Moderate amounts of epistasis are not evolutionarily stable in small populations

Dariya K. Sydykova[1,2], Thomas LaBar[2,3], Christoph Adami[2,4,5], and Claus O. Wilke[1,2,*]

[1]Department of Integrative Biology, The University of Texas at Austin
[2]BEACON Center for the Study of Evolution Action, Michigan State University
[3]Department of Molecular and Cell Biology, Harvard University
[4]Department of Microbiology and Molecular Genetics, Michigan State University
[5]Department of Physics and Astronomy, Arizona State University
[*]wilke@austin.utexas.edu

August 30, 2019

Abstract

High mutation rates select for the evolution of mutational robustness where populations inhabit flat fitness peaks with little epistasis, protecting them from a mutational meltdown. Recent evidence suggests that a different effect protects small populations from extinction via the accumulation of deleterious mutations. In drift robustness, populations tend to occupy peaks with steep flanks and positive epistasis between mutations. However, it is not known what happens when mutation rates are high and population sizes are small at the same time. Using a simple fitness model with both variable epistasis and variable mutational effect size, we show that the equilibrium fitness has a minimum as a function of the parameter that tunes epistasis, implying that this critical point is an unstable fixed point for evolutionary trajectories. In agent-based simulations of evolution at finite mutation rate, we demonstrate that when mutations can change epistasis, trajectories with a subcritical value of epistasis evolve to decrease epistasis, while those with supercritical initial points evolve towards higher epistasis. These two fixed points can be identified with mutational and drift robustness, respectively.

Introduction

When a population is in evolutionary balance, it is able to maintain its fitness while at the same time exploring ways to increase its fit to the environment via adaptive mutations [1].

However, this balance between the evolutionary forces of selection and mutation can sometimes be precarious. When mutation rates become too high, for example, mutations can overpower selection leading to the extinction of populations via mutational meltdown [2]. Similarly, when population size dwindles, selection can become so weak that deleterious mutations cannot be eliminated anymore, leading to extinction via Muller's ratchet [3]. It turns out that populations can adapt to these threats, for example by evolving mutational robustness [4], or by evolving resistance to the effects of drift [5, 6, 7]. Mutational robustness is achieved by a genetic architecture that implies a flat fitness peak, so that a large fraction of mutations are either neutral or have a small fitness effect [8, 9], while giving up wild-type fitness in return. Robustness to drift, on the other hand, appears to involve favoring fitness peaks that have steep flanks, enabled by mutations that are synergistic in their deleterious effect [5, 6], while *reducing* (rather than increasing) the likelihood of mutations with small effect, and increasing the fraction of mutations that are lethal. Interestingly, a recent re-analysis of the "survival-of-the-flattest" effect (that enables mutational robustness) has shown that an increase in the fraction of lethal mutations is also seen in the response to high mutation rates (see also [7]), suggesting that resistance to drift and to mutational meltdown are intertwined.

Here we study the impact of epistasis on both drift and mutational robustness in a simple model fitness landscape and show that whether a population predominantly displays drift or mutational robustness is largely determined by the average value of directional epistasis: populations occupying a peak with synergistic epistasis above a critical value will tend to evolve towards drift-robust peaks (by moving towards peaks with increased positive epistasis), while those inhabiting peaks with sub-critical epistasis will respond by lowering epistasis until mutations are mostly neutral, consistent with mutational robustness. Thus, evolutionary trajectories for populations under evolutionary stress will bifurcate towards drift-robust or mutationally-robust fixed points.

Model

We study a simple fitness landscape in which peaks have unit wild-type fitness, and the fitness of a k-mutant is determined by

$$f(k) = e^{-sk^q} \ , \tag{1}$$

where $s = -\log f(1)$ is the mean effect of a deleterious mutation, q determines the degree of directional epistasis, and genotypes have a finite number of binary loci L (see, e.g., [10])[1]. In such a model $q = 1$ signals absence of epistasis (a multiplicative landscape), $q > 1$ describes a peak with synergy between deleterious mutations, while $q < 1$ is indicative of buffering mutations (antagonistic epistasis). When $q > 1$ we sometimes speak of *negative* epistasis (because the combined two-mutant fitness is lower than the multiplicative expectation), while $q < 1$ indicates *positive* epistasis (the double-mutant is higher in fitness than expected on the basis of the single-mutation effect).

[1]The present model in which fitness declines as a function of genetic distance from the wild-type (modulated by epistasis) gives rise to conclusions similar to what Fisher's geometric model would predict, even though in Fisher's model the distance from wild-type is phenotypic rather than genetic [11].

The equilibrium fitness of a population in mutation-selection balance can be calculated in the weak mutation limit $N\mu \ll 1$, where k-mutants either go to fixation or are lost one-by-one, by averaging over the stationary distribution of k-mutants

$$f_{\rm eq} = \sum_{k=0}^{L} p_k f(k) \ . \tag{2}$$

The stationary distribution p_k is easily calculated using the fixation formula of Sella and Hirsh [12] (which is a more accurate estimate than Kimura's celebrated formula) as

$$p_k = \binom{L}{k} f(k)^{2N-2}/Z \ . \tag{3}$$

In the expression above, N is the population size, and Z is the partition function

$$Z = \sum_{k=0}^{L} \binom{L}{k} f(k)^{2N-2} \ , \tag{4}$$

and we used the Sella-Hirsh fixation probability appropriate for a haploid Wright-Fisher process. For $q = 1$ we can obtain a closed-form expression for the equilibrium fitness,

$$f_{\rm eq} = \left(1 - \frac{1-e^{-s}}{1+e^{(2N-2)s}}\right)^L \approx \left(1 - \frac{se^{-Ns}}{2\cosh(Ns)}\right)^L \ , \tag{5}$$

that shows clearly the steep fitness drop with decreasing population size that is due to genetic drift. But different values for q affect the fitness drop differently. In Fig. 1A, we can see the dependence of $f_{\rm eq}$ on the population size for the multiplicative model ($q = 1$), a model with positive epistasis ($q = 2.0$), as well as the case of negative epistasis ($q = 0.5$), evaluated at $L = 100$. The model suggests that while positive epistasis protects from a fitness drop for moderate population sizes (higher mean equilibrium fitness), the drop becomes severe once populations dwindle below 100. In fact, plotting $f_{\rm eq}$ against q as in Fig. 1B reveals a fitness *minimum* as a function of q, suggesting that fitness loss via drift can be prevented in two different ways: high positive epistasis or high negative epistasis, while populations with weak or no epistasis appear to be the most vulnerable.

Two regimes: selection and neutral drift

The minimum in mean equilibrium fitness apparent in Fig. 1B can be seen as interpolating between two regimes: the *neutral drift* regime defined here by the limit of small $q \ll 1$, and the *selection* regime $q \gg 1$. In the selection regime, an organism's fitness declines rapidly with increasing number of mutations, and this rapid decline effectively limits the maximum number of mutations an organism can carry. By contrast, in the neutral drift regime even organisms with a very large number of mutations have a fitness close to the maximum possible value of 1.

When $q \ll 1$, the fitness landscape is essentially flat, and dynamics are dominated by neutral drift on that landscape. We can estimate the mean equilibrium fitness in the neutral regime by using

$$p_k = \binom{L}{k} / Z \ , \tag{6}$$

that is, the distribution Eq. (3) but with $f_k \equiv 1$. The result is the dashed line in Fig. 2, which agrees with the full model for $q < 1$. Note that in this derivation we set $f_k \equiv 1$ only in Eq. (3) but not in Eq. (2). The idea is that in the drift regime fitness differences are sufficiently small that they have no influence on the mutant distribution p_k. This does not mean, however, that all organisms have a fitness of 1.

On the other hand, when epistasis between deleterious mutations is synergistic ($q \gg 1$), the number of mutations that a population can sustain before it goes extinct is limited to some number $k_{\max}$. We can model this limitation by imposing a maximum number of mutations $k_{\max}$,

$$f_{\text{eq}}^{\text{sel}} = \sum_{k=0}^{k_{\max}} f_k p_k, \tag{7}$$

where $k_{\max} < L$. This truncated mutation model agrees with the full solution for large q (Fig. 2, dotted line).

One of the most striking features of the interplay between the neutral regime and the selection regime is the appearance of a minimum mean fitness (as a function of epistasis) where the drop of fitness is largest. The location of this minimum $q^\star$ (reflecting the amount of directional epistasis that leads to the largest fitness loss) depends on the population size, the mean deleterious effect of mutations, as well as the number of loci (Fig. 3).

To estimate the epistasis coefficient at which the steady-state fitness is at its minimum, we analyze the stationary distribution of fitness (3), which is determined by the combined effects of neutral drift (represented by the term $\binom{L}{k}$) and selection, via the fitness $f_k = e^{-sk^q(2N-1)}$. The binomial coefficient of neutral drift for any k is typically much larger than 1, whereas the selection term is much smaller than 1. To estimate the amount of epistasis where these two components balance each other we maximize neutral drift by setting $k = L/2$, so that the condition for maximal fitness loss (where drift maximally balances selection) becomes

$$1 = \binom{L}{L/2} e^{-s(L/2)^q(2N-1)}. \tag{8}$$

We can solve this equation for q by using the Stirling approximation to expand the binomial coefficient, and obtain for the minimum q^* that

$$q^* \approx \frac{\log\left(L \log 2/2sN\right)}{\log\left(N/2\right)}. \tag{9}$$

We can test this estimate by comparing it to the numerically inferred minimum obtained via numerical simulation, described below. Figure 3 shows that our estimate of q^* is very close to the numerically determined minimum for various population sizes and selection strengths.

Increased mutation rate exacerbates fitness loss in neutral regime

The theoretical results shown above were derived in the weak mutation limit where every mutation is either lost or goes to fixation before another mutation occurs in the population. In this section we study how finite mutation rates modify those results.

We simulate finite populations on a single-peak fitness landscape at finite mutation rates μ using stochastic simulation. The population evolves asexually, and the population size is held constant over time for all simulations. For each combination of mutation rates and epistasis parameters, we simulated populations sizes $N = 10$ and $N = 100$, as well as selection coefficients $s = 0.01$ and $s = 0.001$. We recorded the mean fitness of the population over a period of time after a population reached steady-state (see Methods), as a proxy for this equilibrium fitness f_{eq}. The simulations of the evolutionary process on the fitness landscape defined by Eq. (1) recover the theoretical results well for small mutation rates, as expected. As the mutation rate increases, we see notable departures from the weak mutation limit for the selection regime (larger q), while the neutral (drift) regime is largely unaffected by the increased rates (Fig. 4).

In particular, we notice that the minimum of the equilibrium fitness shifts towards higher q (Fig. 4). Furthermore, while for small mutation rates an increased epistasis protects from the loss of fitness due to genetic drift (mean fitness does not drop appreciably), it is clear that higher mutation rates negate this effect, and instead exacerbate the loss of fitness. Indeed, the increased mutation rate mimics the effect of a smaller population size (see Fig. 1B), which is expected as the effective population size decreases with mutation rate.

While the depressed equilibrium fitness suggests that there are two routes to withstand genetic drift at small population sizes, it is not clear whether evolutionary trajectories could indeed bifurcate.

Bifurcation analysis of survival strategies

The minimum in f_{eq} at $q^\star$ suggests that if q were a dynamical variable, then $q^\star$ represents an unstable fixed point of the evolutionary dynamics. While q is not a dynamical variable in the usual sense, we can simulate it by endowing each genotype with a particular value of q that can be changed via mutation. In such a simulation, the statistics of the mutational process affecting q (the rate of change μ_q as well as the mean change per mutation Δq) matter, so we test multiple different values for each.

It is worth pointing out that a genotype-dependent q appears to contradict the idea of fitness optimization in a landscape with a fixed fitness function such as Eq. (1). Such a function suggests that as a population climbs this peak, the parameters q and s are unaffected by this climb. While this is true for such a simple fitness function, it does not hold for more realistic evolutionary landscapes (for example in digital life [13, 14, 10, 15]), where the mean effect of mutations s and the directional epistasis q are not fixed properties of the landscape, but instead emerge as properties of the local neighborhood in genetic space. As a consequence, moving in this space (via mutations) will affect both s and q. We attempt to simulate part of that dynamics by allowing q to adapt (while keeping s fixed). If selection favors a particular value of epistasis, we should see a gradual change in the mean epistasis $\bar{q}$ of a population.

In Fig. 5, we show how the mean epistasis parameter $\bar{q}$ (averaged over sequences in the population) changes over time when populations are seeded with different seed organisms with fixed initial q. A bifurcation is indicated when trajectories move to different future fixed points given different initial states. While we can see clear signs of a bifurcation when plotting the mean trajectory in q-space over time (Fig. 5), viewing each trajectory separately reveals significant variation among them. In particular, for trajectories that are initialized with a q above the fixed point, some trajectories still move towards the low-q fixed point, which results in the mean of trajectories to appear constant (for example, in Fig. 5C). We discuss this phenomenon below.

Discussion

The dynamics of evolution in asexual population is well-understood in the common population-genetic limits, namely vanishingly small mutation rate and large population (weak mutation and strong selection). When mutation rates are high and selection is weak, the classic theoretical results are undermined by new effects such as mutational robustness (effect of large mutation rate) and drift robustness (effect of small population size), as anticipated by generalized population-genetic models such as "free-fitness" evolution [16, 17, 12, 18].

In most realistic populations, we expect both effects to contribute. For example, when the mutation rate is large, the effective population size is diminished, so that both mutational and drift robustness are bound to be intertwined. The mean directional epistasis between mutations plays a role in both effects. While fitness peaks for mutationally robust populations tend to be flatter with little epistasis between mutations, we also observe something akin to truncation selection [8, 9]. In drift robustness, we observe both an increase in neutral mutations as well as an increase in strongly deleterious and lethal mutations, mediated by strong negative epistasis ($q > 1$).

Here we calculated the mean equilibrium fitness of a population in the limit of small mutation rates using a simple fitness function with variable epistasis and tunable mutation effect-size, and found a minimum as a function of the mean directional epistasis parameter q that depends on population size. Stochastic simulations of adaptation on this landscape suggest that the minimum also depends on mutation rate. The model further suggests that there are two attractive fixed points for evolutionary dynamics, namely small q where mutations become nearly neutral, and large q where deleterious mutations interact synergistically. The low-q fixed point[2] can be identified with mutational robustness ($q \approx 0$). In contrast, the $q > 1$ fixed point is reminiscent of drift robustness. While the existence of a minimum in equilibrium fitness is suggestive of an unstable fixed point $q^\star$ at which evolutionary trajectories bifurcate towards a low-q and a high-q fixed point, an agent-based simulation of such trajectories in a bit-string fitness model implementing Eq. (1) but with variable q paints a more complicated picture.

It is clear from inspection of Eq. (1) that for every single sequence, a reduction of q while keeping k constant increases fitness (as $\partial f/\partial q < 0$), no matter what the mean q of the population. This means that the population will sense an evolutionary pressure to

[2]Note that while technically the low-q fixed point is $q = 0$, this value cannot be attained in any realistic population as such a landscape is completely neutral ($w = 1$) in this limit.

reduce q independently of the mean population epistasis. However, if $q > q^\star$, a secondary selective pressure appears that acts via the fitness distribution of a sequence's offspring. For sequences with $q > q^\star$, sequences with higher q have on average off-spring with higher fitness than those with lower q, leading to a second-order selective pressure to increase q. However, in any particular fitness trajectory, there is a chance that a sequence with $q < q^\star$ is among the offspring. Such a sequence may then go to fixation and abrogate the evolutionary trajectories leading towards a $q > q^\star$, even though the selective pressure towards higher q is still present. We also expect that the likelihood of mutations that create sequences with $q < q^\star$ in the off-spring distribution depends on μ_q as well as Δq. This is precisely what we observe in Figs. 5 and 6: while for small $q < q^\star$ the trend towards the mutationally robust fixed point is evident, at $q > q^\star$ the mean epistasis across 10 replicate experiments often shows a decrease (or remains constant) even though theoretically we expect an approach towards the drift-robust fixed point. The distribution of fitness trajectories shown in Fig. 6 shows that while some trajectories indeed move towards higher q, the possibility of mutating towards $q < q^\star$ leads to trajectories in which the secondary selective pressure towards higher q is muted. Indeed, trajectories towards $q > q^\star$ are absent among the replicates with initial $q < q^\star$, reinforcing the conclusion that a critical amount of epistasis separates a population's response to evolutionary stress either in a mutationally-robust, or a drift-robust manner.

While it is difficult to extrapolate results obtained using the abstract fitness function (1) to more complex landscapes in which many different peaks with different effect sizes and directional epistasis exist at the same time, our results support the notion that mutational robustness and drift robustness are indeed two different effects, which are likely to be intertwined in realistic scenarios.

Methods

Evolutionary model with fixed epistasis

For all simulations, we implemented an individual-based bit string model. Each individual was represented with a number of deleterious mutations it possesses. The number of deleterious mutations is k and $k = 0, 1, ..., L$. L is the maximum number of deleterious mutations allowed in a population. A population was represented as a vector V of a length $L+1$. Each bin V_k within a vector corresponded to the number of individuals with deleterious mutations k. The fitness of an individual could be determined with equation $f_k = e^{-sk^q}$, where s is the selection coefficient, and q is the epistasis coefficient. At each time point, a population replicates and mutates. A replication event increases or decreases the number of individuals within each bin. The number of offspring was drawn from a multivariate distribution, and the probability of replication for each bin was determined by the number of individuals within a bin and a bin's fitness. For all replication events, the population size was held constant. After replication, a mutation event moves individuals up or down a bin. For most simulations, the maximum number of moves up or down a bin was set to 3. When mutation rate was set to 1, the maximum number of moves was set to 4. The probability of mutating was calculated with

$$\sum_{i=\max(0,k-j)}^{\max(L-j,k)} \binom{k}{i}\binom{L-k}{i-k+j} u^{(j-k+2i)} u^{(L+k-j-2i)}. \tag{10}$$

Here, u is per site mutation rate, k is number of mutation an individual has at time t, and j is the number of mutations an individual has at time $t+1$. Using the probabilities calculated from the equation above, the number of individuals that would move or stay were drawn from a multivariate distribution. We simulated population sizes of 100 and 10. We set selection coefficients to 0.01 and 0.001. We use the mutation rates of 1, 0.1, 0.01, 0.001, and 0.0001 (mutation rate is defined as the expected number of mutations per genome per duplication). For each combination of population size, selection coefficient, and mutation rate, we simulated 10 replicates. After a population has reached an equilibrium at $t = 1,500,000$, we calculated equilibrium fitness by taking the mean of population fitness over the next 1,000,000 time steps.

Evolutionary model with evolving epistasis

Similarly to simulations with fixed epistasis, we implemented an individual-based bit-string model to simulate populations with evolving epistasis. A population was represented with two vectors. The length of each of the vectors corresponded to the size of the population, and each element within a vector represented an individual. The first vector contained a number of mutations (k) an individual possesses, and the second vector contained an epistasis coefficient (q) an individual's genome experiences. The fitness of an individual could be determined with equation $f_k = e^{-sk^q}$. At each time point, a population replicated and mutated. We replicated a population by creating a new generation with Wright-Fisher model. We mutated a population with a two step process. First, an individual either gained a mutation ($k+1$), lost a mutation ($k-1$), or did not mutate. The probability of mutating was calculated with equation 10. Second, each individual's epistatis mutated. If epistasis was set to change, it could equally likely increase or decrease by a fixed amount (Δq). We set Δq to remain the same across generations per one trajectory of an evolving population. However, epistasis was only set to evolve after a population has reached an equilibrium (after $t = 200,000$).

Acknowledgements This work was supported in part by the National Science Foundation's BEACON Center for the Study of Evolution in Action, under contract No. DBI-0939454.

References

[1] S. Goyal, D. J. Balick, E. R. Jerison, R. A. Neher, B. I. Shraiman, and M. M. Desai. Dynamic mutation–selection balance as an evolutionary attractor. *Genetics*, 191:1309–1319, 2012.

[2] M. Lynch, R. Bürger, D. Butcher, and W. Gabriel. The mutational meltdown in asexual populations. *J Hered*, 84:339–44, 1993.

[3] J. Haigh. The accumulation of deleterious genes in a population – Muller's ratchet. *Theor. Popul. Biol.*, 14:251—267, 1978.

[4] C.O. Wilke and C. Adami. Evolution of mutational robustness. *Mutation Research*, 522:3–11, 2003.

[5] A. S. Kondrashov. Muller's ratchet under epistatic selection. *Genetics*, 136:1469–73, 1994.

[6] T. LaBar and C. Adami. Evolution of drift robustness in small populations of digital organisms. *Nature Communications*, 8:1012, 2017.

[7] Y. Lan, A. Trout, D. M. Weinreich, and C. S. Wylie. Natural selection can favor the evolution of ratchet robustness over evolution of mutational robustness. BioRxiv 122087, 2017.

[8] C. O. Wilke, J. L. Wang, C. Ofria, R. E. Lenski, and C. Adami. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, 412:331–333, 2001.

[9] J. Franklin, T. LaBar, and C. Adami. Mapping the peaks: Fitness landscapes of the fittest and the flattest. *Artificial Life*, 25:250–262, 2019.

[10] C. O. Wilke and C. Adami. Interaction between directional epistasis and average mutational effects. *Proc. R. Soc. Lond. B*, 268:1469–1474, 2001.

[11] O. Tenaillon, O. K. Silander, J.-P. Uzan, and L. Chao. Quantifying organismal complexity using a population genetic approach. *PLoS One*, 2:e217, 2007.

[12] G. Sella and A. E. Hirsh. The application of statistical physics to evolutionary biology. *Proc. Natl. Acad. Sci. U.S.A.*, 102:9541–9546, 2005.

[13] C. Adami. *Introduction to Artificial Life*. TELOS, Springer-Verlag, New York, 1998.

[14] C. Adami, C. Ofria, and T. C. Collier. Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, 97:4463–4468, 2000.

[15] C. Ofria, D. M. Bryson, and C. O. Wilke. Avida: A software platform for research in computational evolutionary biology. In M. Komosinski and A. Adamatzky, editors, *Artificial Life Models in Software*, pages 3–35. Springer London, 2009.

[16] Y. Iwasa. Free fitness that always increases in evolution. *J Theor Biol*, 135:265–81, 1988.

[17] T. Aita, S. Morinaga, and Y. Husimi. Thermodynamical interpretation of evolutionary dynamics on a fitness landscape in a evolution reactor, i. *Bull. Math. Biol.*, 66:1371–1403, 2004.

[18] N. H. Barton and J. B. Coe. On the application of statistical physics to evolutionary biology. *J Theor Biol*, 259:317–24, 2009.

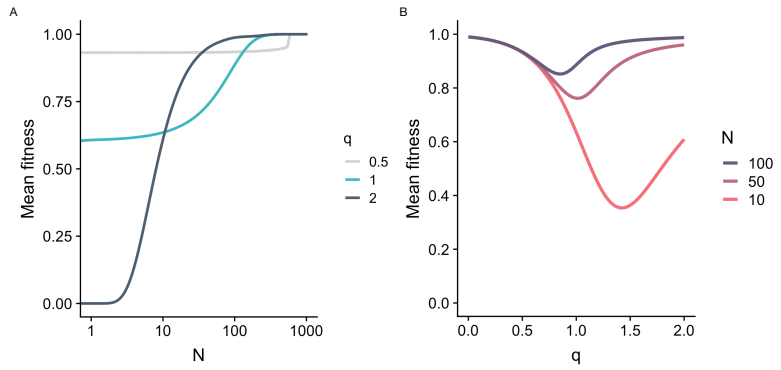

Figure 1: A: Equilibrium fitness Eq. (2) as a function of population size for strong positive epistasis ($q = 2.0$, dark grey), no epistasis ($q = 1.0$, teal), and strong negative epistasis ($q = 0.5$, light grey)B: Equilibrium fitness as a function of epistasis q, for three different population sizes (red: $N = 10$, dark red: $N = 50$, black: $N = 100$.

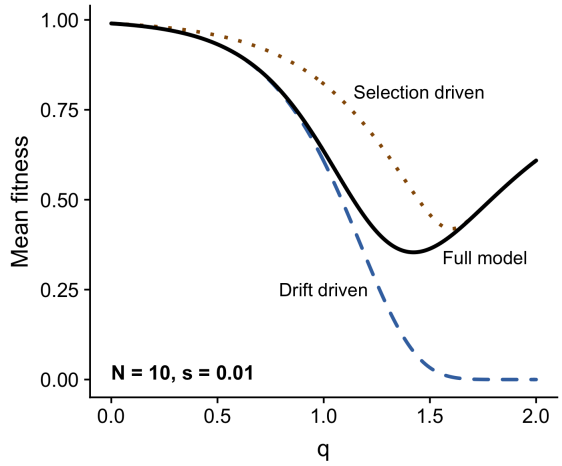

Figure 2: Equilibrium fitness Eq. (2) as a function of population epistasis for $N = 10$ and $s = 0.01$. Full model (solid black line), drift-driven approximation using Eq. (6) (dashed blue line), and selection-driven approximation using Eq. (7) and a maximum number of mutations $k_{\max} = 20$ (dotted brown line).

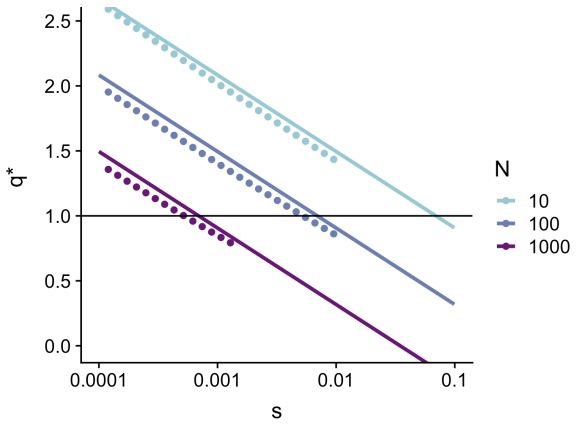

Figure 3: Relationship between critical epistasis q^* and selection (s) for different population sizes (N). The lines represent analytically derived q^*, as given by Eq. (9). The dots represent q^* values obtained by numerically solving Eq. (8).

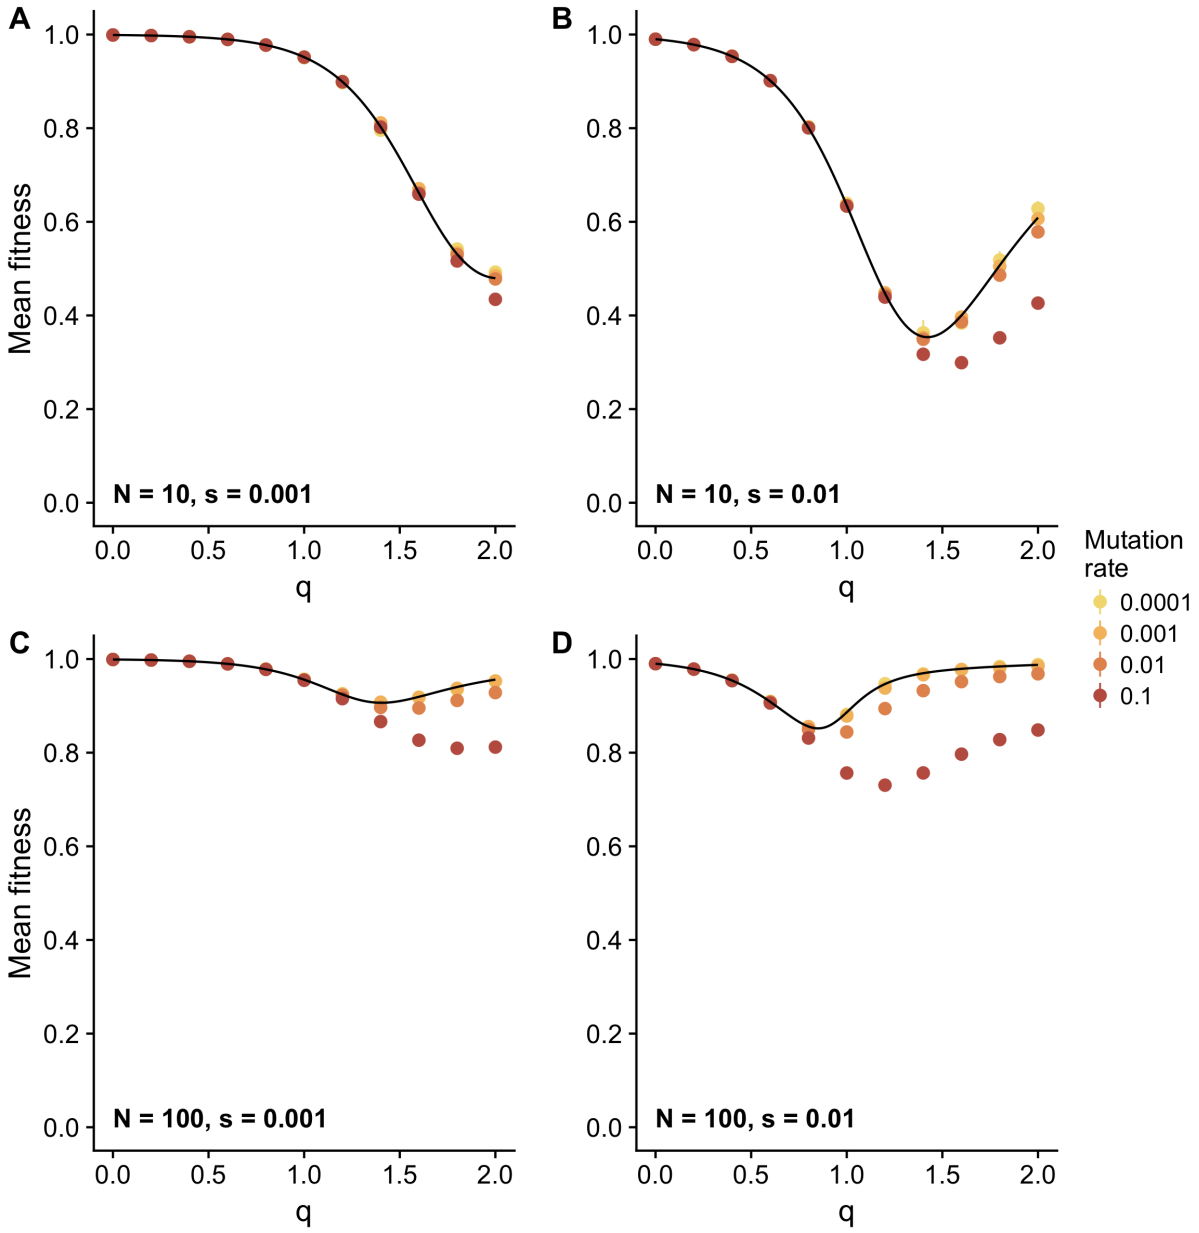

Figure 4: Comparison between theoretical equilibrium fitness and simulated equilibrium fitness. The black line represents theoretical fitness (equation 2) of the population at steady state according to Sella and Hirsh [12]. The dots represent mean fitness of a population at steady state over 10 simulations for different mutation rates (see legend). The error bars represent the standard error. For almost all cases, error bars are smaller than the symbol size. (A-D) Theoretical fitness and simulated fitness for different population sizes (N), epistasis coefficient (q), and selection coefficient (s).

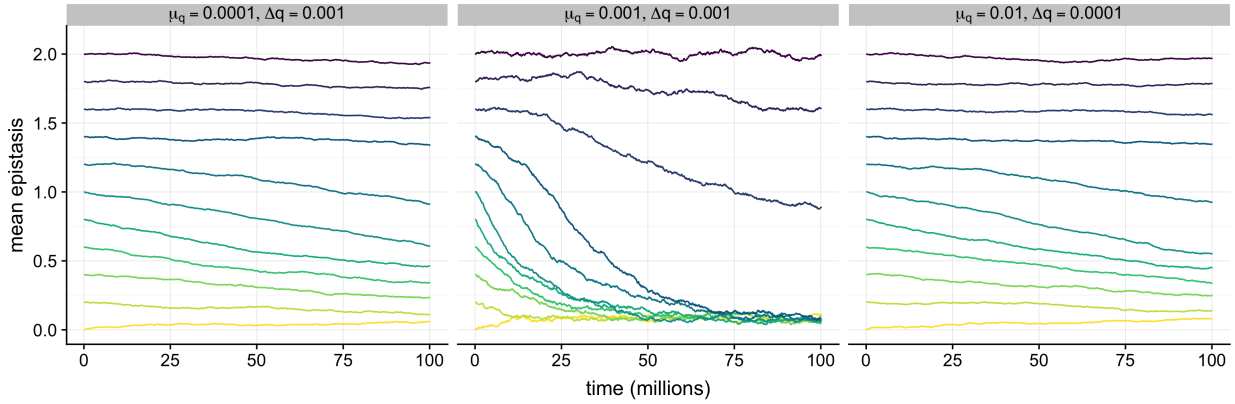

Figure 5: Mean epistasis q as a function of time t for different combinations of q and Δq. Each line corresponds to the mean over 10 replicates. The numbers at the top correspond to the mutation rate of q, and the numbers on the right correspond to Δq. Here, $\mu = 0.01$, $N = 100$, and $s = 0.01$. The different colors correspond to different starting points of q. Each population had a fixed q until it reached an equilibrium, and then q was allowed to evolve after $t = 200,000$. Each line corresponds to one population. Here, $\mu = 0.01$, $N = 100$, and $s = 0.01$. The different colors correspond to different starting points of q.

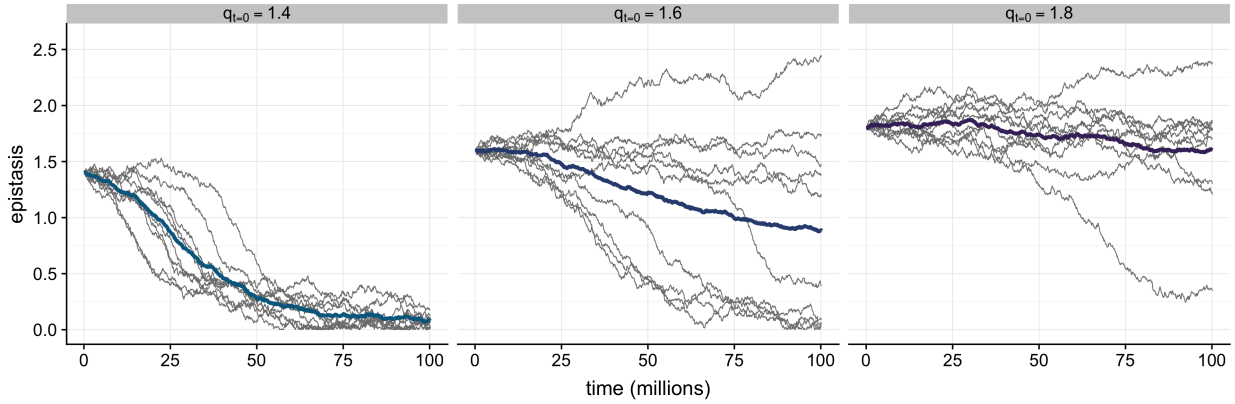

Figure 6: Single population trajectories for evolving epistasis. This figure shows populations in the middle panel in Fig. 5 ($\mu_q = 0.001, \Delta q = 0.001$). A gray line corresponds to the evolution of one population. The blue lines trace the mean among 10 replicate populations. The numbers at the top correspond to the initial population epistasis ($q(t = 0)$) in that panel.