A model for the interplay between plastic tradeoffs and evolution in changing environments

Mikhail Tikhonov[a,b,1], Shamit Kachru[c,d], and Daniel S. Fisher[d,e]

[a]Department of Physics, Washington University in St. Louis, St. Louis, MO 63130; [b]Center for Science and Engineering of Living Systems, Washington University in St. Louis, St. Louis, MO 63130; [c]Department of Physics, Stanford University, Stanford, CA 94305; [d]BioX, Stanford University, Stanford, CA 94305; and [e]Department of Applied Physics, Stanford University, Stanford, CA 94305

Edited by Eugene V. Koonin, National Institutes of Health, Bethesda, MD, and approved March 9, 2020 (received for review September 6, 2019)

Performance tradeoffs are ubiquitous in both ecological and evolutionary modeling, yet they are usually postulated and built into fitness and ecological landscapes. However, tradeoffs depend on genetic background and evolutionary history and can themselves evolve. We present a simple model capable of capturing the key feedback loop: evolutionary history shapes tradeoff strength, which, in turn, shapes evolutionary future. One consequence of this feedback is that genomes with identical fitness can have different evolutionary properties shaped by prior environmental exposure. Another is that, generically, the best adaptations to one environment may evolve in another. Our simple framework bridges the gap between the phenotypic Fisher's Geometric Model and the genotypic properties, such as modularity and evolvability, and can serve as a rich playground for investigating evolution in multiple or changing environments.

evolution | changing environment | performance tradeoff

Performance tradeoffs, caricatured by "you can't be good at everything," are ubiquitous in both ecology and evolution. Sometimes modeled explicitly (e.g., as a fixed total budget of energy or proteome) and often implicitly (e.g., as fitness costs associated with traits), tradeoffs are a staple of ecological and evolutionary modeling (1–6).

The rigid tradeoffs assumed in many models implicitly derive from the assumption that long-acting evolutionary pressures drive organisms to approximate Pareto optimality (7–9) (i.e., to a regime where performance at the relevant tasks cannot all be improved simultaneously). However, even under this assumption (10), the Pareto front should be high-dimensional (11, 12) so that the tradeoff between any subset of traits is not, in fact, rigid. Furthermore, tradeoffs will depend on evolutionary history and themselves evolve (13–18). In laboratory experiments, adapting bacteria to one task can both hinder and improve their performance at another depending on the experimental protocol, the studied strain, and the exact nature of the tasks as well as history of prior exposure (19–21). Some phenomena seem nonintuitive and surprising; for instance, even very weak levels of an antibiotic can induce resistance to much higher levels (22, 23).

In some cases, the specific mechanisms responsible for tradeoffs and their plasticity are known and can be modeled mechanistically (24, 25). This approach can be very informative about a particular case of interest but leaves aside important broader questions. Which of the observed behaviors depend on the details of a given biological system, and which are more general? Which experimental observations should be considered surprising, and which can be captured already by the simplest models? What other behaviors might be expected? For instance, can less frequent exposure to environment X result in better adaptation to it? Can exposure to X result in better fitness in environment Y than exposure to Y itself? More specifically, would a prior exposure to a milder version of a stress facilitate adaptation to its stronger form?

Elucidating which behaviors are surprising vs. general is a key role of theory and simple models. Here, we propose a minimally structured model capturing some key experimentally observed behaviors: namely, the model exhibits performance tradeoffs, but their strength evolves and depends on evolutionary history. This minimal setting proves sufficient to observe nontrivial ways in which tradeoff strength shaped by evolutionary past can predictably influence evolutionary future; in particular, we identify a mechanism whereby achieving highest fitness in one set of environments proceeds through exposure to a different set. Our basic framework can serve as a rich null model for evolution in multiple or changing environments.

Toolbox Model

A widely used framework shaping much of the theoretical intuition about evolutionary adaptation is Fisher's Geometric Model (FGM) (26, 27). Separately, another influential branch of thinking about multienvironment evolution, particularly relevant to the subject of tradeoffs, concerns the modularity of biological function and its relation to the notion of "evolvability" (28, 29). So far, the two approaches have remained largely independent: the classic FGM focuses on the phenotype and cannot naturally accommodate a notion of genome architecture being more or less modular or evolvable. The simple framework that we propose bridges this gap. The presentation below is self-contained, but in effect, we build on FGM by separating genotype from phenotype and incorporating a minimal notion of "regulation" (*SI Appendix*, Fig. S1).

The structure of our model is summarized in Fig. 1*A*. The environment is represented by a target vector $\vec{E}$ in an L-dimensional phenotype space. A genome G is a K-element basis in that space (i.e., a set of K basis vectors $\{\vec{g}_\mu\}$). We posit that an organism can adopt any phenotype realizable as a linear combination of

Significance

Harnessing performance tradeoffs offers a promising path for predicting or controlling evolution. If traits A and B are subject to a rigid tradeoff, selecting for improved A necessarily results in a decline in B. If this observation applied to even a single pair of antibiotics, the problem of evolution of antibiotic resistance would have been resolved. Unfortunately, many relevant tradeoffs are not, in fact, rigid. We present a minimally structured effective model that can help establish a null model intuition for how tradeoff evolution can be expected to depend on key parameters and serve as a baseline for comparison with experimental data.

Author contributions: M.T., S.K., and D.S.F. designed research; M.T. performed research; M.T. analyzed data; M.T. wrote the paper; and S.K. and D.S.F. revised the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Data deposition: The associated code and scripts reproducing all of the figures in the paper are available at Mendeley Data, https://doi.org/10.17632/ykypdppy9n.2.

[1]To whom correspondence may be addressed. Email: tikhonov@wustl.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1915537117/-/DCSupplemental.

First published April 3, 2020.

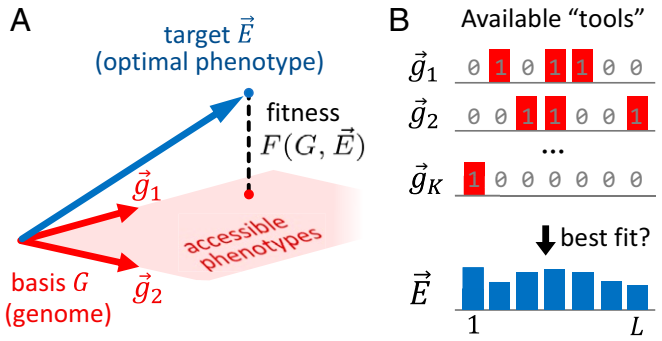

Fig. 1. The toolbox model. (*A*) A genome, *G*, consists of an undercomplete set of *K* basis vectors $\{\vec{g}_\mu\}$ in phenotype space, defining a *K*-dimensional subspace of accessible phenotypes: those that can be represented as linear combinations of $\{\vec{g}_\mu\}$ with positive coefficients (expression levels). The environment is represented by the optimal phenotype $\vec{E}$ (blue dot). The fitness of *G* in this environment is defined by the distance from the optimum to the closest accessible phenotype (red dot). (*B*) For simplicity, we consider $\{\vec{g}_\mu\}$ to be binary vectors of length *L* and assume $\vec{E}$ to be of Euclidean norm 1, with positive components.

its $\{\vec{g}_\mu\}$ with positive coefficients (loosely "expression levels"). An idealized organism with a very flexible physiology capable of independently adjusting each trait as necessary would be able to adopt any phenotype. However, an organism limited to only $K \ll L$ adjustable "knobs" is phenotypically constrained to a K-dimensional subspace. Target phenotypes outside this subspace can only be approximated, and we define fitness $F(G,\vec{E})$ as the (Euclidean) norm of the residual:

$$F(G,\vec{E}) = -\min_{\{a_\mu \geq 0\}} \left\| \vec{E} - \sum_\mu a_\mu \vec{g}_\mu \right\|.$$

In what follows, we will take vectors $\vec{E}$ to be of unit length so that fitness is confined to $-1 \leq F \leq 0$. The simplest interpretation of this model is as a caricature of metabolism, where $\vec{E}$ is the target stoichiometry of a set of L metabolites, and the K knobs [dimensions of internal representation (30)] are the activities of synthesis pathways. Motivated by this example, we take all components of $\vec{E}$ and $\{\vec{g}_\mu\}$ to be positive, and the name "toolbox model" is intended to highlight intriguing conceptual parallels with ref. 31. However, we stress that our setup is not intended as a realistic model of metabolic regulation or any other specific context. Instead, our goal is to construct a minimally structured model that allows the same genotype to be good (or not) in many different environments at the same time. In the setup of Fig. 1*A*, for any set of $N_{\rm env}$ environments with $N_{\rm env} \leq K$, there exists, in principle, a genome G that is perfectly fit ($F = 0$) in all. Conversely, the same environment can be fit by many genomes. For our purposes, we will take $N_{\rm env} \lesssim K \ll L$ so that being fit in $N_{\rm env}$ environments is possible but difficult.

Mathematically, the key feature distinguishing our model from FGM is that the expression levels a_μ are adjusted to be as good as possible in each environment separately. A genome encodes not a single phenotype but a toolbox that allows the phenotype to be environment dependent [phenotypic plasticity (32–35)]. In this sense, our model incorporates a minimal (fast, costless) notion of regulation (*SI Appendix*, Fig. S1). We will see that this has important and nontrivial evolutionary consequences not captured by the classic FGM, such as the evolution of modularity.

Modeling the Evolutionary Process

To study tradeoffs, we must consider more than one environment; we focus here on the simplest case: $N_{\rm env} = 2$. Thus, the central object is an environment pair $\mathcal{P} \equiv \{\vec{E}_A, \vec{E}_B\}$. For simplicity, we consider binary genomes, where all L components of the K vectors $\{\vec{g}_\mu\}$ are either 0 or 1 (Fig. 1*B*) with mutations implemented as bit flips $0 \mapsto 1$ or $1 \mapsto 0$. From each genome, there are thus KL possible mutations. For simplicity, we analyze the strong selection, weak mutation regime where evolution proceeds through a sequence of sweeps (no clonal interference) and only beneficial mutations are relevant (36). This avoids the need to explicitly simulate a population as we only need to track the mutations accumulated by a single adapting lineage.

Given a starting genome G_0 and an environment pair $\mathcal{P} = \{\vec{E}_A, \vec{E}_B\}$, our evolutionary protocol proceeds as follows. One of two environments is chosen equiprobably. The fitness in this environment of all single mutants is evaluated (reoptimizing the expression coefficients every time), and all beneficial mutations are identified. Of these, one "lucky" mutation is drawn with probability proportional to its fitness effect (i.e., to its fixation probability). We refer to this as "one mutational step." After a mutation is accepted, either $\vec{E}_A$ or $\vec{E}_B$ is randomly selected for the next exposure, and the process is repeated. This protocol, adopted for computational simplicity, accepts one mutation per exposure epoch; this approximation does not significantly change our results (*SI Appendix*, Figs. S2 and S3).

The Toolbox Model Exhibits Tradeoffs

We begin by showing that the toolbox model exhibits performance tradeoffs. To do so, we start with $K = 3$ and $L = 6$, small enough that all 2^{KL} genomes can be fully enumerated. Fig. 2*A* shows fitness values F_A, F_B of all genomes in two particular random environments: vectors of length $L = 6$ generated by independently drawing each component from an exponential distribution (to ensure that all components are positive) and normalized to unit length. The Pareto front (Fig. 2*A*, dashed line) owes its existence to the zero/one binarization of genome components, and therefore, its fine structure is sensitive to modeling choices. In subsequent figures, we will focus on early time evolution before the Pareto front is reached, as judged by the mean fitness continuing to increase—the regime where the insights from a simple model are more likely to be generalizable.

Fig. 2*A* shows that a random genome evolving under pressure from one environment becomes mediocre in the other, consistent with the notion of a tradeoff. To define tradeoff strength quantitatively, consider the "mutant cloud" around a given genome (Fig. 2*B*). The scatter plots show the fitness effects δF_A, δF_B of all $KL = 18$ single mutations of several sample genomes evaluated in the two environments. Mutations deleterious in both environments are irrelevant for the evolutionary process and can be ignored. The remaining mutations (beneficial in at least one environment) can be used to define mutational tradeoff strength χ:

$$\chi \equiv -\frac{\sum_{i\in \rm ben} \delta F_{Ai}\, \delta F_{Bi}}{\sqrt{\left(\sum_{i\in \rm ben} \delta F_{Ai}^2\right)\left(\sum_{i\in \rm ben} \delta F_{Bi}^2\right)}},$$

where the index i runs over all mutations beneficial in at least one environment. (Other ways of quantifying tradeoffs are discussed in the *SI Appendix*, *Quantifying Tradeoffs*.) We stress that χ is defined in reference to a particular environment pair $\{\vec{E}_A, \vec{E}_B\}$. This definition ensures that χ ranges from -1 (no tradeoff in identical environments: $\delta F_A = \delta F_B$) to $+1$, the strongest possible tradeoff for which $\delta F_A = -\delta F_B$. A case that will become important shortly is that of a "mutationally modular" genome defined by the property that mutations improving performance in one environment do not affect the other (*SI Appendix*, *Two Definitions of Modularity*). Conveniently, by our definition of χ, modular genomes have mutational tradeoff of zero (Fig. 2*B*).

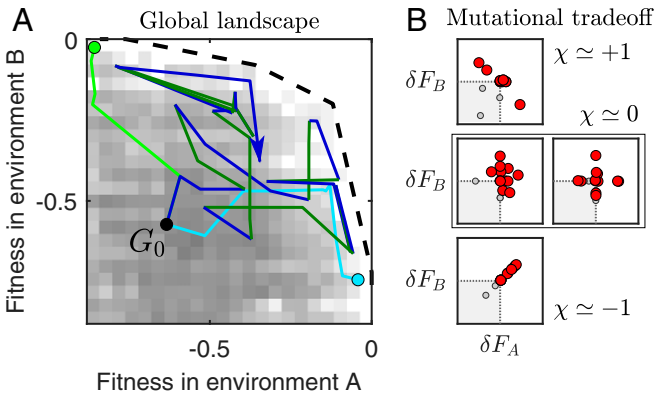

Fig. 2. Quantifying tradeoffs. (*A*) The global fitness landscape computed for $K = 3$ and $L = 6$ (small enough that all genomes can be fully enumerated) in two randomly chosen environments $\vec{E}_A$ and $\vec{E}_B$. The density scatter plot (gray shading) shows the logarithm of the number of genomes per fitness bin; the dashed line roughly traces the Pareto front. A random initial genome G_0 evolving under pressure from one environment becomes mediocre in the other (light blue evolving in *A* or light green evolving in *B*) and runs out of beneficial mutations after $\sim L$ mutation steps. When the same genome evolves in randomly switching environments, the trajectory does not terminate; steps in dark blue and dark green correspond to mutations accepted while exposed to $\vec{E}_A$ and $\vec{E}_B$, respectively; the first 35 steps are shown (some loci flip multiple times). Additional examples are in *SI Appendix*, Fig. S4. (*B*) Examples illustrating our definition of the mutational tradeoff χ. Panels show the relative fitness of all single-step mutants from four example genomes. In red are mutations beneficial in at least one environment. The doubly deleterious mutants (gray) cannot fix and are irrelevant for the adaptive evolutionary process studied herein. The examples show a strong mutational tradeoff (*Top*), the opposite of a tradeoff (*Bottom*), no tradeoff (*Middle Left*), and a mutationally modular genome (*Middle Right*) where mutations beneficial in one environment have little effect on fitness in the other (additional examples in *SI Appendix*, Fig. S5).

Mutational Tradeoff Itself Evolves

What are our expectations for the behavior of the mutational tradeoff χ? First, any notion of tradeoff strength is expected to depend on the difference between environments, which we will quantify by the Euclidean norm $\Delta E = \|\vec{E}_A - \vec{E}_B\|$ ($\vec{E}_{A,B}$ are unit vectors with positive coefficients, and therefore, $0 \le \Delta E \le \sqrt{2}$; in Fig. 2, $\Delta E = 1$). Second, one might expect the χ of an evolving genome to increase with time, as highly evolved genomes should become depleted for jointly beneficial mutations.

To test these expectations, we investigate the dynamics of mutational tradeoff strength for random initial genomes evolving in random environment pairs with various ΔE. As described above, the relevant parameter regime of the toolbox model is $N_{\mathrm{env}} \lesssim K \ll L$. Since our focus is on environment pairs ($N_{\mathrm{env}} = 2$), from now on we will use $K = 4$ and $L = 100$. Our starting genomes will be random binary matrices with each entry set to 1 with probability $p = 0.25$ (so that initially, on average, each trait is affected by one regulator: $pK = 1$). For each trial, the components of both environment vectors are independently drawn from an exponential distribution as before (using a Gaussian distribution does not qualitatively change the results); the vectors are then normalized and rotated toward or away from each other to achieve a desired ΔE (*SI Appendix*, *Parameterizing Environment Pairs*). The results are presented in Fig. 3*A*.

Plotting these data as a function of ΔE (Fig. 3*B*) appears to confirm much of our intuition. First, for random genomes (timepoint 0), mutational tradeoff χ is strongest when ΔE is largest: as expected, being good at two tasks is harder when they differ more (*SI Appendix*, Fig. S6). Second, genomes evolved for 250 mutation steps do indeed exhibit a larger mutational tradeoff. For a sense of scale, recall that our genomes have $KL = 400$ loci, and therefore, 250 steps are sufficient for over half of the bits to flip.

Curiously, however, Fig. 3*A* also shows that, after ~ 100 mutation steps, genomes evolving at a large ΔE consistently exhibit a counterintuitive decline in χ. This behavior is not a peculiarity of our definition of mutational tradeoff; other measures of tradeoff strength exhibit the same phenomenon (*SI Appendix*, Fig. S7). Examining the evolved genomes in more detail reveals that this decline in χ is associated with genomes becoming more modular as reflected in the "plus"-like shape of the example mutant cloud shown in Fig. 3*C* (compare with Fig. 2*B*). Quantitatively, we can define mutational modularity M (again, in reference to a particular environment pair):

$$M = 1 - \frac{\sum_i \rho_i |\sin 2\phi_i|}{\sum_i \rho_i},$$

where (ρ_i, ϕ_i) are the mutation effects transformed to polar coordinates: $\rho = \sqrt{\delta F_A^2 + \delta F_B^2}$, $\phi = \arctan \frac{\delta F_B}{\delta F_A}$, and i runs over all KL mutations, including deleterious ones. At $M = 1$ (perfect modularity), each mutation only affects fitness in one environment (polar angle $\phi = 0, \pm\frac{\pi}{2}$, or π) (*SI Appendix*, Fig. S8). Using this definition, Fig. 3*D* shows the same evolutionary trajectories in the "mutational tradeoff vs. mutational modularity" plane. We observe that evolution quickly drives the tradeoff strength to its maximum before slowly increasing modularity, and the modularity increase is more pronounced at larger values of ΔE as summarized in Fig. 3*E*.* Reassuringly, what holds for the entire cloud also applies to the top mutants (most relevant for evolution). Denote by $\delta F_A^{\mathrm{best}}$, $\delta F_B^{\mathrm{best}}$ the effect on F_A of the best mutation in environment B and vice versa (Fig. 3*C*). For a mutationally modular genome, one expects both to be close to 0: $\delta F_{A,B}^{\mathrm{best}} \approx 0$. Fig. 3*F* compares earlier- and later-timepoint histograms of $\delta F_{A,B}^{\mathrm{best}}$ over 100 independent replicates (random starting points, random environment pairs with $\Delta E = 0.9$; the histogram does not distinguish $\delta F_A^{\mathrm{best}}$ and $\delta F_B^{\mathrm{best}}$ because of the symmetry swapping labels A and B). The later timepoint shows a clear enrichment of modular genomes (the peak at $\delta F_{A,B}^{\mathrm{best}} \approx 0$).

One explanation for this enrichment might be that high-fitness genomes are generally more modular. However, we will now demonstrate that evolution at a large ΔE specifically promotes tradeoff weakening: even conditioned on having the same high fitness, the χ of evolved genomes remains atypical (37). Specifically, we will show that 1) high-fitness genomes can exhibit a whole range of mutational tradeoff values, 2) prior evolutionary history predictably pushes genomes into different regions of this high-fitness space, and 3) the resulting genomes, although sharing the same fitness, can differ dramatically in their properties (for instance, in their ability to evolve further).

Evolutionary History Shapes Mutational Tradeoff

Fig. 3*B* demonstrated that mutational tradeoff strength of evolved genomes is not typical of all genomes (the red curve is clearly distinct from the blue curve). Establishing whether evolution specifically promotes mutational tradeoff weakening requires a stronger statement, namely that the χ arising through

*Note, however, that both χ and M are explicitly properties of both the genome and the environment pair, and one should be cautious not to overinterpret a direct comparison of these metrics across different values of ΔE. For instance, the trajectories shown in Fig. 3*F* all start from similar (random) genomes, but the initial χ and M are different because each trajectory is being evaluated in a different environment pair. The effect of evolution at large ΔE on promoting modularity in our model is a hypothesis suggested by Fig. 3*F* but verified in Fig. 4, where all genomes are evaluated in the same environment, and direct comparisons of their fitness, mutational tradeoff, and/or mutational modularity become meaningful.

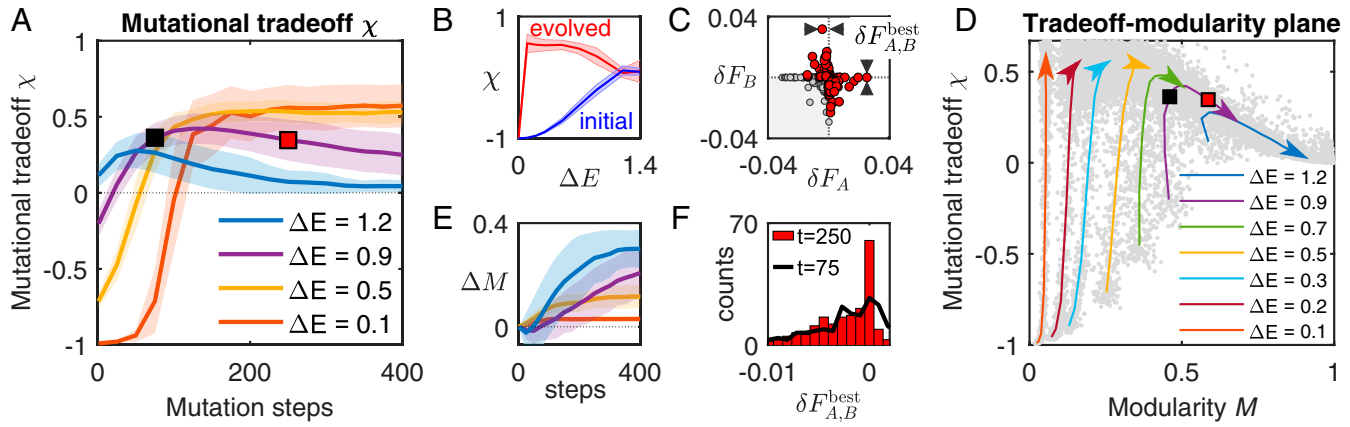

Fig. 3. Evolved tradeoffs are not typical. (*A*) Mutational tradeoff (compare with Fig. 2*B*) as a function of the number of mutation steps (a proxy for time) for evolution in random environment pairs differing by a given amount quantified by ΔE. Each curve shows mean ± 1 SD over 100 independent trials (shaded); for each trial, a new random initial genome was evolved in a new random environment pair ($K = 4$, $L = 100$). Initially, more different tasks are associated with stronger mutational tradeoff, and tradeoff strength generally increases with time, with a notable exception at large ΔE. (*B*) Tradeoffs after 250 mutational steps vs. environment difference ΔE. For $\Delta E = 0$ (two identical environments), the mutational tradeoff is necessarily -1. However, when the environments differ, evolution drives tradeoff strength to strongly atypical values. Shown is mean ± 1 SD over 100 independent trials for each ΔE. (*C*) The mutant cloud around an example genome evolved for 250 steps in two very different environments ($\Delta E = 0.9$) exhibits mutational modularity: the best mutant in one environment has only a weak effect (denoted δF_A^{best}, δF_B^{best}) in the other. (*D*) Evolutionary trajectories of *A* replotted against the modularity score *M* (in the text). Arrows trace the mean over 100 random initial genomes, each evolving in a random environment pair with the specified ΔE; points show the individual genomes (every 25 mutation steps). A perfectly modular genome ($M = 1$) necessarily has $\chi = 0$; conversely, $\chi = \pm 1$ entails $M = 0$. As a result, the datapoints fill out a triangular allowed zone. The two squares on the $\Delta E = 0.9$ trajectory mark the same timepoints as in *A*. (*E*) Change in modularity score *M* relative to the initial timepoint for the trajectories of *A* (compare with *D*). Shown is mean ± 1 SD; colors are the same as in *A*. The plot confirms that evolution at large ΔE is associated with an increase in modularity. (*F*) Histograms of $\delta F_{A,B}^{\text{best}}$ over 100 trajectories evolved at $\Delta E = 0.9$ at the two timepoints highlighted in *A* and *D*. The later timepoint (red) shows a clear enrichment of modular genomes (the peak at $\delta F_{A,B}^{\text{best}} \approx 0$) compared with the earlier timepoint (black).

evolution is not even typical of high-fitness genomes. To show this, the most direct approach would be to compare the evolved χ with the typical values observed when sampling high-fitness genomes in an unbiased way. Unfortunately, we have no procedure for such an unbiased sampling other than a complete enumeration, which is only viable for extremely small K and L (compare with Fig. 2). Instead, we will reach the same conclusion by showing that different ways of evolving high fitness lead to genomes with different mutational tradeoff strength.

Consider the following computational experiment. Generate one random environment pair $\mathcal{P}^* \equiv \{\vec{E}_A^*, \vec{E}_B^*\}$; we used a pair with $\Delta E = 0.9$. Define fitness in the environment pair as the average over fitness in $\vec{E}_A^*$ and $\vec{E}_B^*$:

$$F(G, \mathcal{P}^*) \equiv \frac{F\left(G, \vec{E}_A^*\right) + F\left(G, \vec{E}_B^*\right)}{2}.$$

Throughout the experiment, we will only be concerned with fitness and mutational tradeoff as measured in this pair $\mathcal{P}^*$. Under our protocol, the success of any mutation depends only on its fitness effect in the one environment to which the genome is exposed at the time. However, as exposures alternate, the average fitness in the pair will typically also increase.

What other evolutionary protocols could lead to increased fitness in $\mathcal{P}^* = \{\vec{E}_A^*, \vec{E}_B^*\}$? One cannot expect that evolving in some random other pair will increase the mean fitness in $\mathcal{P}^*$, but one can consider evolving a genome in similar environments or the average of $\vec{E}_A^*$ and $\vec{E}_B^*$. In our model, for any $\{\vec{E}_A^*, \vec{E}_B^*\}$, one can create similar pairs of environments by rotating the vectors toward or away from each other (in log space to preserve positivity of components) (*SI Appendix*). This yields a one-parameter family of environment pairs $\mathcal{P}(\Delta E)$ indexed by their difference ΔE (Fig. 4*A*). For a concrete analogy, if $\vec{E}_A$ and $\vec{E}_B$ represent hot and cold seasons, then $\mathcal{P}(\Delta E)$ is a family of environment pairs where the intensity of seasonal variation is more mild or more severe (quantified by ΔE). If we preevolve the same 20 random starting genomes in these conditions, with ΔE exaggerated or softened as described, what effect will this have on the fitness and mutational tradeoff strength as measured in the pair of interest, $\mathcal{P}^*$?

The results are presented in Fig. 4*B*, showing the mutational tradeoff vs. average fitness, both measured in $\mathcal{P}^*$. The black solid line in Fig. 4*B* corresponds to the simplest scenario, where the 20 initial genomes are evolved directly in $\mathcal{P}^*$. The evolutionary time runs left to right, reflected in increasing fitness. If our hypothesis is correct, evolving the same initial genomes at a larger ΔE (red in Fig. 4 *B* and *C*) should promote a more modular genome architecture and a weaker tradeoff; this is indeed what we observe (Fig. 4 *B* and *C* and *SI Appendix*, Fig. S9). Conversely, if instead of exaggerating the differences, we soften them to $\Delta E = 0$ (i.e., replace the environment pair $\mathcal{P}^*$ with a single environment, their mean), we obtain the trajectory in blue in Fig. 4 *B* and *C*. (Note that the mutational tradeoff evaluated in an environment "pair" with $\Delta E = 0$ is always -1, but Fig. 4 *B* and *C* shows the mutational tradeoff evaluated in the original pair of interest $\mathcal{P}^*$.)

Fig. 4 *B* and *C* directly demonstrates that, within our model, manipulating evolutionary history predictably pushes evolving genomes toward stronger or weaker modularity and mutational tradeoff. In particular, if we pick similarly performing genomes from the right-hand side of this plot, we will find that they all attain the same mean performance in different ways, with a wide range of tradeoff values. We can now ask: how does this difference affect their near-term evolutionary future?

Tradeoff Strength Shapes Evolution

To address how the evolutionary history and current tradeoff strength affect future evolution, we used the protocol of Fig. 4*B* (computationally preevolving the same 20 initial genomes in

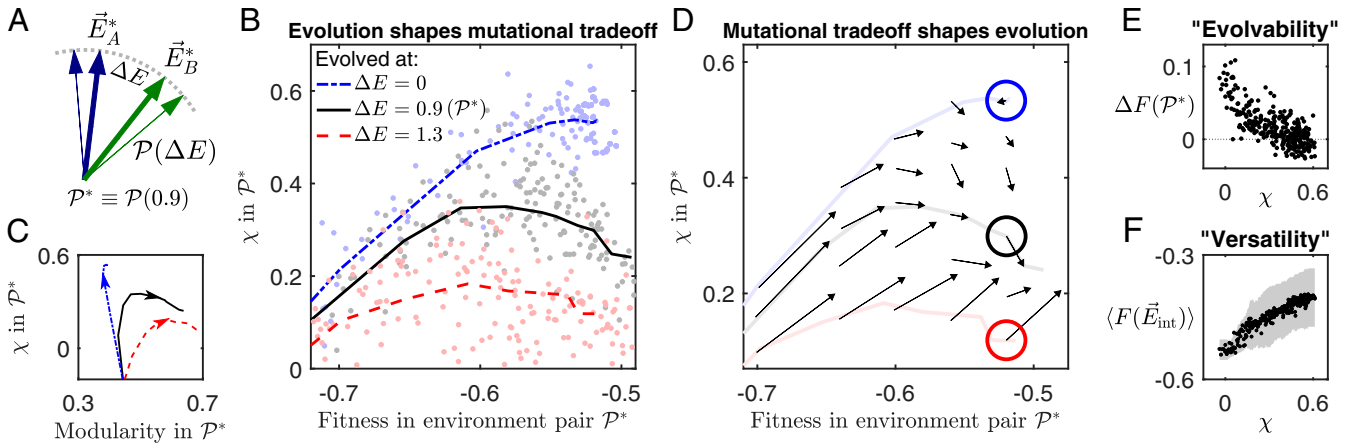

Fig. 4. Evolutionary history shapes tradeoff structure and vice versa. (*A*) For a randomly chosen pair of environments $\vec{E}_A^*$, $\vec{E}_B^*$, the difference between them can be softened or exaggerated by rotating the vectors toward or away from each other (in the log space to preserve positivity) (*SI Appendix*). This family of pairs is denoted $\mathcal{P}(\Delta E)$, with $\mathcal{P}^*$ the original pair that we chose to have $\Delta E = 0.9$. (*B*) The same 20 random initial genomes ($K = 4$, $L = 100$) were evolved for 250 mutational steps: either directly in $\mathcal{P}^*$ (solid black line) or in another environment pair from the $\mathcal{P}(\Delta E)$ family, with $\Delta E = 0$ (the average environment) or $\Delta E = 1.3$ (exaggerated differences). The mutational tradeoff measured in the original environment pair $\mathcal{P}^*$ is plotted against the mean fitness in $\mathcal{P}^*$ (i.e., genomes are evolved in different environment pairs but are all evaluated in the same pair). Data points show the individual genomes every 25 mutational steps; each line is the average over the 20 trajectories evolved in the same environment pair. We observe that different evolutionary histories consistently drive genomes to different tradeoff strengths, even when compared at the same fitness. (*C*) Trajectories of *B* shown on the tradeoff vs. modularity plane; colors are the same as in *B*. Evolution at a low ΔE yields high-tradeoff, low-modularity genomes, while an exaggerated ΔE leads to modularity increase (*SI Appendix*, Fig. S9). (*D*) Mutational tradeoff shapes the near-term evolutionary future. Here, sets of 20 genomes, prepared by preevolving for varying lengths of time at different ΔE as in *B*, were "transferred" to continue evolution in the pair $\mathcal{P}^*$. For each set, an arrow shows the average change of fitness and mutational tradeoff in $\mathcal{P}^*$ over the next 20 mutation steps. Trajectories from *B* are shown for reference. The contrast in outcomes is particularly striking for the highest-fitness genomes (colored circles; see also *E*). (*E*) Evolvability of genomes with different mutational tradeoff. The 20 genomes of *B* were each evolved under 15 different values of ΔE until reaching fitness -0.5 ± 0.001 when evaluated in $\mathcal{P}^*$. This yielded a set of 300 genomes with the same fitness in $\mathcal{P}^*$ but different mutational tradeoff χ. Each was then evolved in $\mathcal{P}^*$ for 20 steps; the mean fitness gain (over the two environments in the pair $\mathcal{P}^*$) is plotted against the initial value of mutational tradeoff χ. The genomes with the weakest χ evolve fastest, whereas genomes with the highest χ actually decline in mean fitness (compare with *D*). (*F*) The most evolvable genomes are also the least "versatile." The same 300 genomes as in *E* are evaluated in random "intermediate" environments $\vec{E}_{\rm int}$ between $\vec{E}_A^*$ and $\vec{E}_B^*$ (in the text). Plotted is mean fitness $\langle F(\vec{E}_{\rm int})\rangle$ over 100 intermediate environments $\pm$ 1 SD.

15 different environment pairs for a varying length of time) to obtain a collection of genomes evenly populating a region of the fitness vs. mutational tradeoff plane. Each was then evolved for 20 further mutational steps in our pair of interest $\mathcal{P}^*$. The result is presented in Fig. 4*D*, where each arrow describes the change in the mean fitness and tradeoff strength in $\mathcal{P}^*$ observed over these 20 mutation steps (averaged over the 20 genomes in the set).

First, notice that arrows in Fig. 4*D* in the vicinity of the black trajectory are tangent to it. This, of course, is exactly what we expect, since the black trajectory (copied from Fig. 4*B*) traces the evolution in this same pair $\mathcal{P}^*$. However, arrows starting elsewhere in Fig. 4*D* on this plane (preevolved at a different ΔE, then transferred into $\mathcal{P}^*$) show different behaviors. It is especially interesting to compare the arrows on the far right side of the plot (colored circles in Fig. 4*D*). As a reminder, under our protocol, all genome sets were obtained from the same 20 initial genomes and differ only by evolutionary history. Yet, evolving them in the same environment $\mathcal{P}^*$ for the same amount of time (20 mutation steps) leads to very different outcomes. The low-tradeoff genomes obtained by preevolving at an exaggerated ΔE (red circle in Fig. 4*D*) exhibit a dramatically faster speed of subsequent evolution in $\mathcal{P}^*$ compared with similar fitness genomes that never left this environment pair (black circle in Fig. 4*D*). In contrast, attempting to further evolve the highest-tradeoff genomes (blue circle in Fig. 4*D*) actually leads to a decrease in mean fitness. For these genomes, mutations beneficial in one environment of the pair are so strongly deleterious in the other that the mean fitness declines. To summarize, genomes with different evolutionary history have predictably different evolutionary future.

This observation is further illustrated in Fig. 4*E*, where the mean fitness gain over 20 generations is shown for 300 individual genomes, all with the same initial mean fitness -0.5 ± 0.001, plotted against their initial mutational tradeoff strength. The plot directly confirms that weak-tradeoff genomes evolve fast, whereas for the strongest-tradeoff genomes, the fitness gain dips into the negative.

The relation observed in Fig. 4*E* is consistent with the previously proposed idea that modular architecture is more "evolvable" (28, 38, 39). Intriguingly, however, our framework allows the benefits of modularity to be nuanced. For example, rather than continuing evolution in the pair $\mathcal{P}^* = \{\vec{E}_A^*, \vec{E}_B^*\}$, consider instead generating a large number of intermediate environments $\vec{E}$, where each component is independently and uniformly drawn to be between the respective component of $\vec{E}_A^*$ and $\vec{E}_B^*$. Fig. 4*F* shows the average fitness of the exact same 300 genomes in the intermediate environments. The trend is now reversed: the "best" genomes as judged by Fig. 4*E* perform worst in this test. Intuitively, while a modular architecture (for a given environment pair) facilitates continuing adaptation to that same pair (Fig. 4*E*), it also suffers from a form of "overfitting": the nonmodular architectures may be more versatile when evaluated across a range of similar environments (Fig. 4*F*).

In summary, Fig. 4 demonstrates that evolution shapes mutational tradeoff and conversely, that mutational tradeoff shapes evolution. We will conclude by exhibiting a nontrivial phenomenon arising as a consequence of this feedback loop.

Best Adaptation for One Environment Pair Evolves in Another

Consider the problem in which a pair $\mathcal{P}^*$ is given, and we would like to evolve a random starting genome G_0 toward a high mean fitness in this particular pair of environments. The most natural approach, of course, would be to evolve G_0 directly in $\mathcal{P}^*$.

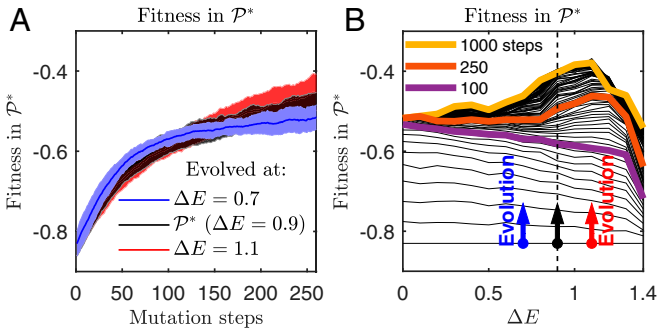

Fig. 5. Best adaptation for one environment evolves in another. (*A*) The same 20 random initial genomes ($K = 4$, $L = 100$) are evolved at different ΔE and evaluated at $\Delta E = 0.9$ (in $\mathcal{P}$*). Shown is mean fitness $\pm$ 1 SD over the 20 trajectories. Remarkably, the genomes evolving directly in the environment pair of interest (black) are typically less fit than those evolved from the same initial genomes and for the same number of mutation steps but in one of the other environment pairs with softened (blue) or exaggerated (red) differences. (*B*) Same as *A* replotted vs. ΔE. The same initial genome is evolved in different environment pairs; contour lines spaced by 10 mutation steps show the mean fitness evaluated in the environment pair of interest ($\mathcal{P}$* with $\Delta E = 0.9$). Plot is averaged over 20 random instances of the initial genome. The three arrows correspond to the three trajectories shown in *A*. We see that, generically, the genomes with the highest fitness in $\mathcal{P}$* were evolved in other environment pairs: after 100 mutation steps, the genomes performing best at $\Delta E = 0.9$ are those evolved at $\Delta E = 0$, while at 250 mutation steps, the best were evolved at $\Delta E = 1.1$. See text for explanation.

However, the results of Fig. 4 suggest that evolving instead in a different environment pair, with differences exaggerated or softened, might offer a more efficient path toward $\mathcal{P}^*$-fitness increase. Fig. 5*A* confirms this by replotting the trajectories from Fig. 4*A*, showing fitness in the environment pair of interest $\mathcal{P}^*$ (with $\Delta E = 0.9$) as a function of time for genomes evolving under three different values of ΔE. Importantly, in this example, genomes evolving directly in $\mathcal{P}^*$ are typically less fit than those evolved from the same initial genomes and for the same number of mutation steps but in a different environment pair (blue or red in Fig. 5*A*).

A more detailed visualization of this phenomenon is presented in Fig. 5*B*. In the hot and cold season analogy, think of the ΔE axis in Fig. 5*B* as a transect through a continent, running from a region with no seasonal variation ($\Delta E = 0$) through regions where it is increasingly more severe; we are focusing on a particular temperate region $\mathcal{P}^*$ with moderate seasonal variation ($\Delta E^* = 0.9$). Imagine populating this continent with an initially clonal population (the initial genome G_0). Initially neglecting migration, we let these genomes independently evolve in their respective environments. However, every 10 mutation steps, we consider potential immigrants to the temperate region of interest: namely, we evaluate the performance of all of the lineages in the pair $\mathcal{P}^*$ and plot it on the y axis.

By construction, initially all genomes are the same, and their performance in $\mathcal{P}^*$ is also the same. However, 100 mutation steps later (purple in Fig. 5*B*), the highest fitness in $\mathcal{P}^*$ is exhibited by genomes that evolved at $\Delta E = 0$. These genomes were evolved under selection pressure from a single averaged environment and cannot, of course, develop any hot- or cold-specific adaptations, but at the early stages, evolution proceeds very efficiently as there are no conflicting pressures from the two environments.

As evolution proceeds further, the behavior inverts. After 250 mutation steps (orange in Fig. 5*B*), the highest fitness in $\mathcal{P}^*$ is exhibited by genomes that evolved at $\Delta E \approx 1.1$. Initially, the increase in mean fitness under such harsh seasonal variations was very slow, with beneficial mutations in one environment undoing the gain made in the other. However, eventually, the lineages evolving under this large ΔE develop a weakened mutational tradeoff (compare with Fig. 4): this enables them to gain fitness more efficiently, overtaking the lineages evolving in $\mathcal{P}^*$. In this model, modularity is thus good for evolvability, but it takes time and the right conditions to itself evolve.

It is worth noting that this effect can persist even in the long run (shown in yellow in in Fig. 5*B* for 1,000 steps; additional discussion is in *SI Appendix*, Fig. S10). One implication of this in our spatial/seasonal metaphor is that, in the presence of migration, the phenomenon described will lead to a qualitative change in the expected genealogy structure of the long-term surviving lineages. In particular, the environment pairs with large variations turn into sources of evolutionary innovation, promoting the development of $\vec{E}_A$- and $\vec{E}_B$-specific adaptations capable of invading other environments and outcompeting the resident lineages. Importantly, in our model, this is true even though we have explicitly excluded the effects of diversity that would in general be created in each environment.

Discussion

Understanding evolution in multiple or changing environments requires developing an understanding of which phenomena observed in the laboratory (or the wild) depend on specific details, which are more general, which experimental outcomes are truly surprising, and which can be found in a very simple model. We presented a toy model able to capture the key feedback loop of evolutionary tradeoff plasticity, whereby organisms evolving in different environments are constrained by performance tradeoffs, but such tradeoffs themselves depend on the evolutionary history. By defining fitness through a regulatory or physiological optimization problem, our approach is reminiscent of asking evolutionary questions from within the framework of flux balance analysis (40–46). However, our "toolbox" model abstracts away any specifically metabolic (or other) detail, retaining only the flexibility of regulation (i.e., the fact that genome-encoded tools can be used in an environment-dependent manner). Remarkably, this simple model already exhibits qualitatively rich phenomena, including effects that change sign during the course of evolution.

One of the effects that we studied by example shows how the best adaptation for one environment may be expected to first emerge in another. Examples of this phenomenon are known experimentally: for instance, the fastest way to evolve resistance to a high dose of antibiotic is through a series of exposures to increasing doses rather than direct pressure from the environment of interest (19, 23). Our results suggest that, for evolution considered across multiple environments, this scenario may well be generic. If true, this is likely to impose a strong constraint on the predictive power of single-environment evolutionary models: successful lineages coming from elsewhere are beyond their scope, yet successful invaders can have a profound impact, including on genealogical structures, a key observable used for inferences from data. This highlights the need for simple illustrative models able to incorporate explicit dependence on environment to help develop null model phenomenological expectations, against which experimental results can be compared (21).

The phenomena reported here did not require fine tuning of model parameters; they arise because the space of high-fitness genomes is naturally large and diverse so that different evolutionary histories generically bias evolved genomes toward different corners of the high-fitness space. In our model, this property is ensured by choosing the number of tools that can be independently regulated to be larger than the number of environments probed ($K > N_{\mathrm{env}}$), but the observation that there are many ways to be fit should surely not be model specific.

Establishing the generality of our observations, extending them to multiple environments (rather than just pairs), and relating the predictions of our framework to experiments all constitute productive directions for future work. Furthermore, our simple model assumed that organisms sense their environment perfectly and regulate their physiology optimally. Relaxing these assumptions offers natural directions in which our model could be extended.

Materials and Methods

All simulations were performed in MATLAB (Mathworks, Inc.). The associated code and scripts reproducing all of the figures in the paper are available at Mendeley Data, https://doi.org/10.17632/ykypdppy9n.2 (47).

ACKNOWLEDGMENTS. We thank A. Agarwala, D. Finlay, S. Leibler, R. Monasson, J. Moran, and A. Murugan. This work was supported in part by NSF Grants PHY-1607606, PHY-1545840, and PHY-1720397 and a Simons Investigator Award (to S.K.).

1. D. A. Roff, D. J. Fairbairn, The evolution of trade-offs: Where are we? *J. Evol. Biol.* **20**, 433–447 (2007).
2. H. S. Callahan, H. Maughan, U. K. Steiner, Phenotypic plasticity, costs of phenotypes, and costs of plasticity: Toward an integrative view. *Ann. N. Y. Acad. Sci.* **1133**, 44–66 (2008).
3. S. Klumpp, T. Hwa, Growth-rate-dependent partitioning of rna polymerases in bacteria. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 20245–20250 (2008).
4. C. T. Kremer, C. A. Klausmeier, Coexistence in a variable environment: Eco-evolutionary perspectives. *J. Theor. Biol.* **339**, 14–25 (2013).
5. J. R. Meyer, I. Gudelj, R. Beardmore, Biophysical mechanisms that maintain biodiversity through trade-offs. *Nat. Commun.* **6**, 6278 (2015).
6. A. Posfai, T. Taillefumier, N. S. Wingreen, Metabolic trade-offs promote diversity in a model ecosystem. *Phys. Rev. Lett.* **118**, 028103 (2017).
7. K. A. Mooney, R. Halitschke, A. Kessler, A. A. Agrawal, Evolutionary trade-offs in plants mediate the strength of trophic cascades. *Science* **327**, 1642–1644 (2010).
8. O. Shoval *et al.*, Evolutionary trade-offs, pareto optimality, and the geometry of phenotype space. *Science* **336**, 1157–1160 (2012).
9. D. T. Fraebel *et al.*, Environment determines evolutionary trajectory in a constrained phenotypic space. *eLife* **6**, e24669 (2017).
10. G. A. Parker, J. Maynard Smith, Optimality theory in evolutionary biology. *Nature* **348**, 27–33 (1990).
11. A. F. Bennett, R. E. Lenski, An experimental test of evolutionary trade-offs during temperature adaptation. *Proc. Natl. Acad. Sci. U.S.A.*, **104** (suppl. 1), 8649–8654 (2007).
12. Y. Li, D. A. Petrov, G. Sherlock, Single nucleotide mapping of trait space reveals Pareto fronts that constrain adaptation. *Nat. Ecol. Evol.* **3**, 1539–1551 (2019).
13. M. Novak, T. Pfeiffer, R. E. Lenski, U. Sauer, S. Bonhoeffer, Experimental tests for an evolutionary trade-off between growth rate and yield in *E. coli*. *Am. Nat.* **168**, 242–251 (2006).
14. A. Buckling, M. A. Brockhurst, M. Travisano, P. B. Rainey, Experimental adaptation to high and low quality environments under different scales of temporal variation. *J. Evol. Biol.* **20**, 296–300 (2007).
15. J. N. Jasmin, C. Zeyl, Evolution of pleiotropic costs in experimental populations. *J. Evol. Biol.* **26**, 1363–1369 (2013).
16. A. Rodriguez-Verdugo, D. Carrillo-Cisneros, A. Gonzalez-Gonzalez, B. S. Gaut, A. F. Bennett, Different tradeoffs result from alternate genetic adaptations to a common environment. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 12121–12126 (2014).
17. R. S. Satterwhite, T. F. Cooper, Constraints on adaptation of escherichia coli to mixed-resource environments increase over time. *Evolution* **69**, 2067–2078 (2015).
18. M. Baym, L. K. Stone, R. Kishony, Multidrug evolutionary strategies to reverse antibiotic resistance. *Science* **351**, aad3292 (2016).
19. L. W. McGee *et al.*, Synergistic pleiotropy overrides the costs of complexity in viral adaptation. *Genetics* **202**, 285–295 (2016).
20. L. M. Bono, L. B. Smith, D. W. Pfennig, C. L. Burch, The emergence of performance trade-offs during local adaptation: Insights from experimental evolution. *Mol. Ecol.* **26**, 1720–1733 (2017).
21. P. Yen, J. A. Papin, History of antibiotic adaptation influences microbial evolutionary dynamics during subsequent treatment. *PLoS Biol.* **15**, e2001586 (2017).
22. S. Khan, T. K. Beattie, C. W. Knapp, The use of minimum selectable concentrations (mscs) for determining the selection of antimicrobial resistant bacteria. *Ecotoxicology* **26**, 283–292 (2017).
23. E. Wistrand-Yuen *et al.*, Evolution of high-level resistance during low-level antibiotic exposure. *Nat. Commun.* **9**, 1599 (2018).
24. R. H. Y. Louie, K. J. Kaczorowski, J. P. Barton, A. K. Chakraborty, M. R. McKay, Fitness landscape of the human immunodeficiency virus envelope protein that is targeted by antibodies. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E564–E573 (2018).
25. I. N. Berezovsky, E. I. Shakhnovich, Physics and evolution of thermophilic adaptation. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 12742–12747 (2005).
26. R. A. Fisher, *The Genetical Theory of Natural Selection* (Clarendon Press, Oxford, UK, 1930).
27. O. Tenaillon, The utility of Fisher's geometric model in evolutionary genetics. *Annu. Rev. Ecol. Evol. Syst.* **45**, 179–201 (2014).
28. G. P. Wagner, M. Pavlicev, J. M. Cheverud, The road to modularity. *Nat. Rev. Genet.* **8**, 921–931 (2007).
29. L. M. Bono, J. A. Draghi, P. E. Turner, Evolvability costs of niche expansion. *Trends Genet.* **36**, 14–23 (2019).
30. B. K. Xue, P. Sartori, S. Leibler, Environment-to-phenotype mapping and adaptation strategies in varying environments. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 13847–13855 (2019).
31. S. Maslov, S. Krishna, T. Y. Pang, K. Sneppen, Toolbox model of evolution of prokaryotic metabolic networks and their regulation. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 9743–9748 (2009).
32. S. M. Scheiner, Genetics and evolution of phenotypic plasticity. *Annu. Rev. Ecol. Systemat.* **24**, 35–68 (1993).
33. M. Pigliucci, *Phenotypic Plasticity: Beyond Nature and Nurture* (Johns Hopkins University Press, Baltimore, MD, 2001).
34. T. J. DeWitt, S. M. Scheiner, *Phenotypic Plasticity: Functional and Conceptual Approaches* (Oxford University Press, Oxford, UK, 2004).
35. R. J. Fox, J. M. Donelson, C. Schunter, T. Ravasi, J. D. Gaitán-Espitia, Beyond buying time: The role of plasticity in phenotypic adaptation to rapid environmental change. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **374**, 20180174 (2019).
36. J. H. Gillespie, A simple stochastic gene substitution model. *Theor. Popul. Biol.* **23**, 202–215 (1983).
37. A. Agarwala, D. S. Fisher, Adaptive walks on high-dimensional fitness landscapes and seascapes with distance-dependent statistics. *Theor. Popul. Biol.* **130**, 13–49 (2019).
38. N. Kashtan, E. Noor, U. Alon, Varying environments can speed up evolution. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 13711–13716 (2007).
39. D. M. Lorenz, A. Jeng, M. W. Deem, The emergence of modularity in biological systems. *Phys. Life Rev.* **8**, 129–160 (2011).
40. J. D. Orth, I. Thiele, B. O. Palsson, What is flux balance analysis?. *Nat. Biotechnol.* **28**, 245–248 (2010).
41. R. Schuetz, N. Zamboni, M. Zampieri, M. Heinemann, U. Sauer, Multidimensional optimality of microbial metabolism. *Science* **336**, 601–604 (2012).
42. W. R. Harcombe, N. F. Delaney, N. Leiby, N. Klitgord, C. J. Marx, The ability of flux balance analysis to predict evolution of central metabolism scales with the initial distance to the optimum. *PLoS Comput. Biol.* **9**, e1003091 (2013).
43. T. Grosskopf *et al.*, Metabolic modelling in a dynamic evolutionary framework predicts adaptive diversification of bacteria in a long-term evolution experiment. *BMC Evol. Biol.* **16**, 163 (2016).
44. Q. Zhao, A. I. Stettner, E. Reznik, I. C. Paschalidis, D. Segre, Mapping the landscape of metabolic goals of a cell. *Genome Biol.* **17**, 109 (2016).
45. D. Bajic, J. C. C. Vila, Z. D. Blount, A. Sanchez, On the deformability of an empirical fitness landscape by microbial evolution. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 11286–11291 (2018).
46. M. Thommes, T. Wang, Q. Zhao, I. C. Paschalidis, D. Segre, Designing metabolic division of labor in microbial communities. *mSystems* **4**, e00263-18 (2019).
47. M. Tikhonov, A model for the interplay between tradeoff plasticity and evolution in changing environment. Mendeley Data. https://doi.org/10.17632/ykypdppy9n.2. Deposited 20 March 2020.