSM approximation to analytically describe the $v = 0$ state.

A key result is that within the MSSM regime (which we discuss below), the fixation probability of a new mutation is given by

$$p_{\text{fix}}(s) \approx \frac{e^{T_c s}}{N}, \quad [2]$$

valid for both positive and negative s . Here T_c is a derived quantity whose definition and relationship to the parameters N , U and $\rho(s)$ we reproduce in *SI Appendix*, and which approximately equals $\langle T_2 \rangle / 2$ —one-half the average time since two randomly chosen individuals share a common ancestor (i.e., a *coalescence* timescale). We note that an identical result for $p_{\text{fix}}(s)$ was derived in a somewhat different context by Hallatschek (41). We also note that in much of the population genetics literature, the average pairwise coalescent time $\langle T_2 \rangle$ is identified with an effective population size N_e . We avoid this language here because in general, the evolutionary dynamics in rapidly evolving populations are not equivalent to those in a neutrally evolving population for any choice of N_e (42).

Eq. (2) differs from the standard formula for the fixation probabilities of independently evolving loci (43), which in our notation can be written as

$$p_{\text{fix}}(s) = \frac{2T_c}{N} \times \frac{2s}{1 - e^{-4T_c s}}. \quad [3]$$

Eq. (3) has been used by Whitlock (21) to address similar questions, although that work treats T_c (referred to as $N_e/2$) as an independent parameter, instead of considering how it depends on the population parameters N , U , and $\rho(s)$. For the sake of comparison, we discuss the predictions following from Eq. (3) alongside our results below. The predictions are qualitatively (and even quantitatively) similar in some respects, but they break down in other cases. This is unsurprising in light of recent work that has shown that Eq. (3) fails to adequately describe the fixation probabilities of mutations in the presence of widespread linked selection, particularly when mutations confer fitness effects on a wide range of scales (44).

Eq. (2) immediately implies that if we scale fitness effects to the coalescence timescale by defining $\gamma \equiv T_c s$, the $v = 0$ surface is defined by the concise equation

$$\int \tilde{\rho}(\gamma) \gamma e^{\gamma} d\gamma = 0. \quad [4]$$

Eq. (4) implies that we can characterize the $v = 0$ surface given only the distribution of “scaled” fitness effects, $\tilde{\rho}(\gamma)$ (as well as validity of the MSSM approximation, which we discuss below). We emphasize that this is not by itself sufficient to determine how the $v = 0$ surface depends on the underlying parameters, because T_c depends in a nontrivial way on N , U , and $\rho(s)$. We return to this dependence in more detail below. However, in the next section we first analyze key properties of the $v = 0$ surface in the space of scaled selective effects, focusing particular attention on two specific choices of the DFE as representative examples.

The $v = 0$ surface in the space of scaled effects. In the scaled parameter space, the $v = 0$ surface depends only on $\tilde{\rho}(\gamma)$, and not on the population size N or the mutation rate U (which enter only through their effect in determining T_c). To gain qualitative insight, we begin by considering the simple case in which all beneficial mutations confer a single scaled effect, $\gamma_b \equiv T_c s_b$, and all deleterious mutations confer a (potentially different) single scaled effect, $-\gamma_d \equiv -T_c s_d$ (with $s_b, s_d > 0$ by convention). Specifically, we have $\tilde{\rho}(\gamma) = [\eta/(1 + \eta)] \delta(\gamma - \gamma_b) + [1/(1 + \eta)] \delta(\gamma + \gamma_d)$, where we have defined $\eta \equiv U_b/U_d$ as the ratio of beneficial to deleterious mutation rates. Plugging this into Eq. (4), we find that within the 3-dimensional parameter




Fig. 1. Cross sections of the $v = 0$ surface in the space of scaled effects γ_d vs. γ_b , for the four values of η denoted above. Each point corresponds to a simulated parameter combination, colored by its measured value of $T_c^2 v$. Solid line denotes the prediction for the $v = 0$ curve given by Eq. (4). Dashed lines denotes the prediction for the $v = 0$ curve obtained using Eq. (3). Dotted lines denote the lines on which $\eta = \gamma_d/\gamma_b$. Simulated parameter combinations lie on the grids of logarithmically spaced Ns_d and Ns_b values depicted in Fig. 2, with DFEs consisting of a single beneficial effect and a single deleterious effect, and with $NU = 10^4$.

space spanned by γ_b , γ_d , and η , the $v = 0$ constraint is a 2-dimensional surface given by

$$\eta|_{v=0} = \frac{\gamma_d e^{-\gamma_d}}{\gamma_b e^{\gamma_b}}. \quad [5]$$

In Fig. 1 we validate this prediction for the $v = 0$ surface. To do so, we conducted Wright-Fisher simulations for populations whose parameters lie on a grid with varying η , Ns_b and Ns_d . Each simulated population is plotted using its corresponding value of T_c , measured by observing its pairwise neutral heterozygosity π_{neu} averaged over simulation runs (with T_c taken as $\langle T_2 \rangle / 2 = \pi_{\text{neu}} / (4U_n)$, where U_n is the neutral mutation rate used in simulations). The prediction in Eq. (5) qualitatively (and except perhaps for $\gamma_d \gg 1$, quantitatively) describes the $v = 0$ surface in the space of scaled fitness effects γ_b and γ_d . The simulations represented in Fig. 1 are all conducted for populations with $NU = 10^4$. In Fig. S1, we present the results of additional simulations which include the cases $NU = 10^3$ and $NU = 10^2$; similar agreement is obtained.

Several qualitative features of Eq. (5) are notable. If the selective effects of both beneficial and deleterious mutations are small compared to $1/T_c$ (i.e. $\gamma_b \ll 1$ and $\gamma_d \ll 1$), we see from Eq. (5) that the $v = 0$ surface is defined by $\eta|_{v=0} \approx \gamma_d/\gamma_b$ (dotted line in Fig. 1; note this converges with the $v = 0$ surface observed in simulations when $\gamma_b \ll 1$ and $\gamma_d \ll 1$). This corresponds to $U_b s_b = U_d s_d$, the surface on which $v = 0$ if beneficial and deleterious mutations accumulate neutrally, such that each mutation fixes with probability $1/N$. The surface $\eta = \gamma_d/\gamma_b$ can be thought of as an upper bound to the actual $v = 0$ surface; as the strength of selection (i.e. γ_b and/or γ_d) is increased, the actual fraction η required to have $v = 0$ will always be smaller than γ_d/γ_b .

More generally, if we increase γ_b at fixed γ_d , $\eta|_{v=0}$ decreases: a smaller ratio η of beneficial to deleterious mutations is required to be in the $v = 0$ state. This makes intuitive sense: increasing γ_b increases both the fixation probability of

beneficial mutations and the fitness benefit they provide to the population upon fixing. The effect of changing γ_d is more subtle, because while increasing γ_d decreases the fixation probability of deleterious mutations, it also increases the fitness cost to the population they incur upon fixing. This means that the expected fitness costs to the population of deleterious mutations are not monotonic with effect size: for a fixed value of γ_b (below a threshold value described below), the population will adapt for small γ_d , decline in fitness for intermediate γ_d , and adapt for large γ_d . This is because for sufficiently small γ_d , deleterious mutations fix routinely but do not confer large enough effects to counteract the beneficial mutations which fix, while for sufficiently large γ_d , deleterious mutations are purged by selection too efficiently to counteract the fixation of beneficial mutations. Instead, deleterious mutations are maximally impactful (in the sense that $\eta|_{v=0}$ is maximized) at the intermediate scaled effect size $\gamma_d^* = 1$. For sufficiently large γ_b and/or η , deleterious mutations cannot lead to decline in fitness for *any* value of γ_d (although $T_c v$ will still be minimized at $\gamma_d = 1$). For example, at a given η , the population will always adapt provided that $\gamma_b > \gamma_b^*$, where $\gamma_b^* e^{\gamma_b^*} = 1/(e\eta)$. We can think of the curve $(\gamma_d, \gamma_b, \eta) = (\gamma_d^*, \gamma_b^*, \eta)$, parameterized by η , as a “ridgeline” of the $v = 0$ surface, on which η is maximized as a function of γ_d .

The above analysis can be extended straightforwardly to a full distribution of fitness effects. As a simple example, we consider the case in which both beneficial and deleterious mutations are drawn from exponential distributions with mean scaled effects γ_b and γ_d respectively, and a ratio $\eta = U_b/U_d$ of beneficial to deleterious mutations (in the *SI Appendix*, we extend these results to the case of gamma-distributed DFEs, and comment further on arbitrary DFEs). Specifically, we have $\bar{\rho}(\gamma) = [\eta/(1+\eta)/\gamma_b] e^{-\gamma/\gamma_b} + [1/(1+\eta)/\gamma_d] e^{\gamma/\gamma_d}$. Plugging this scaled DFE into Eq. (4), we find that the $v = 0$ surface is defined by

$$\eta = \frac{\gamma_d}{\gamma_b} \left(\frac{1 - \gamma_b}{1 + \gamma_d} \right)^2. \quad [6]$$

We note that in this case, a requirement that $\gamma_b < 1$ dynamically emerges given validity of the MSSM approximation (and thus the η given by Eq. (6) is positive for all relevant γ_b and γ_d) while $\gamma_d > 1$ is permitted. As in the case of a two-effect DFE, a $v = 0$ “ridgeline” exists on which η is maximized as a function of γ_d , and on which $\gamma_d^* = 1$ and $\gamma_b^* = \eta (\sqrt{1 + 1/\eta} - 1)^2$. The behavior of the $v = 0$ surface for large γ_d is qualitatively different, however. Instead of the dependence $\eta \propto \gamma_d e^{-\gamma_d}$, we have the much slower falloff $\eta \propto 1/\gamma_d$ for large γ_d . The dependence $1/\gamma_d$ in turn approximates the fraction of deleterious effects with $|\gamma| < 1$ —that is, the fraction of deleterious effects with a reasonable chance of fixing—as opposed to the fixation probability $\propto e^{-\gamma_d}$ of the average deleterious effect. As a result, compared to the case of a two-effect DFE, $v = 0$ curves for a full DFE have a much more broad decay of γ_b toward zero at large γ_d . In Fig. S1 we compare the prediction in Eq. (6) to the results of simulations for the same η values considered in Fig. 1, and for NU values ranging from 10^2 to 10^4 . Our results are qualitatively similar to the case of single fitness effects shown in Fig. 1, although agreement begins to break down—particularly for $\gamma_d \sim \gamma_d^*$ —for smaller NU and smaller η (i.e. for smaller values of NU_b).

Computing T_c in terms of the population genetic parameters, N and $U\rho(s)$. Above, we described the $v = 0$ constraint on the distribution of scaled effects $\gamma = T_c s$. Expressed in this way, the $v = 0$ constraint takes on the simple analytical form in Eq. (4). In certain cases, the distribution of scaled effects γ is more readily probed than the distribution of unscaled effects s . For instance, given DNA sequencing data from a population at a fixed point in time (e.g. an observation of its site frequency spectra, of both synonymous mutations and nonsynonymous mutations), its distribution of scaled fitness effects can be inferred (45–47). To infer the distribution of unscaled fitness effects s then requires an independent estimate of the coalescence timescale, which is typically confounded with estimates of the neutral mutation rate (48). However, in certain cases—such as in the context of experimentally evolved populations (49, 50), the distribution of unscaled effects s may be more practical to measure. For this reason, it may be more useful to obtain a $v = 0$ constraint on the parameters N and $U\rho(s)$; this can also be useful in building an intuitive understanding for how shifts in these underlying parameters affect adaptation or fitness decline. As we will see in the Discussion, a $v = 0$ constraint on unscaled effects s may also be more useful in working out the long-term implications of simple patterns of fitness-mediated epistasis observed empirically, for which the parameter $U\rho(s)$ may vary in a more straightforward way, over the course of an evolutionary trajectory, than the distribution of scaled effects.

In general, the coalescence timescale T_c depends in a complicated way on the parameters N and $U\rho(s)$, and except in special cases, the mapping between a distribution of unscaled effects Ns and the distribution of scaled effects $T_c s$ is not particularly clear. The MSSM approximation yields a relation between T_c and the underlying parameters N and $U\rho(s)$ —and can thus be used to obtain a $v = 0$ constraint on the parameters N and $U\rho(s)$. We reproduce the relation between T_c , N and $U\rho(s)$ under the MSSM approximation in the *SI Appendix*. The key result is

$$\log Nx_c \approx T_c(x_c - U) - \frac{vT_c^2}{2} + U \int \frac{\rho(s)ds}{s} (e^{T_c s} - 1), \quad [7]$$

along with equations for v and x_c in terms of $U\rho(s)$ and T_c which we reproduce in the *SI Appendix*. Although it is in general straightforward to solve these equations numerically, the $v = 0$ constraint induces a particular simplification in that it becomes possible to solve for $T_c|_{v=0}$ in terms of the single parameter $\rho(s)$, using Eq. (4). A $v = 0$ constraint between N and $U\rho(s)$ then follows straightforwardly from the Eq. (7), although the nature of these equations precludes writing down simple analytical expressions for the $v = 0$ constraint such as Eq. (5).

Validity of the MSSM approximation. The results for $p_{\text{fix}}(s)$ in Eq. (2) and particularly for T_c in Eq. (7) both depend on the validity of the MSSM approximation. We review the conditions of validity of the MSSM approximation in the *SI Appendix*. Roughly speaking, validity of the MSSM approximation requires that selection on individual mutations is at most *moderate* (in that typical fixed fitness effects are smaller than the characteristic range in relative fitnesses from which future common ancestors typically descend) while mutation is *strong* (in that the the dynamics by which individuals fix are




Fig. 2. Cross-sections of the $v = 0$ surface in the space of unscaled effects Ns_d vs. Ns_b , for $NU = 10^4$ and four values of η , with populations colored by simulated values of $T_c^2 v$. Populations in the top row are subject to a two-effect DFE; populations in the bottom row are subject to a two-exponential DFE. Solid lines denote predictions of the $v = 0$ surface obtained by connecting predictions obtained using the MSSM approximation and N_e -based heuristic (that is, using Eq. (9)). Dotted lines denote predictions obtained using the MSSM approximation, for $Ns_d > (Ns_d)^*$; dashed lines denote predictions of the N_e -based heuristic. Stars denote points at which $(\gamma_d, \gamma_b) = (\gamma_d^*, \gamma_b^*)$, as computed by MSSM approximation. Gray squares denote parameter combinations not simulated because $U_d s_d < U_b s_b$, or parameter combinations eliminated from consideration because fewer than 10 epochs were reached during the simulation runtime.

strongly nonneutral and the population is rapidly evolving). Given $U\rho(s)$ and a solution for T_c , it is straightforward to determine whether these conditions are met for a particular point on the $v = 0$ surface.

The points $(\gamma_d^*, \gamma_b^*, \eta)$ lying along the $v = 0$ “ridgeline” are convenient landmarks of $v = 0$ curves at which the validity of the MSSM approximation can be assessed. At these points, selection is sufficiently strong that the MSSM predictions are nontrivial, and the $v = 0$ surface is guaranteed to deviate substantially from the surface $\eta = \gamma_d/\gamma_b$ on which $v = 0$ assuming neutral accumulation. Furthermore, validity of the MSSM approximation at $(\gamma_d^*, \gamma_b^*, \eta)$ implies validity of the approach for smaller values of γ_d and γ_b along the same $v = 0$ curve, up to and including $(\gamma_d^*, \gamma_b^*, \eta)$, for a given η . Note that in the limit $\gamma_b \rightarrow 0$ and $\gamma_d \rightarrow 0$, the population dynamics become purely neutral, and “strong mutation” condition of the MSSM approximation breaks down; however, by this point, the $v = 0$ surface will simply approach its neutral expectation $\eta = \gamma_d/\gamma_b$.

In the *SI Appendix*, we assess the validity of the MSSM approximation at the points $(\gamma_d^*, \gamma_b^*, \eta)$, for the cases of a two-effect DFE and a two-exponential DFE. We find that in both of these cases, the conditions of validity of the MSSM approximation are met provided that $T_c U_b \gg 1$ at $(\gamma_d^*, \gamma_b^*, \eta)$; note that we assume $U_b \leq U_d$ throughout, so that $T_c U_d \gg 1$ is also implied. The condition $T_c U_b \gg 1$ has a relatively simple dynamical interpretation—essentially, that a given lineage will acquire multiple beneficial mutations over the coalescence timescale—and is satisfied at $(\gamma_d^*, \gamma_b^*, \eta)$ for sufficiently large NU (since $T_c U$ increases with NU , if γ_b , γ_d and η are held fixed). This suggests that the MSSM approximation is of broad use in the describing the $\gamma_d < \gamma_d^*$ portion of $v = 0$ curves for rapidly evolving populations in which interference

is important.

While the MSSM approximation is expected to be valid in describing the $\gamma_d < \gamma_d^*$ portion of the $v = 0$ surface, it may break down in the presence of strong deleterious effects $\gamma_d \gg \gamma_d^*$, even if $T_c U_b \gg 1$. Here, however, we will see that an alternative heuristic approach works well: we take

$$T_c \approx N e^{-U_d} \int \frac{\rho_d(s)}{s} (1 - e^{-T_c s}) ds. \quad [8]$$

The form of T_c in Eq. (8) can be motivated by analyzing how strongly deleterious mutations contribute to the determination of T_c in the MSSM approximation (that is, through the final term in Eq. (7)). We note that in using Eq. (8) we neglect any effect of beneficial mutations on the determination of T_c , which is reasonable when beneficial effects are sufficiently weak and deleterious effects sufficiently strong. When deleterious mutations all confer a single deleterious effect with $T_c s_d \gg 1$, T_c in Eq. (8) reduces to the well-known quantity $N e^{-U_d/s_d}$. This is often referred to as N_e , and under certain conditions gives the number of deleterious-mutation-free individuals in the population at equilibrium (51), and in turn the coalescence timescale (52). A similar interpretation can be given if $T_c s \gg 1$ for all possible deleterious effects (53, 54). When instead both weak and strong deleterious effects are possible, the factor $(1 - e^{-T_c s})$ essentially picks out those deleterious effects with $T_c s \gg 1$ (which are important in reducing the effective population size and related coalescence timescale). Here, we refer to this heuristic approach to estimating T_c as an “ N_e -based heuristic”; by substituting this T_c into Eq. (3) we can obtain predictions for $v = 0$ curves in the same way $v = 0$ curves are obtained above using the MSSM approximation.

The $v = 0$ constraint in terms of population genetic parameters. We compare our predictions for the $v = 0$ constraint to




Fig. 3. Comparison between simulations and MSSM predictions for, in panels A and B, the extremal point $((N_{s_b})^*, (N_{s_d})^*)$ and, in panels C and D, the ridgeline point (γ_b^*, γ_d^*) . Each point is obtained from the simulation results depicted in a particular panel of Fig. 2 or Fig. S2 (with values of η denoted on the horizontal axis and values of $NU \in \{10^2, 10^3, 10^4\}$, and with populations subject to a two-effect DFE); for details on how ridgeline and extremal points are extracted from a given panel of simulation results, see *Materials and Methods*. In panels A and B, the top theory curve corresponds to $NU = 10^4$, the middle theory curve corresponds to $NU = 10^3$, and the bottom theory curve corresponds to $NU = 10^2$. Points are colored according to their values of $T_c U_b$ at $(\gamma_d^*, \gamma_b^*, \eta) = (\gamma_d^*, \gamma_b^*, \eta)$, with T_c measured in simulations through levels of pairwise neutral heterozygosity.

the results of simulations in Fig. 2. We can see that the small N_{s_d} and large N_{s_d} portions of $v = 0$ curves are well-described by the MSSM approximation and our N_e -based heuristic, respectively. For concreteness, we connect these two approaches by taking the result of the MSSM approximation for N_{s_d} up to and including $(N_{s_d})^*$, the maximal value of N_{s_d} on a $v = 0$ curve, given η and NU . We therefore consider the following prediction for $v = 0$ curves:

$$N_{s_b} = \begin{cases} (N_{s_b})^{MSSM} & N_{s_d} \leq (N_{s_d})^* \\ \min[(N_{s_b})^*, (N_{s_b})^{NE}] & N_{s_d} > (N_{s_d})^* \end{cases} \quad [9]$$

where $(N_{s_b})^{MSSM}$ and $(N_{s_b})^{NE}$ denote predictions for N_{s_b} along $v = 0$ curves, given the MSSM approximation and N_e -based heuristic, respectively; $(N_{s_d})^*$ and $(N_{s_b})^*$ are predictions for maximal N_{s_d} value and corresponding N_{s_b} at which $v = 0$, obtained using the MSSM approximation. In Fig. 2 we can see that Eq. (9) adequately predict $v = 0$ curves across a range of N_{s_d} values and η values for both the cases of two-effect and two-exponential DFEs. These simulations are conducted with $NU = 10^4$; the results of simulations with $NU = 10^2$ and $NU = 10^3$ are presented in Fig. S2. As is observed for $v = 0$ curves in the space of scaled effects, we note that the $v = 0$ curves are much less “sharp” at large values of N_{s_d} for the cases of two-exponential DFEs, as compared to cases of two-effect DFEs. Note that because at fixed η and NU , T_c varies with N_{s_b} and N_{s_d} , the point in parameter space at which $(\gamma_d, \gamma_b) = (\gamma_d^*, \gamma_b^*)$ is not the same as the point at which $(N_{s_b}, N_{s_d}) = ((N_{s_b})^*, (N_{s_d})^*)$. However, we can see in Fig. 2 and Fig. S2—with the point at which $(\gamma_d, \gamma_b) = (\gamma_d^*, \gamma_b^*)$ denoted by a star marker—that for all cases considered, these points nearly coincide. Thus, validity of the MSSM approximation up to one of these points essentially implies validity up to the other point, which motivates us to use the predictions of the MSSM approximation for N_{s_d} up to $(N_{s_d})^*$ in Eq. (9).

To summarize the accuracy of the MSSM approximation in predicting $v = 0$ curves across the entire range of NU and η values simulated, in Fig. 3 we compare our predictions for $(N_{s_b})^*$, $(N_{s_d})^*$, γ_b^* and γ_d^* to the corresponding quantities obtained from simulations (with details of how these quantities are extracted from simulations provided in *Materials and Methods*). Fig. 3 includes populations which are subject to two-effect DFEs; we provide the same comparison for populations instead subject to two-exponential DFEs in Fig. S3. We can see that these quantities (and thus the weaker-selection portion of $v = 0$ curves) are well-predicted as long as $T_c U_b \gg 1$ at a given point along the $v = 0$ ridgeline.

In Fig. 2, we have chosen to illustrate cross sections of the parameter space spanned by the axes N_{s_b} and N_{s_d} because of the striking nonmonotonicity of $v = 0$ curves in this space (at fixed η and NU). To gain a more complete qualitative understanding of the $v = 0$ constraint, we can consider small perturbations of parameters from the $v = 0$ constraint along each of the two remaining axes, η and NU . The behavior with η is simple: $N^2 v$ increases monotonically with η (with other parameters held fixed), so $v = 0$ curves are shifted (and distorted, to some extent) towards lower N_{s_b} values as η is increased (i.e. with larger η , smaller values of N_{s_b} are needed to have $v = 0$). The behavior with NU is less immediately clear, but also straightforward: $v = 0$ curves are shifted toward larger N_{s_b} values as NU is increased. This reflects the fact that a larger NU value implies more frequent interference, and thus less efficient selection for beneficial mutations and against deleterious mutations. From this dependence, the behavior with U (at fixed η , N , s_b and s_d) also follows: $v = 0$ curves are shifted toward larger N_{s_b} values as U is increased. Note also that these patterns are reflected in the dependence of $(N_{s_b})^*$ on η and NU as shown in Fig. 3, as well as the dependence of full $v = 0$ curves on η and NU , which can be seen in Fig. S2.

Patterns of Molecular Evolution. Despite the fact that the $v = 0$ surface involves no change in the mean fitness of a population over time, the evolutionary dynamics of populations on the $v = 0$ surface can be far from neutral. These populations lie in a dynamic steady state involving accumulation of both beneficial and deleterious mutations, at rates which may differ substantially from the accumulation rate of neutral mutations. The fixation probabilities $p_{fix}(s) \approx e^{T_c s}/N$, along with the MSSM approximation for determining T_c , can be used to characterize the expected total rate F of (selected) mutation fixation of a population, both lying on or off of the $v = 0$ surface. In particular, with analogy to our characterization of a $v = 0$ surface, an $F = U$ surface can be characterized on which selected mutations, on average, accumulate/fix as if they were entirely neutral (with faster-than-neutral accumulation of beneficial mutations precisely balancing slower-than-neutral accumulation of deleterious mutations). If we assume that synonymous mutations are neutral and nonsynonymous mutations are selected, the $F = U$ surface can also be thought of as a $dN/dS = 1$ surface (which would typically be interpreted as evidence for neutral, or nearly neutral, evolution (55, 56)).

The $F = U$ surface is described by the equation $\int \tilde{p}(\gamma)(e^\gamma - 1)d\gamma = 0$, from which it follows that

$$\eta|_{F=U} = \frac{1 - e^{-\gamma_d}}{e^{\gamma_b} - 1}, \quad [10]$$

for the special case of a two-effect DFE. Note that $\eta|_{F=U}$




Fig. 4. Cross sections of $F = U$ and $\alpha = 0$ surfaces. (A) Populations shown in Fig. 2A (which are subject to a two-effect DFE with $NU = 10^4$) are colored by their values of $T_c(F - U)$ measured in simulations. (B) Populations shown in Fig. 2A are colored by measured values of α in simulations. In A and B, solid lines depict predictions of the MSSM approximation for the $F = U$ and $\alpha = 0$ surface, respectively, within its regime of validity; dotted lines denote MSSM predictions beyond its regime of validity. In B, dashed lines denote the corresponding predictions obtained using the N_e -based heuristic described above.

increases monotonically with γ_d , and decreases monotonically with γ_b ; thus, in contrast to $v = 0$ curves, $F = U$ curves do not attain extrema (at least in the space of scaled effects). In contrast to the $v = 0$ surface, γ_b does not tend to 0 as $\gamma_d \rightarrow \infty$, but instead tends to $\log(1 + 1/\eta)$. As a result, to obtain $F = U$ curves in the space of unscaled effects, the “ N_e -based heuristic” described above, in which T_c is given by Eq. (8), does not apply; even at large γ_d , the dynamics are not driven primarily by strong purifying selection on deleterious mutations. Instead, we can use T_c obtained with the MSSM approximation (that is, using Eq. (7)) over a larger range of N_{s_d} values. In Fig. 4, we plot a grid of fixation rates obtained from the same simulations (of populations subject to two-effect DFEs) considered in Fig. 2, along with corresponding $F = U$ theory curves obtained using the MSSM approximation. Our theory curves quite accurately distinguish simulated parameter combinations in which $F > U$ from those in which $F < U$. The same comparison is provided in Fig. S4, including populations with different NU values and populations subject to a two-exponential DFE. We note that, as compared to the $F = U$ surface in the space of scaled effects, the behavior is slightly more complex in this case, and the $F = U$ surface can be somewhat non-monotonic in certain cases (due to variation of T_c with N_{s_b} and N_{s_d}).

Measurements of within-population genetic *diversity* enable another way of characterizing and potentially drawing inferences from a population. In particular, an imbalance in the number of nonsynonymous (and synonymous) polymorphisms observed, relative to the number of nonsynonymous (and synonymous) mutations fixed since divergence of two populations, is often interpreted as a signature of selection—either positive or negative—through the McDonald Kreitman test (57). We can quantify a related imbalance through the statistic α ,

defined as

$$\alpha = 1 - \frac{U_n \pi_{\text{sel}}}{F \pi_{\text{neu}}}, \quad [11]$$

which resembles the McDonald-Kreitman statistic α_{MK} , making an analogy between synonymous mutations and neutral mutations, and between nonsynonymous mutations and selected mutations. Here π_{neu} and π_{sel} denote levels of pairwise heterozygosity of neutral and selected mutations, respectively (with pairwise heterozygosity simply the average number of polymorphisms observed in a sample of two individuals). The statistic α_{MK} has often been used to estimate the fraction of substitutions in a population which are adaptive (58). However, neither α_{MK} nor the statistic α defined here need be positive; $\alpha_{MK} < 0$ is often observed and interpreted as evidence that purifying selection plays a dominant role in the evolution of a population (59). The $\alpha = 0$ surface is thus a third surface on which the dynamics are at least ostensibly neutral in some capacity, and which can also be described using the MSSM approximation. We reproduce results of the MSSM approximation for π_{neu} and π_{sel} in the *SI Appendix*. Using these results, the $\alpha = 0$ surface follows as

$$\int \tilde{\rho}(\gamma) \left(\frac{\gamma e^\gamma + 1 - e^\gamma}{\gamma} \right) d\gamma = 0, \quad [12]$$

a consequence of which is that

$$\eta|_{v=0} < \eta|_{\alpha=0} < \eta|_{F=U}, \quad [13]$$

for any $\tilde{\rho}_b(\gamma)$ and $\tilde{\rho}_d(\gamma)$. Thus, $v > 0$ and $F < U$ on the $\alpha = 0$ surface (and $\alpha < 0$ and $F < U$ on the $v = 0$ surface). Fig. 4 includes a comparison between predictions for the $\alpha = 0$ surface—obtained using Eq. (12), along with Eq. (7) or Eq. (8)—and the values of α obtained in simulations, for populations subject to two-effect DFEs. To obtain predictions for the $\alpha = 0$ surface in the space of unscaled effects, the

“ N_e -based heuristic” is again useful for large N_{s_d} and the MSSM approximation (i.e. Eq. (7)) is useful for small N_{s_d} in relating T_c to N , although the specific patching we have employed in Eq. (9) to obtain the $v = 0$ surface does not quite carry over. We provide a comparison between $\alpha = 0$ predictions and simulated α values for populations subject to a two-exponential DFE in Fig. S5; except for cases when η is small and N_{s_d} is large (and the MSSM approximation breaks down) agreement is qualitatively similar.

Discussion

Previous efforts to treat the dynamics of both beneficial and deleterious mutations have largely done so by treating the two types of mutations in fundamentally different ways (15, 19, 20, 22, 60). These efforts typically start by making one of several assumptions about whether and how often deleterious mutations can fix, and how they impact the fixation probabilities of beneficial mutations. We lack an adequate understanding of the interplay between the two types of mutation in general, or even of which type of mutation will be more important in shaping the fitness trajectory of a population, given a particular distribution of fitness effects. Here, we have used our recently developed MSSM approximation, along with a simple N_e -based heuristic, to characterize this balance between beneficial and deleterious mutations. Our description of the $v = 0$ surface applies under quite general conditions we have described above—essentially, as long as NU_b is sufficiently large that beneficial mutations enter the population sufficiently frequently. As we discuss in more detail below, the $v = 0$ surface is particularly relevant to the fate of a population at long evolutionary times: the $v = 0$ surface limits the ability of evolution to climb fitness landscapes—and thus, under certain conditions, determines the extent to which evolution acts as an optimization process.

We have found that the $v = 0$ constraint is concisely expressed in terms of the distribution of scaled fitness effects $T_c s$ available to a population. Expressed as such, the population size N is relevant only via its impact on the coalescence timescale T_c . Alternatively, given a fixed U and $\rho(s)$, a particular N_0 can be identified such that $v > 0$ for $N > N_0$ and $v < 0$ for $N < N_0$. This has long been recognized in the context of *mutational meltdown* models (61), which typically assume that decreases in fitness imply decreases in a population’s size, and thus further decreases in its fitness (or the opposite increase in its size, if the population instead increases in fitness initially). Our analysis immediately yields a critical *effective* population size N_e at which $v = 0$. For example, for the case of a single beneficial effect and a single deleterious effect, rearranging Eq. (5) yields $N_e|_{v=0} = 2/(s_b + s_d) \log[(U_d s_d)/(U_b s_b)]$. A similar critical effective population size is identified by Whitlock (21). Because the MSSM approximation provides a relation between T_c and N , however, our analysis can also yield a critical *census* population size N_0 at which $v = 0$, for a given $U\rho(s)$.

Previous work has considered the balance between accumulation of beneficial mutations and deleterious mutations under more limiting assumptions (22, 62–64). Notably, Held et. al. (63) analyze the $v = 0$ balance between beneficial and deleterious mutations that emerges under a biophysically grounded model of selection on traits at multiple genes, finding that selection coefficients are tuned to a state of marginal relevance in which $T_c s \sim \mathcal{O}(1)$, under assumptions of negative epistasis

on individual traits and sufficiently frequent mutations. Goyal et. al. (22) compute the fraction $U_b/(U_b + U_d)$ of beneficial mutations to total mutations at which $v = 0$, given a single effect size $s_b = s_d = s$ of both beneficial mutations and deleterious mutations. Goyal et. al. (22) also provide bounding arguments for the case in which the effect sizes of beneficial mutations and deleterious mutations differ, and briefly discuss the case in which distributions of fitness effects are more broad. The key idea of this and other single-effect approaches is that the single effect s that is modeled must be chosen as the most-likely (or in a sense, typical) effect size of a *fixed* mutation (17, 18). Thus, use of a single- s approach to obtain a $v = 0$ constraint requires that the most-likely effect size of a fixed beneficial mutation at least roughly matches the most-likely effect size of a fixed deleterious mutation.

Rice et. al. (64) have found that, under the assumption of no epistasis (i.e., assuming a genome of finite size where mutations at the different loci do not interact epistatically) this is precisely to be expected after long evolutionary timescales. Their basic argument is simple: if beneficial mutations of a particular effect size are more likely to fix than deleterious mutations of that effect size, those mutational opportunities will be depleted faster (or vice versa). This will continue until the distribution of fixed beneficial effects precisely matches the distribution of fixed deleterious effects, at which point adaptation will come to a halt ($v = 0$) and the DFE can be described as “evolutionarily stable”. Therefore, a single- s approach is perhaps appropriate in describing the approach to an evolutionary attractor at long times resulting from a population running out of beneficial mutations. In the presence of epistasis, however, a mutation not only enables a back mutation of the opposite effect, but can also alter the full distribution of fitness effects available to an individual. As a result, the distributions $\rho_b(s)$ and $\rho_d(s)$ can change in independent ways over the course of evolution. Our analysis—which makes no assumption that $\rho_b(s)$ and $\rho_d(s)$ are similar in scale—is thus more applicable to a description of the $v = 0$ constraint in the presence of widespread epistasis. We emphasize that the existence of the $v = 0$ state does not depend on the existence of any particular form of epistasis, and is expected to hold under fairly general conditions. However, the form of the equilibrium point will depend on epistasis, and as we now discuss, our analysis provides a framework to work out the implications of various proposed patterns of *fitness-mediated epistasis* (31), in which $\rho(s)$ varies systematically with the fitness of a population.

To do so requires an additional assumption of how the parameters of a population depend on fitness. For instance, in the presence of *diminishing-returns* epistasis (in which beneficial effects become systematically *weaker* as the fitness of a population increases) populations are constrained to lie on a horizontal line in the space N_{s_d} vs. N_{s_b} depicted in Fig. 2 and Fig. 4. If a population starts out in the $v > 0$ region, its mean fitness will increase and its value of N_{s_b} will subsequently decrease, until the population converges to the $v = 0$ surface. Qualitatively, then, the implications of diminishing-returns epistasis are similar to the implications of a declining fraction of beneficial mutations (which are considered by Goyal et. al. (22)): as evolution proceeds, the population will approach the $v = 0$ surface (either from the $v < 0$ region or the $v > 0$ region); the $v = 0$ surface is thus a stable evolutionary attractor.




Fig. 5. (A) Simulated trajectories of evolving populations subject to both diminishing-returns and increasing-costs epistasis, overlaid on panel C of Fig. 2. Populations are subject to two-effect DFEs with $NU = 10^4$ and $\eta = 0.1$. In all cases, the product $(N_{s_b})(N_{s_d})$ is assumed constant as a population evolves; depending on a population's initial values of N_{s_b} and N_{s_d} it may or may not approach the $v = 0$ surface at long times. The color of a filled circle denotes the time (in generations) a population had a given set of parameters. (B) Mutation accumulation and (C) fitness trajectories of the simulated populations depicted in A. Despite a declining rate of fitness change observed in all cases, rates of (selected) mutation accumulation remain roughly constant.

On the other hand, a pattern of *increasing-costs* epistasis (in which deleterious fitness effects systematically increase with fitness) corresponds to a vertical line in the space N_{s_d} vs. N_{s_b} . The behavior in this case is more complex, depending on the location of the starting point in relation to the $v = 0$ ridgeline: if $N_{s_b} > N_{s_b}^*$, a long term evolutionary attractor does not exist, while if $N_{s_b} < N_{s_b}^*$, two fixed points exist, one of which is stable and one of which is unstable. One could also imagine a pattern of *decreasing-costs* epistasis, in which deleterious fitness effects become systematically *smaller* as a population increases in fitness (e.g. if deleterious mutations are thought of as reversions of beneficial mutations for which a pattern of diminishing-returns epistasis exists). The behavior in this case is similar to that of increasing-costs epistasis, but the stability (and instability) of the two fixed points is swapped. The combination of decreasing-costs epistasis and diminishing-returns epistasis (i.e. negative epistasis), which is present in the biophysically grounded model studied by Held et. al. (63), can also be considered, and results in a stable fixed point.

It is not yet entirely clear which, if any, of these simple patterns best describe the dominant patterns of epistasis in natural populations. More generally, a curve through the parameter space—parameterized by fitness—can be assumed, on which populations are constrained to lie. The intersection(s) of these curves with the $v = 0$ surface determine the long-term evolutionary fixed points, the stability of which can be determined straightforwardly. In principle, these curves may covary along several dimensions (e.g. with quantities such as the overall mutation rate, the relative fraction of beneficial to total mutations, and other quantities such as the shapes of beneficial and deleterious DFEs). In Fig. 5, as a schematic we illustrate the consequences of a relatively simple pattern involving both diminishing-returns epistasis and increasing-costs epistasis; as recently argued by Lyons et. al. (65) and Reddy and Desai (66), these two trends both emerge from a simple null model of pervasive microscopic epistasis. Populations are simulated starting from a range of five initial conditions in the parameter space; in each case, s_b decreases

with the fitness of an individual and s_d increases with fitness of an individual, with the product $s_b s_d$ assumed constant (see *Materials and Methods* for details). Note that for two of the five initial conditions, the population would never approach the $v = 0$ surface (but its rate of adaptation does slow down over time). For the remaining three initial conditions, the population approaches the $v = 0$ surface, either from higher fitness or from lower fitness, and in particular approaches the $\gamma_d < \gamma_d^*$ portion of the $v = 0$ surface described by the MSSM approximation.

In all cases shown in Fig. 5, selected mutation accumulation proceeds throughout at a roughly constant rate, although, consistent with fixation rates measured in simulations and shown in Fig. 4, those populations which approach the $v = 0$ surface have a higher long-term steady state fixation rate of new mutations. Those populations which approach the $v = 0$ surface can be thought to undergo rapid molecular evolution at steady state, in that $F \sim U_b + U_d$ is possible. In contrast, those populations which do not approach the $v = 0$ surface can instead end up with a much smaller fixation rate $F \sim U_b$ (with $F = U_b$ if deleterious mutations are strong enough that all are purged by selection, and beneficial mutations weak enough that they accumulate entirely neutrally). Broadly speaking, these patterns—a rate of fitness increase which declines over time (67), and in particular, the maintenance of a roughly constant rate of mutation accumulation despite a declining rate of fitness increase (27, 68, 69)—have been observed in multiple microbial evolution experiments. Our analysis provides a way to identify regions of the parameter space in which these and similar observed patterns are possible, or alternatively, to yield constraints on the dominant modes of fitness-mediated epistasis given an observed fitness and/or mutation accumulation trajectory.

Using arguments along the lines described above, our analysis can be used to predict the flow of a population through parameter space, given a complete characterization of any form of fitness-mediated epistasis—that is, of $\rho(s|X)$. Crucially, we make the assumption of *slow epistasis*, such that $\rho(s|X)$ can be treated as uniform within the population, and constant in time, in identifying its corresponding rate of fitness increase. This assumption means that we neglect the possibility that individual mutations could lead to specific shifts in $U\rho(s)$, which could then themselves be subject to selection. For example, a lineage could arise that has access to more (or stronger-effect) beneficial (or deleterious) mutations than other individuals within the population, and this lineage could then be subject to second-order selection. Addressing this effect is an interesting topic for other work, and has recently been addressed by (70) using a related theoretical framework.

Materials and Methods

To validate our predictions, we conducted individual-based Wright-Fisher simulations. Simulations were performed using code available at <https://github.com/mjmel/mssm-sim> and used by (24). Simulations consist of a mutation step and a reproduction step repeated each generation. In the mutation step, each individual acquires a Poisson-distributed number of mutations with mean U ; the effect of each mutation is independently drawn from the distribution $\rho(s)$ and increments (or decrements) an individual's log-fitness X . Purely neutral mutations are also introduced at rate U_n . The identities and fitness effects of the mutations carried by each individual are tracked. In the reproduction step, individuals are resampled with replacement with probabilities proportional to their fitnesses e^X .

Populations are initialized clonally and the number of generations which elapses before the first fixation of a mutation is recorded, setting the *epoch length*. Simulations are run for up to 100 epochs, with the mean fitness and heterozygosities—both of neutral mutations and of selected mutations—recorded at each epoch. Simulations which for which fewer than 10 epochs have been reached after a runtime of 24 hours are discarded from our analysis.

For a given value of η and NU , we simulated parameter combinations lying on a grid of N_{s_b} and N_{s_d} values (depicted in Fig. 2, for example). For each value of η and NU , we extracted the coordinates of its ridgeline point, both in the space of scaled effects and in the space of unscaled effects, for comparison with theory. To do so in the space of scaled effects, we take the point with the largest γ_b value such that $v < 0$; the γ_b and γ_d of this point are then recorded as γ_b^* and γ_d^* , respectively. In the space of unscaled effects, we take $(N_{s_b})^*$ as the largest value of N_{s_b} such that $v < 0$ for some parameter combination; $(N_{s_d})^*$ is then taken as the median value of N_{s_d} for those parameter combinations such that $N_{s_b} = (N_{s_b})^*$ and $v < 0$.

We also conducted simulations in which the available fitness effects s_b and s_d depend on fitness. These simulations are identical to those described above, except after the occurrence of each mutation, the effect magnitudes of the next available beneficial and deleterious mutations are updated accordingly. In all cases, we held the product $s_b s_d$ constant with $s_b \propto e^{-X/5}$, where X denotes an individual's log-fitness. A similar functional dependence of s_b on X is found by Wiser et. al. (71) to describe *E. coli* populations in the LTEE experiment. These simulations were run for a total of 5000 generations, with measurements of fitness and the number of fixed mutations recorded every 100 generations.

ACKNOWLEDGMENTS. We thank Benjamin Good and Ivana Cvijovic for many helpful comments and suggestions. MMD acknowledges support from NSF Grant PHY-1914916 and NIH grant R01GM104239. Simulations were conducted on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

1. Wright S (1932) *The roles of mutation, inbreeding, crossbreeding, and selection in evolution.* (na) Vol. 1.
2. Miralles R, Gerrish PJ, Moya A, Elena SF (1999) Clonal interference and the evolution of rna viruses. *Science* 285(5434):1745–1747.
3. De Visser A, Zeyl CW, Gerrish PJ, Blanchard JL, Lenski RE (1999) Diminishing returns from mutation supply rate in asexual populations. *Science* 283(5400):404–406.
4. Desai MM, Fisher DS, Murray AW (2007) The speed of evolution and maintenance of variation in asexual populations. *Current Biology* 17(5):385–394.
5. Kao KC, Sherlock G (2008) Molecular characterization of clonal interference during adaptive evolution in asexual populations of *saccharomyces cerevisiae*. *Nat Genet* 40(12):1499–1504.
6. Lang GI, et al. (2013) Pervasive genetic hitchhiking and clonal interference in forty evolving yeast populations. *Nature* 500(7464):571–574.
7. Nourmohammad A, Otwinowski J, Łuksza M, Mora T, Walczak AM (2019) Fierce selection and interference in b-cell repertoire response to chronic hiv-1. *Molecular Biology and Evolution* 36(10):2184–2194.
8. Strelkowa N, Lässig M (2012) Clonal interference in the evolution of influenza. *Genetics* 192(2):671.
9. Barroso-Batista J, et al. (2014) The first steps of adaptation of *escherichia coli* to the gut are dominated by soft sweeps. *PLOS Genetics* 10(3):1–12.
10. Gerrish PJ, Lenski RE (1998) The fate of competing beneficial mutations in an asexual population. *Genetica* 102(0):127.
11. Muller HJ (1964) The relation of recombination to mutational advance. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis* 1(1):2–9.
12. Felsenstein J (1974) The evolutionary advantage of recombination. *Genetics* 78(2):737.
13. Rouzine IM, Brunet É, Wilke CO (2008) The traveling-wave approach to asexual evolution: Muller's ratchet and speed of adaptation. *Theoretical Population Biology* 73(1):24–46.
14. Neher RA, Shraiman BI (2012) Fluctuations of fitness distributions and the rate of muller's ratchet. *Genetics* 191(4):1283–1293.
15. Bachtrog D, Gordo I (2004) Adaptive evolution of asexual populations under muller's ratchet. *Evolution* 58(7):1403–1413.
16. Desai MM, Fisher DS (2007) Beneficial mutation–selection balance and the effect of linkage on positive selection. *Genetics* 176(3):1759–1798.
17. Fisher DS (2013) Asexual evolution waves: fluctuations and universality. *Journal of Statistical Mechanics: Theory and Experiment* 2013(01):P01011.
18. Good BH, Rouzine IM, Balick DJ, Hallatschek O, Desai MM (2012) Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proceedings of the National Academy of Sciences* 109(13):4950–4955.
19. Johnson T, Barton NH (2002) The effect of deleterious alleles on adaptation in asexual populations. *Genetics* 162(1):395.

20. Good BH, Desai MM (2014) Deleterious passengers in adapting populations. *Genetics* 198(3):1183–1208.
21. Whitlock MC (2000) Fixation of new alleles and the extinction of small populations: Drift load, beneficial alleles, and sexual selection. *Evolution* 54(6):1855–1861.
22. Goyal S, et al. (2012) Dynamic mutation–selection balance as an evolutionary attractor. *Genetics* 191(4):1309–1319.
23. Schiffls S, Szöllösi GJ, Mustonen V, Lässig M (2011) Emergent neutrality in adaptive asexual evolution. *Genetics* 189(4):1361.
24. Melissa MJ, Good BH, Fisher DS, Desai MM (2022) Population genetics of polymorphism and divergence in rapidly evolving populations. *Genetics* 221(4):iyac053.
25. Mustonen V, Lässig M (2009) From fitness landscapes to seascapes: non-equilibrium dynamics of selection and adaptation. *Trends in Genetics* 25(3):111–119.
26. Chou HH, Chiu HC, Delaney NF, Segrè D, Marx CJ (2011) Diminishing returns epistasis among beneficial mutations decelerates adaptation. *Science* 332(6034):1190–1192.
27. Kryazhimskiy S, Rice DP, Jerison ER, Desai MM (2014) Global epistasis makes adaptation predictable despite sequence-level stochasticity. *Science* 344(6191):1519–1522.
28. Bakerlee CW, Phillips AM, Nguyen Ba AN, Desai MM (2021) Dynamics and variability in the pleiotropic effects of adaptation in laboratory budding yeast populations. *eLife* 10:e70918.
29. Johnson MS, Martsul A, Kryazhimskiy S, Desai MM (2019) Higher-fitness yeast genotypes are less robust to deleterious mutations. *Science* 366(6464):490–493.
30. Johnson MS, Desai MM (2022) Mutational robustness changes during long-term adaptation in laboratory budding yeast populations. *eLife* 11:e76491.
31. Good BH, Desai MM (2015) The impact of macroscopic epistasis on long-term evolutionary dynamics. *Genetics* 199(1):177–190.
32. Hegreness M, Shores N, Hartl D, Kishony R (2006) An equivalence principle for the incorporation of favorable mutations in asexual populations. *Science* 311(5767):1615.
33. Smith JM, Haigh J (1974) The hitch-hiking effect of a favourable gene. *Genetical Research* 23(1):237–35.
34. Peck JR (1994) A ruby in the rubbish: beneficial mutations, deleterious mutations and the evolution of sex. *Genetics* 137(2):597.
35. Charlesworth B (1994) The effect of background selection against deleterious mutations on weakly selected, linked variants. *Genetics Research* 63(3):213–227.
36. Hill WG, Robertson A (1966) The effect of linkage on limits to artificial selection. *Genetical Research* 8(3):269–294.
37. Hallatschek O (2011) The noisy edge of traveling waves. *Proceedings of the National Academy of Sciences* 108(5):1783–1787.
38. Neher RA, Shraiman BI, Fisher DS (2010) Rate of adaptation in large sexual populations. *Genetics* 184(2):467–481.
39. Neher RA, Hallatschek O (2013) Genealogies of rapidly adapting populations. *Proceedings of the National Academy of Sciences* 110(2):437–442.
40. Rouzine IM, Wakeley J, Coffin JM (2003) The solitary wave of asexual evolution. *Proceedings of the National Academy of Sciences* 100(2):587–592.
41. Hallatschek O (2018) Selection-like biases emerge in population models with recurrent jackpot events. *Genetics* 210(3):1053–1073.
42. Neher RA (2013) Genetic drift, selective interference, and population genetics of rapid adaptation. *Annual Review of Ecology, Evolution, and Systematics* 44(1):195–215.
43. Crow JF, Kimura M, et al. (1970) An introduction to population genetics theory. *An introduction to population genetics theory*.
44. Messer PW, Petrov DA (2013) Frequent adaptation and the mcdonald–kreitman test. *Proceedings of the National Academy of Sciences* 110(21):8615.
45. Sawyer SA, Hartl DL (1992) Population genetics of polymorphism and divergence. *Genetics* 132(4):1161–1176.
46. Eyre-Walker A, Keightley PD (2007) The distribution of fitness effects of new mutations. *Nature Reviews Genetics* 8:610 EP –.
47. Keightley PD, Eyre-Walker A (2010) What can we learn about the distribution of fitness effects of new mutations from dna sequence data? *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 365(1544):1187–1193.
48. Charlesworth B (2009) Effective population size and patterns of molecular evolution and variation. *Nature Reviews Genetics* 10(3):195–205.
49. Frenkel EM, Good BH, Desai MM (2014) The fates of mutant lineages and the distribution of fitness effects of beneficial mutations in laboratory budding yeast populations. *Genetics* 196(4):1217–1226.
50. Levy SF, et al. (2015) Quantitative evolutionary dynamics using high-resolution lineage tracking. *Nature* 519(7542):181–186.
51. Haigh J (1978) The accumulation of deleterious genes in a population—muller's ratchet. *Theoretical population biology* 14(2):251–267.
52. Charlesworth B, Morgan M, Charlesworth D (1993) The effect of deleterious mutations on neutral molecular variation. *Genetics* 134(4):1289–1303.
53. Söderberg RJ, Berg OG (2007) Mutational interference and the progression of muller's ratchet when mutations have a broad range of deleterious effects. *Genetics* 177(2):971–986.
54. Good BH, Walczak AM, Neher RA, Desai MM (2014) Genetic diversity in the interference selection limit. *PLOS Genetics* 10(3):1–1.
55. Yang Z, Bielawski JP (2000) Statistical methods for detecting molecular adaptation. *Trends in Ecology & Evolution* 15(12):496–503.
56. Kryazhimskiy S, Plotkin JB (2008) The population genetics of dn/ds. *PLOS Genetics* 4(12):e1000304–.
57. McDonald JH, Kreitman M (1991) Adaptive protein evolution at the adh locus in drosophila. *Nature* 351(6328):652–654.
58. Eyre-Walker A (2006) The genomic rate of adaptive evolution. *Trends in Ecology & Evolution* 21(10):569–575.
59. Charlesworth J, Eyre-Walker A (2008) The mcdonald–kreitman test and slightly deleterious mutations. *Molecular biology and evolution* 25(6):1007–1015.
60. McFarland CD, Korolev KS, Kryukov GV, Sunyaev SR, Mirny LA (2013) Impact of deleterious passenger mutations on cancer progression. *Proceedings of the National Academy of Sciences*

- 110(8):2910–2915.
61. Lynch M, Bürger R, Butcher D, Gabriel W (1993) The mutational meltdown in asexual populations. *Journal of Heredity* 84(5):339–344.
62. McFarland CD, Mirny LA, Korolev KS (2014) Tug-of-war between driver and passenger mutations in cancer and other adaptive processes. *Proceedings of the National Academy of Sciences* 111(42):15138–15143.
63. Held T, Klemmer D, Lässig M (2019) Survival of the simplest in microbial evolution. *Nature Communications* 10(1):2472.
64. Rice DP, Good BH, Desai MM (2015) The evolutionarily stable distribution of fitness effects. *Genetics* 200(1):321–329.
65. Lyons DM, Zou Z, Xu H, Zhang J (2020) Idiosyncratic epistasis creates universals in mutational effects and evolutionary trajectories. *Nature Ecology & Evolution*.
66. Reddy G, Desai MM (2021) Global epistasis emerges from a generic model of a complex trait. *eLife* 10:e64740.
67. Couce A, Tenaillon OA (2015) The rule of declining adaptability in microbial evolution experiments. *Frontiers in Genetics* 6:99.
68. Barrick JE, et al. (2009) Genome evolution and adaptation in a long-term experiment with *Escherichia coli*. *Nature* 461(7268):1243–1247.
69. Good BH, McDonald MJ, Barrick JE, Lenski RE, Desai MM (2017) The dynamics of molecular evolution over 60,000 generations. *Nature* 551:45 EP –.
70. Ferrare JT, Good BH (2023) Evolution of evolvability in rapidly adapting populations. *bioRxiv*.
71. Wiser MJ, Ribeck N, Lenski RE (2013) Long-term dynamics of adaptation in asexual populations. *Science* 342(6164):1364–1367.

DRAFT