

**Enhancing Evolution:
When Spatial Structure and Environmental Change Improve Adaptation**

Joshua Richard Nahum

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy
University of Washington

2013

Reading Committee:

Benjamin Kerr, Chair

Carl Bergstrom

Harvey Bradshaw

Program Authorized to Offer Degree:

Department of Biology

©Copyright 2013

Joshua Richard Nahum

University of Washington

Abstract

Enhancing Evolution:

When Spatial Structure and Environmental Change Improve Adaptation

Joshua Richard Nahum

Chair of the Supervisory Committee:

Tenured Professor Benjamin Kerr

Department of Biology

Evolution by natural selection is the primary creative force in biology. In order for populations to thrive, or even survive, requires a continual generative process. An understand of the process of natural selection, its strengths and weaknesses, is necessary for both predictions about the future course of life, as well as its history. Evolution by natural selection requires the combination of three elements: variation, heredity, and selection. In natural populations, variation can be introduced by random mutation, recombination (such as by sex), and migration from differing populations. Heredity, or the correlation between traits of the parents with that of their offspring, is mediated by genetic material. Selection is the process

whereby the traits of some members of the population are more successful (fit) and hence contribute more offspring to future generations. When these three factors are present, populations can (over the course of many generations) increase in fitness to become better adapted to their environment. The ability to adapt is not restricted to populations in natural environments, as artificial populations such as microbes grown in laboratory settings can also evolve. In fact, simulated, digital populations that possess the three features needed for evolution by natural selection can also adapt. In this dissertation, I utilize both microbes evolving in artificial settings and digital organisms, which exist only in a computer's memory, to make general statements about the process of evolution. Of course, such generalizations must be taken with a grain of salt, as the specifics of any system may interfere with assumptions made in the follow models. But, I present the following work as forays into the possible; I demonstrate that certain processes can affect the course of evolution, and such processes should be taken into account when seeking an understand of evolution in other (perhaps more natural) situations. In Chapter 1, I investigate how a non-transitive system of microbial competitors evolves, and demonstrate that the phenomenon of "survival of the weakest" requires the presence of population structure. In Chapter 2, I explore the effect of structure on evolving populations, and I find that structure can enhance the rate of adaptation in certain circumstances. Lastly in Chapter 3, I describe and demonstrate circumstances where environmental change can enhance the rate of adaption.

Chapter 1: The evolution of restraint in a structured rock-paper-scissors community

Previously published as (1)

Abstract

It is not immediately clear how costly behavior that benefits others evolves by natural selection. By saving on inherent costs, individuals that do not contribute socially have a selective advantage over altruists if both types receive equal benefits. Restrained consumption of a common resource is a form of altruism. The cost of this kind of prudent behavior is that restrained individuals give up resources to less restrained individuals. The benefit of restraint is that better resource management may prolong the persistence of the group. One way to dodge the problem of defection is for altruists to interact disproportionately with other altruists. With limited dispersal, restrained individuals persist due to interaction with like types, while it is the unrestrained individuals that must face the negative long-term consequences of their rapacity. Here, we study the evolution of restraint in a community of three competitors exhibiting a non-transitive (“rock-paper-scissors”) relationship. The non-transitivity ensures a form of negative feedback whereby improvement in growth of one competitor has the counterintuitive consequence of *lowering* the density of that ‘improved’ player. This negative feedback generates detrimental long-term consequences for unrestrained growth. Using both computer simulations and evolution experiments with a non-transitive community of

Escherichia coli, we find that restrained growth can evolve under conditions of limited dispersal in which negative feedback is present. This research thus highlights a set of ecological conditions sufficient for the evolution of one form of altruism.

Keywords: bacteriocin, ecological feedback, experimental evolution, positive assortment, survival of the weakest

Wisely and slow. They stumble that run fast.

William Shakespeare

The conflict between individual and group interests is a common element in many social dilemmas. Consider the rate at which an organism consumes shared resources. Prudent use of common resources promotes the longevity or fecundity of the group; however, any individual that exhibits restraint suffers in competition with those utilizing resources rapidly. Rapacity is selectively favored and the displacement of prudent types by their unrestrained contemporaries occurs despite harmful consequences for the group (2, 3). Restraint in the use of common resources is a form of altruism: behavior that is self-sacrificial and pro-social. Like other types of altruistic behavior, restraint faces a fundamental problem of subversion (4, 5). How can restrained types persist in the midst of would-be cheaters—individuals that have a competitive edge *because* they are unrestrained? In this article, we address this question directly by outlining ecological conditions sufficient to favor the evolution of restraint.

One ingredient found in most explanations for the evolution of altruism, and thus relevant to the evolution of restraint, is *positive assortment*. Altruism stands a better chance when altruistic individuals disproportionately help those possessing the genes for altruism (6–10). One of the most obvious ways to achieve positive assortment is through interactions between genetic relatives (11). In such a case, altruistic individuals disproportionately experience beneficial social environments

(engineered by their kin); whereas selfish individuals tend to face a milieu lacking pro-social behavior (as their kin tend to be less altruistic). Interaction with kin can occur actively, via the choice of relatives as social contacts, or passively, via the interaction with neighbors in a habitat with limited dispersal. There is now a large body of literature on the effect of active and passive assortment on the evolution of altruism (6, 10, 12–18). At a fundamental level, this research focuses on the distribution of *interactions* among altruistic and selfish individuals. However, in many systems, these individuals are also interacting with other members of their community (e.g., competing species, predators, prey, mutualists, etc.). It is less common to consider the role of broader ecological interactions on the evolution of various forms of altruism.

Here we consider the evolution of restraint in communities where ecological interactions generate a type of negative feedback. One of the simplest communities with this property involves three members engaged in non-transitive competition reminiscent of the children’s game “rock-paper-scissors.” In this game, each strategy beats one of the other two and is beaten by the third (e.g., paper covers rock, but is cut by scissors). Imagine a community with three competitors forming a non-transitivity—a situation in which no player is competitively superior to all others. For convenience, call the players Rock, Paper and Scissors. Each type has a rate at which it displaces its victim (e.g., Rocks “crush” Scissors at some rate). Next, imagine a less restrained variant of Rock, call it Rock*, that displaces Scissors at a faster rate. In a Rock*-Paper-Scissors community, the abundance of Scissors

decreases due to the increased prowess of Rock*. As a consequence, Scissors' victim (Paper) is liberated, which can displace Rock*. In an ironic twist, the "improved" Rock* decreases in abundance due to the expansion of its victim's victim. This form of negative feedback ensures that a higher displacement rate results in *decreased* abundance (19–22). Thus, more restrained players may be less prone to extinction, a phenomenon termed 'survival of the weakest' (19). A complication arises when considering a community with multiple variants present simultaneously (e.g., Rock *and* Rock* with Paper and Scissors). The same traits that allow Rock* to displace Scissors faster may render Rock* a better competitor against Rock. In this case, restraint has a selective disadvantage despite its positive effects on abundance. How then can restraint evolve in a non-transitive community?

Spatial structure can play a critical role promoting restraint in non-transitive systems. Returning to our Rock-Paper-Scissors community, limitation of dispersal results in a patchwork of the three players. A patch of any one player chases its victim and is chased by its enemy (23, 24). Within any patch, an unrestrained variant (Rock*) will replace its restrained counterpart (Rock). However, patches of unrestrained variants are more likely to go extinct. This difference in patch viability favors restraint. Limited dispersal ensures a type of positive assortment where restrained and unrestrained individuals tend to be surrounded by like types. This means that the long-term negative consequences of faster displacement are visited disproportionately on the less restrained type. Consequently, restraint can be maintained evolutionarily in a structured non-transitive community. This outcome

has been shown theoretically (25, 26), but there is little empirical work on this topic. This is despite the fact that non-transitive dynamics have been described in natural communities ranging from microbes to animals to plants (27–32).

One well-studied non-transitive system involves strains of *Escherichia coli* that produce antimicrobial proteins termed colicins (2, 3). Colicin-producing cells possess a plasmid housing the colicin gene, as well as a gene coding for a colicin-specific immunity protein. Cells that lack the plasmid, and thus lack immunity, are sensitive to the colicin. However, sensitive cells can experience mutations yielding resistance to colicins. Resistance is due to alteration or loss of membrane proteins that bind or translocate the colicin. As these same membrane components are involved in nutrient acquisition, resistance is often costly in the absence of colicins (measured by a reduced growth rate relative to sensitive cells) (33, 34). However, in some cases, the producer incurs even greater costs to carry the colicin plasmid and express immunity constitutively. Thus, these three players constitute a non-transitive community: the sensitive strain outgrows the resistant strain, the resistant strain outgrows the producer, and the producer kills the sensitive strain. Previous work with the three members of the colicin E2 system has demonstrated non-transitivity both *in vitro* (23) and *in vivo* (35). Nevertheless, there have been no experimental studies of the evolution of restraint in this system.

In this article we describe experiments with bacteria that explore how positive assortment and negative ecological feedback influence the evolution of restraint. Of

the three players (sensitive, resistant, and producer), we focus on the resistant strain. The mutations that define the resistant strain are costly and there is evidence from numerous systems that secondary mutations can compensate for the initial costs of antimicrobial resistance (36–40). Thus, we predict that this strain is the most likely to *increase* its growth rate, making it the most attractive candidate to study factors that would hinder such increase. We place the community in a metapopulation, structured into many subpopulations. We manipulate the pattern of migration within the metapopulation, which affects the degree of positive assortment. Migrations are either restricted to occur between neighboring subpopulations (*Restricted* treatment) or could occur between any subpopulations (*Unrestricted* treatment). The evolution of the resistant strain can be compared across migration treatments to gauge the effect of population structure on the evolution of restraint. To identify the role of negative feedback, the evolution of the resistant strain in the full community is compared to the evolution of the resistant strain evolving *alone* (*Community* and *Alone* treatments, respectively). By monitoring the resistant strain in three different types of metapopulations (*Restricted Community*, *Unrestricted Community*, and *Restricted Alone*), we assess the impact of both positive assortment and negative ecological feedback on the evolution of restraint.

Results

Presence of non-transitivity

As detailed in the *Methods*, we constructed a strain that produced two colicins (Producer), a strain sensitive to both colicins (Sensitive) and a strain resistant to both colicins (Resistant). The double colicin producer was used to decrease the likelihood of *de novo* resistance arising from the sensitive population during the evolution experiment. These three constructed strains are henceforth referred to as the “ancestors.” To confirm the non-transitive relationship, we performed pairwise competitions among the ancestral strains. Each competition was initiated with a ratio matching the proportions of two competitors when they first meet via migration within the metapopulation. The resistant ancestor was out-competed by the sensitive ancestor (one-sample t test; $t_5 = -5.78$, $p = 0.0022$). The producer ancestor was outgrown by the resistant ancestor (one-sample t test; $t_5 = -3.62$, $p = 0.015$). The sensitive ancestor was always driven to extinction when mixed with the producer (giving a relative fitness of zero in all five replicates). As each player was competitively inferior to a second player (but superior to the third) these three strains form a non-transitive system (Figure 1).

Ecological dynamics

We propagated our bacteria as metapopulations using 96-well microtiter plates, where each well constituted a distinct subpopulation. We initialized the metapopulations with the non-transitive community (*Community* treatment) or the resistant strain alone (*Alone* treatment). Every 12 hours, all subpopulations were diluted into fresh growth medium and migrations between subpopulations occurred. Within each metapopulation, migrations occurred between neighboring

wells (*Restricted* treatment) or among any wells (*Unrestricted* treatment). We measured the abundances of all strains every six transfers. All three players were maintained in the *Restricted Community* and *Unrestricted Community* treatments for the duration of the experiment (Figure 2a,b). The resistant strain persisted at a constant level in the *Restricted Alone* treatment for the length of the experiment (Figure 2c).

Evolution of the resistant strain

We randomly sampled eight resistant isolates from the last transfer of the experiment. Each of these isolates was competed against a marked variant of the common resistant ancestor. To avoid pseudoreplication, we averaged relative fitness across isolates within each of five replicates of each treatment. We found that isolates from the *Restricted Community* treatment had the lowest competitive ability (single factor ANOVA; $F_{2,12}=9.36$, $p=0.0036$, multiple comparisons by Tukey's HSD). This is consistent with the evolution of a restrained growth rate. Resistant cells in a full community evolved a significantly higher competitive ability under unrestricted migration than under restricted migration (*Unrestricted Community* versus *Restricted Community* in Figure 3). Resistant cells propagated alone evolved a significantly higher competitive ability than resistant cells in a non-transitive community (*Restricted Alone* versus *Restricted Community* in Figure 3). Thus, both population structure and the presence of the full community were important to the evolution of competitive restraint.

Simulation of eco-evolutionary dynamics

To better understand the evolutionary behavior of our system, we modeled the bacterial metapopulations using a lattice-based simulation (see *Methods* and *SI Methods* for details). Each metapopulation was initialized with the three ancestral strains in a spatially clumped pattern. The basic algorithm consisted of a cycle of three stages: (i) growth/competition within wells, (ii) dilution of wells, (iii) migration among wells. Thus, a simulated cycle corresponds to a transfer within our experiment. Every cycle, mutations to growth rate were permitted in resistant subpopulations. We simulated evolution within metapopulations in each of the three treatments described above (*Restricted Community*, *Unrestricted Community* and *Restricted Alone*).

While diversity was maintained in the *Restricted Community* treatment, the community tended to lose players in the *Unrestricted Community* treatment in the long run (e.g., after 100 transfers). Consequently, the *Unrestricted Community* treatment was excluded from analysis. The loss of diversity was robust to changes in several different parameters of the model and suggests that the *Unrestricted Community* treatment in the laboratory may have lost strains had it been run for more transfers. This result is also consistent with previous work on the importance of limited dispersal to coexistence in this system (23, 24). After evolving the metapopulations in each treatment, we determined the mean relative fitness of the resistant population. Consistent with our empirical results, we found the average

growth rate of resistant strains from the *Restricted Community* treatment to be significantly lower than the average growth rate from the *Restricted Alone* treatment (see *SI Figure 3*).

To confirm the importance of positive assortment in the evolution of restraint, we ran an additional treatment: *Restricted Community with Permutation*. This treatment was identical to the *Restricted Community* treatment, except that at the beginning of each cycle, wells containing only resistant cells (ancestor or mutants) were randomly permuted. This operation allowed for mixing between the patches of resistant wells (capturing an element from the *Unrestricted* treatment). The average growth rate of resistant strains from the *Restricted Community* treatment was significantly lower than the average growth rate from the *Restricted Community with Permutation* treatment (Figure 4).

The rate of displacement by fitter variants within any population will be slowed by population subdivision. We were curious if the lower growth rate of our *Restricted Community* treatment could be explained entirely by the fact that the evolving resistant population was divided into semi-isolated patches. To explore this possibility, we ran an additional simulation treatment: *Restricted Alone with Shadowing*. In this treatment, a *Restricted Alone* metapopulation evolved alongside a standard *Restricted Community* metapopulation, with the caveat that the *Restricted Alone* metapopulation's migrations and spatial distribution was forced to match the resistant portion of its paired *Restricted Community* metapopulation. In this way,

the *Restricted Alone* “shadowed” the *Restricted Community*. This meant that the *Restricted Alone* metapopulation was divided into patches. However, because mutation occurred independently in the *Restricted Alone* shadow and its *Restricted Community* master, mutations within a given patch in the shadow world had no effect on the survival of the patch in that world. We found that division into semi-isolated patches accounted for some, but not all, of the effect of lowering growth rate in the short term (Figure 4a, single factor ANOVA; $F_{2,217}=58.31$, $p<0.001$, multiple comparisons by Tukey’s HSD). However, simulations that ran for longer (Figure 4b) show that the *Shadowing* treatment converges to the *Permutation* treatment (single factor ANOVA; $F_{2,217}=93.76$, $p<0.001$, multiple comparisons by Tukey’s HSD). We find the same patterns when we run simulations that exactly match the metapopulation size and number of transfers used in our experiment (see *SI Figure 5*). Thus, apparently the connection between the presence of fast growing variants within a patch and a greater probability of patch extinction was an important ingredient in explaining the evolution of restraint in the *Restricted Community* treatment.

Discussion

For the resistant isolates considered here, the evolution of the lowest competitive ability occurred in the treatment in which migration was restricted and all three members of the non-transitive community were present (Figure 3). If either migration was unrestricted or the resistant strain evolved alone, final competitive ability was significantly higher. The low competitive ability in the *Restricted*

Community treatment presumably reflects a relatively low growth rate. There are a few possible explanations for this outcome. First, if the number of resistant cell divisions in the *Restricted Community* treatment was less than the number of divisions in the other treatments, isolates from the *Restricted Community* treatment might not have enough opportunity to evolve a higher growth rate. However, we find no significant difference among the treatments in the total number of resistant cell divisions (*SI Methods* and *SI Figure 6*). A second explanation is that restricted migration slows the spread of any advantageous mutant (41). In this case, resistant mutants with a higher growth rate reach a lower frequency in the *Restricted Community* treatment than in the *Unrestricted Community* treatment by the end of the experiment. However, the resistant isolates with the *highest* growth rate came from *Restricted Alone* treatment: thus, a restriction to migration does not uniformly hinder the advent of fast-growing resistant mutants. A third explanation is that the presence of producers constrains the manner in which a resistant strain can compensate for the cost of resistance (e.g., reversion to sensitivity is not an option). This would limit the set of evolutionary options for resistant cells in the *Restricted Community* treatment relative to the *Restricted Alone* treatment. However, the growth rate of isolates from the treatment with the highest level of interaction between resistant cells and producers (*Unrestricted Community*) was similar to that of the treatment without producers (*Restricted Alone*). Additionally, not a single resistant isolate from any treatment reverted to sensitivity; thus, reversion did not explain competitive differences. Lastly, the *Restricted Community* treatment's resistant population was divided into discontinuous regions by barriers consisting

of the other strains (illustrated in Figure 5), and such barriers would inhibit the spread of advantageous mutants. Our simulation-based treatment, *Restricted Alone with Shadowing*, where the resistant type was restricted to the patchy spatial distribution of *Restricted Community*, evolved a lower growth rate, indicating that a population subdivision may contribute to the low growth rate in the *Restricted Community*. Nonetheless, subdivision does not fully account for the restraint found in the *Restricted Community* treatment (see Figure 4 and *SI Figure 5*). Thus, we do not find complete support for any of these explanations and instead favor the following alternative.

In the *Restricted Community* treatment, the non-transitivity of the full community provides a form of negative feedback and the restricted migration ensures a form of positive assortment. We suggest that it is these two factors, negative feedback and positive assortment, that set the stage for the evolution of restraint. In the *Restricted Community* treatment, we have a set of patches chasing one another (see Figure 5). A faster growing resistant mutant has a competitive advantage within a resistant patch, but a fast-growing resistant patch is more likely to burn through its victim (the producer) and consequently face its enemy (the sensitive strain). This sequence of events is shown in Figure 5 for a *Restricted Community* simulation in which wells with a faster growing resistant mutant were labeled black. Limited migration ensures that it is the unrestrained mutants that reap the negative long-term consequences (patch extinction) of their myopic strategy. When assortment is eradicated by shuffling the contents of multiple patches (as in the simulation-based treatment *Restricted Community with Permutation*) restraint is not maintained

(Figure 4). Without the negative feedback of the full community (e.g., in the *Restricted Alone* treatment) or without the positive assortment resulting from limited migration (e.g., in the *Unrestricted Community* treatment), the evolution of restraint is not expected.

We have explored a bacterial system under laboratory conditions, but our findings carry potential implications for other systems. Non-transitive relationships have been described in a diverse set of organisms, including yeast (28), plants (32), coral reef invertebrates (27) and lizards (29). At the moment, it is not clear whether non-transitivities are common in natural ecosystems (42). In contrast, the ubiquity of spatial structure is widely recognized. Structure may be most pronounced in sessile organisms (e.g., plants, some marine invertebrates, microbes in biofilms); however, even populations of motile organisms can possess some degree of structure due to spatial limitations to dispersal and interaction. The spatial scale of ecological processes has been shown to be an important factor in the invasion of rare types (43, 44), coexistence of multiple types (45), the stability of communities (46), and the evolutionary trajectory of community members (47). We have shown that limited migration in a non-transitive community can promote the evolution of restraint. However, spatial structure can be important for the evolution of restraint in other types of communities as well.

As one example, limited dispersal can promote restraint within victim-exploiter communities (48, 49). An inherent form of negative feedback exists when one species (e.g., predator, parasite, herbivore) exploits another for critical resources

(e.g., prey, host, plant). To see this, consider a simple version of the Lotka-Volterra model where the dynamics of exploiters (at density E) and victims (at density V) are described by:

$$\begin{aligned}\frac{dV}{dt} &= \beta V - \lambda VE, \\ \frac{dE}{dt} &= \lambda VE - \delta E,\end{aligned}$$

where b is the birth rate of victims, l measures the attack rate of the exploiter, and d is the death rate of the exploiter. The non-trivial equilibrium for this community is:

$$\left(\hat{V}, \hat{E}\right) = \left(\delta/\lambda, \beta/\lambda\right)$$

As the exploiter reduces its attack rate, its equilibrium abundance increases (as l drops, $\hat{E} = \beta/\lambda$ grows). Nonetheless, an exploiter with a higher attack rate will displace a second exploiter exercising restraint (50). Selection for rapacious exploitation that results in community collapse constitutes an example of the “tragedy of the commons” (51). Limited dispersal ensures that any tragedy of the commons that results from overexploitation befalls primarily the unrestrained exploiters. Several theoretical studies have explored the role of spatial structure in promoting restraint in victim-exploiter interactions (52, 53). There have also been experimental demonstrations that limited dispersal favors restraint in host-parasite communities, in the form of reduced parasite virulence and/or infectivity (48, 49, 54).

A second example involves the role of structure in promoting restraint in hypercycle “communities.” A hypercycle is a series of self-replicative molecules that are

cyclically linked, where each molecule catalyzes the replication of the next molecule in the cycle. Unstructured hypercycles are plagued by 'parasitic' molecules, which receive greater catalytic activity from the previous molecule in the cycle while withholding catalytic support for the next molecule in the cycle. Boerlijst and Hogeweg (55) showed theoretically that hypercycles in an incompletely mixed medium could keep parasitic molecules at bay. In a structured habitat, the hypercycle community organizes into a collage of rotating spirals. A parasitic molecule originating at the center of a spiral can lead to spiral demise and replacement by other spirals. Thus, short-term payoffs to the parasite (displacement within a spiral) can generate negative long-term consequences (spiral extinction) in a structured world. This favors the evolution of "restrained" molecules that avoid the immediate gains of parasitism.

Spatial structure and ecological feedback can also favor mutualistic behavior between species (56). Recently, Harcombe (57) studied a case of bacterial cross-feeding. In lactose medium, *Salmonella enterica* consumes the acetate waste products of a mutant strain of *E. coli*. The *E. coli* mutant was a methionine auxotroph and could grow if *S. enterica* excreted methionine. Harcombe demonstrated that even though methionine excretion was intrinsically costly, a mutant of *S. enterica* that exported an excess of methionine was able to displace wild-type *S. enterica* (which did not excrete methionine) when these types were grown on lactose plates with *E. coli*. The cooperative excretion by *S. enterica* was favored through a combination of ecological feedback (acetate was produced when

E. coli obtained methionine) and spatial structure (ensuring that excreting cells had disproportionate access to acetate). When Harcombe destroyed either the feedback (by growing the community on acetate plates such that *S. enterica* did not rely on *E. coli*) or the structure (by growing the community in lactose flasks), the excreting *S. enterica* mutant was outcompeted by the wild-type. This work demonstrates that ecological feedback and positive assortment can be important ingredients in other forms of cooperation.

In all the communities described above, a form of altruism exists. The elements that we have underlined as important to the evolution of restraint connect readily to prominent theoretical frameworks used to understand the evolution of altruism. In our non-transitive system, limited dispersal results in a preponderance of interaction between relatives. Kin selection arguments often focus on the coefficient of relatedness between interacting individuals (58, 59). In our system, limited dispersal results in higher coefficients of relatedness than under conditions of unlimited dispersal, a form of positive assortment (60). The multilevel selection framework describes altruism as a behavior opposed by within-group selection, but favored by between-group selection (61, 62). In the patchwork of a structured community, a restrained variant is at a local disadvantage (e.g., within its patch), but patches of restrained types may persist longer due to the negative feedback from rapid growth. We propose that multiple frameworks have relevance for understanding restraint in our system because each framework focuses on (different) important elements underlying the evolution of altruism (63).

Overall, we observe that a form of altruism can evolve in microbial metacommunities. With limited migration, similar types associate into patches that chase one another. The negative feedback resulting from the non-transitivity in our system means that patches filled with unrestrained variants are more prone to extinction. Thus, we see that altruistic restraint is favored precisely when those that “run fast” tend to “stumble.”

Methods

Community players

The bacterial community consisted of three players: a toxin-producing strain (P), a toxin-sensitive strain (S), and a toxin-resistant strain (R). The producer expressed two toxins (colicin E2 and colicin D). This strain was constructed by transforming the colicin D plasmid (Col D) into chemically competent BK2 cells (that already possessed Col E2). This double producer was then marked with resistance to phage T5. Strain BK10 (*E. coli* K-12, strain BZB1011) is identical to BK2, except that the former lacks the Col E2 plasmid. The sensitive strain was constructed by transforming the pACYC184 plasmid into competent BK10 cells. This transformation accomplished two things: (i) the sensitive player had a positive selectable marker (resistance to tetracycline) and (ii) the competitive differences between the sensitive strain and other strains could be amplified by adding a low concentration tetracycline to the media. Finally, the resistant strain was constructed by a series of sequential selections: first, selecting BK10 that was

resistant to colicin E2, then selecting for resistance to colicin D, then selecting for resistance to phage T6. We note that some of these markers are known to be costly (e.g., T6 resistance). However, use of these costly markers was intentional: we were attempting to construct a community with a pronounced rock-paper-scissors dynamic. We note that even before the transformations and marker addition, these strains exhibit a rock-paper-scissors relationship (sensitive outgrows resistant, which outgrows producer, which kills sensitive). However, through marker addition and medium manipulation (addition of a low concentration of tetracycline), the growth rate differences were magnified leading to an intensified non-transitive relationship. For this reason, in both the evolution experiment and assays, the growth medium was lysogeny broth (LB-Miller) supplemented with 0.25mg/mL tetracycline.

Experimental treatments

The evolution experiments involved propagating metapopulations of bacteria with two factors manipulated. The first experimental factor was the identity of the players in the metapopulation. Either the full community (S-R-P) was used or the resistant strain (R) was propagated alone (the *Community* or *Alone* treatments, respectively). In the *Community* treatments, each metapopulation consisted of two microtiter plates (192 wells, each with 200 μ l growth medium). In the *Alone* treatments, each metapopulation consisted of a single microtiter plate (96 wells, each with 200 μ l of growth medium). The difference in the number of wells reflected our attempt to balance the total number of resistant cells across treatments (see

Figure 2). The second factor manipulated was the pattern of migration within the metapopulations. Migration was either restricted to occur between wells directly bordering each other along cardinal directions or was unrestricted (the *Restricted* or *Unrestricted* treatments, respectively). In both treatments, every well had 1/3 probability of experiencing an immigration event from one random well in its neighborhood. In the *Restricted* treatment, this neighborhood included the wells directly north, east, south or west of the focal well (using wrap-around boundaries to eliminate edge-effects). In the *Unrestricted* treatment, the neighborhood included all wells minus the focal. Migration events directly followed dilution of the entire metapopulation into fresh growth medium. Every 12 hours, 40-fold dilution was accomplished using a 96 slot-pin multi-blot replicator (5µl into 200µl). Immediately following dilution, a BioRobot 8000 liquid-handling robot (Qiagen) executed the migrations, where each migration involved transferring 5µl from the source well within the exhausted plate into the destination well within the fresh plate. Between transfers, plates were incubated (37°C) and shaken (350 rpm using a microtiter shaker). For the *Alone* treatment, the metapopulation was initiated with the resistant strain in every well. For the *Community* treatment, the initial spatial arrangement of strains was obtained from the 100th transfer of a 192-point lattice-based simulation with a restricted neighborhood (see *SI Methods*). Each metapopulation was propagated for a total of 36 transfers. The abundance of each strain was gauged every 6 transfers by selective plating (using tetracycline, phage T5, and phage T6). There were five replicates of each of three treatments: (1) *Restricted Community*, (2) *Unrestricted Community*, and (3) *Restricted Alone*.

Competition assay

We picked eight random resistant isolates from the last transfer of each metapopulation (here we denote any one of these strains as R_E). We marked our ancestral resistant strain with resistance to phage T5 (we denote this marked ancestor as R_A). Before the competition, R_E and R_A are grown separately in 200 μ l of growth medium for two 12-hour cycles (with 40-fold dilution at transfer). After this acclimation phase, we added 5 μ l of R_E and 5 μ l of R_A to a well containing 200 μ l fresh growth medium. The titer of each strain was assessed (through selective plating with and without phage T5) immediately after the competition was initiated and again after 12 hours. If $R_i(t)$ is the titer of strain R_i at time t , then the fitness of the evolved strain relative to its ancestor is given by:

$$w(R_E, R_A) = \frac{\ln(R_E(12)/R_E(0))}{\ln(R_A(12)/R_A(0))}$$

The same competitive assay was used to establish the non-transitive dynamic between the three ancestral players (simply with different selective plating schemes).

Simulation

We model the metapopulation as an $L \times W$ regular square lattice with periodic boundaries subjected to a cycle of three phases: (i) growth, (ii) dilution, and (iii) migration. Each lattice point i at time t is described by the vector:

$$x_i(t) = \langle s_i(t), p_i(t), r_i^0(t), r_i^1(t), \dots, r_i^K(t) \rangle,$$

where $s_i(t)$, $r_i^0(t)$, and $p_i(t)$ are the abundances of sensitive, resistant and producer ancestors, respectively. The variables $r_i^1(t), r_i^2(t), \dots, r_i^K(t)$ are the abundances of each of K types of mutant resistant strains. These abundances are expressed in units of the limiting nutrient concentration (see *SI Methods*).

During the growth phase, the dynamics of each strain (y) of each lattice point is described by the following differential equation (see *SI Methods*):

$$\dot{y} = y_i G_Y(n_i)$$

where $n_i = 1 - \sum y_i$ and we assume a Monod growth curve:

$$G_Y(n_i) = \frac{\mu_Y n_i}{\kappa_Y + n_i}.$$

The parameters m_Y and k_Y give the maximum growth rate and Monod constant (nutrient concentration yielding $\frac{1}{2}$ maximum growth rate) of player Y . Each growth phase lasts T time units.

Dilution at time t is given by:

$$x_i(t') = \phi x_i(t),$$

where ϕ is the dilution factor and t' marks the post-dilution state.

Migration happens with a uniform probability a . If a migration event occurs, a point within the focal point's neighborhood is chosen at random. For the *Restricted* treatment, the neighborhood is the four nearest lattice points (the von Neumann neighborhood). For the *Unrestricted* treatment, the neighborhood is the entire lattice minus the focal point. In the case of migration, let the chosen neighbor

of the focal point i be designated j . The state following migration (signified by t'') is given by

$$x_i(t'') = (1 - \phi)x_i(t') + \phi x_j(t).$$

Removal occurs next. At point i , any player whose abundance is less than or equal to a critical value (a_{crit}) is removed. Also, the sensitive player is removed if the producer is present. A removed player has its abundance set to zero. In the simulation, the dilution, migration and removal are assumed to be instantaneous and followed by a new growth phase. Lastly mutation can occur with probability p . In the case of a mutational event, a fraction g of the total abundance of the resistant players (ancestral and mutant) of a point is converted to a random resistant type.

We initialize lattice point i with the starting abundances of each ancestral player ($s_i(0)$, $r_i^0(0)$, and $p_i(0)$) using the same method as in the bacterial experiment (see *SI Methods*). After C growth cycles, we measured the average fitness of the resistant mutants (the R^k strains) in competition with the resistant ancestor (R^0). Letting R^* represent all resistant strains in an evolved metapopulation, the mean relative resistant fitness is:

$$\bar{w}(R^*, R^0) = \sum_{k=0}^K f(k) w(R^k, R^0),$$

where

$$f(k) = \frac{\sum_{i=1}^{L \times W} r_i^k(CT)}{\sum_{k=0}^K \sum_{i=1}^{L \times W} r_i^k(CT)},$$

and

$$w(R^k, R^0) = \frac{\ln(r^k(T) / r^k(0))}{\ln(r^0(T) / r^0(0))}.$$

In Table 1, we give the values for all the simulation parameters, which are tailored to our experimental conditions or estimated from assays (e.g., see *SI Methods*). For Figure 4, Figure 5, and *SI Figure 3*, we assume $L=100$, $W=100$. For *SI Figure 4* and *SI Figure 5* (small lattice simulations) we assume $L=16$, $W=12$, and $C=36$, the values corresponding to our laboratory experiment.

Acknowledgements

We thank P. Aulakh, C. Eshelman, H. Lindsey, S. Smith, and J. West for help in the laboratory; B. Bohannan, S. Forde and M. Riley for discussion surrounding the early conception of this work; and M. Clifford, J. Cooper, S. Drescher, J. Gallie, C. Glenney, S. Heilmann, H. Lindsey, B. Miner, and J. West for comments on a previous draft of this manuscript. BK thanks J. Avise, F. Ayala, D. Queller, and J. Strassmann for the invitation to participate in the Sackler colloquium “In the Light of Evolution V: Cooperation”. This material is based in part upon work supported by the National Institutes of Health (07-004309 B-03 AM03) and work supported by the National Science Foundation under Cooperative Agreement No. DBI-0939454, a CAREER award to BK (DEB0952825), and a GRF to JN. Any opinions and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the National Institutes of Health.

Chapter 1 Figures




Figure 1: Pair-wise competitions between the ancestral bacterial strains demonstrate non-transitivity. The asterisks signify that relative fitness is significantly less than one and the bars give the standard error of the mean. The resistant ancestor is dominated by the sensitive ancestor and the ancestral producer is outgrown by the resistant ancestor. The sensitive strain is killed by producer in all replicates, yielding a uniform relative fitness of zero. As each strain out-competes one other strain, but is out-competed by the third strain, a non-transitive relation holds.

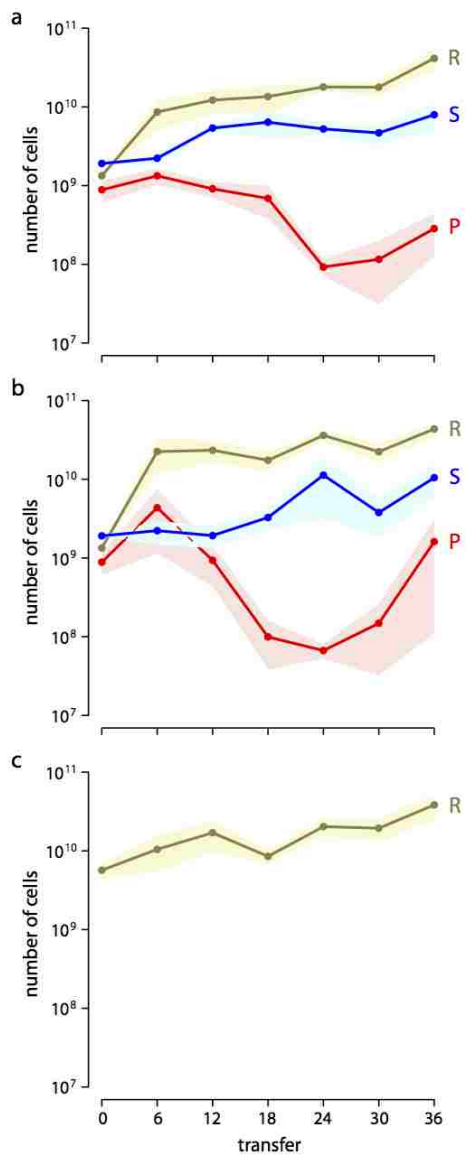


Figure 2: Bacterial abundance in (a) the *Restricted Community* treatment, (b) the *Unrestricted Community* treatment, and (c) the *Restricted Alone* treatment. Points represent mean abundance of the sensitive strain (blue, S), resistant strain (yellow, R) and producer strain (red, P). Shading gives the standard error of the mean. All three players coexisted in the *Community* treatments for the duration of the

experiment and the density of the resistant strain was comparable across all three treatments.




Figure 3: The fitness of evolved resistant isolates relative to their common ancestor. Mean relative fitness of each treatment is shown and bars give the standard error of the mean. The fitness of isolates from the *Restricted Community* treatment was significantly lower than the fitness of isolates from the other treatments. Letters distinguish treatments significantly different using post hoc comparisons. This pattern is consistent with the evolution of restrained growth in the *Restricted Community* treatment.

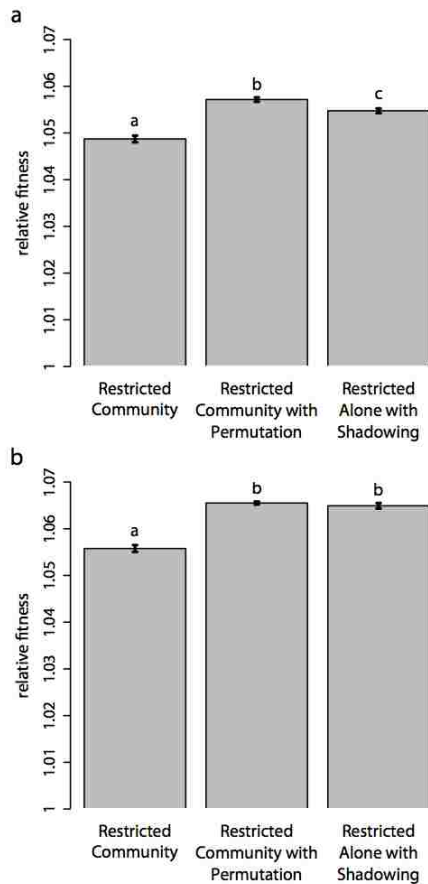


Figure 4: The mean resistant fitness relative to the resistant ancestor after simulated evolution in multiple treatments. Fitness values after (a) 100 and (b) 200 cycles are shown. Mean relative fitness of each treatment is shown and bars give the standard error of the mean. Letters distinguish significantly different treatments by post hoc comparisons. The fitness of resistant populations from the *Restricted Community* treatment was significantly lower than that of the other treatments at both time points. This pattern is consistent with the evolution of restrained growth in the *Restricted Community* treatment.




Figure 5: Snapshots of a metapopulation from an illustrative *Restricted Community* simulation recorded every 20 cycles (a-h). The metapopulation was initialized with the three bacterial strains: sensitive (blue), resistant (yellow) and producer (red), in addition to a small patch of a mutant resistant strain (black) with an increased growth rate. The mutant initially outcompetes nearby ancestor patches (a-e), but is extinguished after outcompeting neighboring patches of the producer (f-h).

SUPPORTING INFORMATION:

METHODS

Community Model. Within a single population, we track the densities of sensitive cells (S), resistant cells (R), producers (P), and mutant resistant cells (M) (we consider only a single mutant class here, but the model easily extends to cover an arbitrary number of mutant classes). Below we measure density by absorbance in a spectrophotometer, which has a linear relation to cell abundance. For convenience we refer to the density of the player as well as its type by an italicized capital letter. In the first version of our model, we also track the concentration of nutrients (n).

The system is described by the following set of ordinary differential equations:

$$\begin{aligned}\dot{S} &= SG_S(n), \\ \dot{R} &= RG_R(n), \\ \dot{P} &= PG_P(n), \\ \dot{M} &= MG_M(n), \\ \dot{n} &= -\varepsilon_S SG_S(n) - \varepsilon_R RG_R(n) - \varepsilon_P PG_P(n) - \varepsilon_M MG_M(n),\end{aligned}$$

where ε_Y is the amount of nutrients needed to shift the absorbance of strain Y by a single unit and $G_Y(n)$ is the growth rate of strain Y . We use a change of variables, where, for each player Y :

$$\varepsilon_Y Y = y,$$

such that bacterial density is expressed in terms of nutrient concentration. Thus, we have

$$\begin{aligned}
\dot{s} &= sG_S(n), \\
\dot{r} &= rG_R(n), \\
\dot{p} &= pG_P(n), \\
\dot{m} &= mG_M(n), \\
\dot{n} &= -mG_M(n) - sG_S(n) - rG_R(n) - pG_P(n).
\end{aligned}$$

For convenience, the community is initialized with:

$$s(0) + r(0) + p(0) + m(0) + n(0) = 1.$$

Since

$$\frac{d(s + r + p + m + n)}{dt} = 0,$$

we know that $s(t) + r(t) + p(t) + m(t) + n(t) = 1$ for all t . Thus, we can rewrite the

original system of five differential equations as a system of four:

$$\begin{aligned}
\dot{s} &= sG_S(1 - m - s - r - p), \\
\dot{r} &= rG_R(1 - m - s - r - p), \\
\dot{p} &= pG_P(1 - m - s - r - p), \\
\dot{m} &= mG_M(1 - m - s - r - p).
\end{aligned}$$

Growth Parameter Estimation. In the previous section, the growth rate (G_Y) of strain Y is a function of limiting nutrient concentration (n). A simple way to assess this function would be to measure growth rate at different nutrient concentrations and then fit a curve to yield $G(n)$. Our experiment was conducted in a rich medium (LB); therefore, there was no clear single nutrient to vary. Thus, we took an approximate approach. We substituted fractions of the growth medium (LB) with saline (0.86% NaCl), while maintaining the concentration of tetracycline at a constant level. Each bacterial strain was grown in a microtiter well with 200ml of a


given concentration of impoverished medium for a full 12 hours and then diluted 4-fold into 200ml of medium of the same concentration and grown for 1 hour. This actively growing bacterial culture was then diluted 4-fold into 200ml of medium of the same concentration and incubated in a spectrophotometer (VersaMax, Molecular Devices). Absorbance (600nm) was measured at 2-minute intervals for an hour. Let $A_{Y,f}(t)$ be the absorbance of strain Y in impoverished medium with a fraction f of LB (and $1-f$ saline) at time t . Let $G_{Y,f}$ be the slope of the least-squares line $\ln(A_{Y,f}) = G_{Y,f}t + b$; thus, $G_{Y,f}$ is the Malthusian growth parameter corresponding to exponential growth. We used the data from time point 1/15 to 7/10 (in hours) to estimate $G_{Y,f}$. For a given strain, we measured $G_{Y,f}$ at a number of different f values (i.e., different concentrations of LB). Using the Monod growth model, we computed the parameters for the least-squares curve:

$$G_{Y,f} = \frac{\mu_Y f}{\kappa_Y + f},$$

where m_Y is the maximal growth rate and k_Y is the fraction of LB necessary to grow at half the maximal rate for strain Y . We first determined the least-squares value for m_S and k_S . We then used m_S to constrain the m_R value ($m_R \leq m_S$) and the least-squares m_R value to constrain the m_P value ($m_P \leq m_S$).

Initializing the Metapopulation. Each community's starting spatial layout of the three ancestral strains was determined by running a lattice-based model, in which each lattice point corresponded to one well (subpopulation) within the metapopulation. Using preliminary competition data of the ancestral strains, SI

Figure 1 shows the number of dilution/growth cycles needed for one strain to displace another.



SI Figure 1: The frequencies of two competing strains are given as pie charts. Sensitive cells (S, blue) displace resistant cells (R, yellow) over 6 dilution/growth cycles. Resistant cells displace producers (P, red) over 4 dilution/growth cycles. Producers kill sensitive cells within a single dilution/growth cycle. The letters and subscript under each pie chart give each state the name used in the state transition matrix.

If we record the state of a well (lattice point) directly after an incubation period, we can describe community dynamics (approximately) by using the following discrete state set:

$$\{S, R, P, SR_1, SR_2, SR_3, SR_4, SR_5, RP_1, RP_2, RP_3\}$$

For a well that does not experience an immigration event, the following transitions occur for distinct strains X and Y over a dilution/growth cycle:

$$\begin{aligned}
& X \rightarrow X, \\
XY_i \rightarrow & \begin{cases} X & \text{if } i \text{ is its maximum value,} \\ XY_{i+1} & \text{otherwise.} \end{cases} \quad [S1]
\end{aligned}$$

When migration occurs, the contents of wells in two (potentially different) states are combined. Each entry in SI Table 1 gives the final state for a well that starts in the row state, is diluted, experiences an immigration event from a well in the column state, and then grows for one period.

SI Table 1: Transition Matrix

	S	R	P	SR₁	SR₂	SR₃	SR₄	SR₅	RP₁	RP₂	RP₃
S	S	SR ₁	P	SR ₂	SR ₃	SR ₄	SR ₅	S	RP ₂	RP ₃	R
R	SR ₁	R	RP ₁	SR ₁	SR ₁	SR ₁	SR ₁	SR ₁	RP ₂	RP ₃	R
P	P	RP ₁	P	RP ₁	RP ₁	RP ₁	RP ₁	RP ₁	RP ₁	RP ₁	RP ₁
SR₁	SR ₂	SR ₁	RP ₁	SR ₂	SR ₂	SR ₂	SR ₂	SR ₂	RP ₂	RP ₃	R
SR₂	SR ₃	SR ₁	RP ₁	SR ₂	SR ₃	SR ₃	SR ₃	SR ₃	RP ₂	RP ₃	R
SR₃	SR ₄	SR ₁	RP ₁	SR ₂	SR ₃	SR ₄	SR ₄	SR ₄	RP ₂	RP ₃	R
SR₄	SR ₅	SR ₁	RP ₁	SR ₂	SR ₃	SR ₄	SR ₅	SR ₅	RP ₂	RP ₃	R
SR₅	S	SR ₁	RP ₁	SR ₂	SR ₃	SR ₄	SR ₅	S	RP ₂	RP ₃	R
RP₁	RP ₂	RP ₂	RP ₁	RP ₂	RP ₂	RP ₂	RP ₂	RP ₂	RP ₂	RP ₂	RP ₂
RP₂	RP ₃	RP ₃	RP ₁	RP ₃	RP ₃	RP ₃	RP ₃	RP ₃	RP ₂	RP ₃	RP ₃
RP₃	R	R	RP ₁	R	R	R	R	R	RP ₂	RP ₃	R

The lattice was initialized by randomly assigning each lattice point to the S, R, or P state. At each transfer, each point experienced an immigration event with probability 1/3. If an immigration event occurred, a point (representing the source of a migration) within the focal point's neighborhood was chosen at random. As this


simulation emulated the *Restricted Community* treatment, the neighborhood was restricted to the four nearest lattice points. In the event of migration, the transition matrix in SI Table 1 was consulted. If migration did not occur, the transitions outlined in [S1] were followed. The entire lattice was updated synchronously, using the previous lattice as a source for all migrations. The resulting arrangement of states after 100 cycles was used to initialize the experimental metapopulations of the *Community* treatments (using a 12×16 lattice) as well as the *Community* simulations (using a variety of lattice dimensions).

Calculating the number of cell divisions. We recorded the abundances of cells every $J=6$ days. Let the number of resistant cells on day t be $R(t)$. Thus, we have recorded $R(Ji)$ for $i \in \{0,1,2,3,\dots,L\}$ (where $L=6$ is the total number of recorded intervals). For all positive t values less than JL for which we did not record abundance, we linearly interpolate between the nearest known R values to estimate the $R(t)$ value. Thus, given a dilution factor of $f=1/40$, the total number of cell divisions (D) is:

$$D = \sum_{i=0}^{L-1} \left\{ \left(\frac{(J+1) - \phi(J-1)}{2} \right) R(J(i+1)) + \left(\frac{(J-1) - \phi(J+1)}{2} \right) R(Ji) \right\}$$


RESULTS

Growth Curves.




SI Figure 2: The growth curves for the three ancestral strains, sensitive (S), resistant (R), and producer (P) are shown. The least-squares parameters for the corresponding Monod functions are given in Table 1 in the Methods. The simulations described in the Methods used these fitted growth curves.

Simulations.




SI Figure 3: The average fitness of resistant cells relative to their ancestor after 100 cycles of simulated evolution (as described in the main text) in a 100×100 lattice. Mean relative fitness of each treatment is shown and bars give the standard error of the mean. The asterisk indicates a significant difference (Welch's two sample t test; $t_{74.39}=39.44$, $p<0.001$).

Small Lattice Simulations.




SI Figure 4: The average fitness of resistant cells relative to their ancestor after 36 cycles of simulated evolution in a 12×16 lattice. Simulation runs in the *Unrestricted Community* treatment in which the three strains did not coexist were excluded. Mean relative fitness of each treatment is shown and bars give the standard error of the mean. Letters distinguish treatments significantly different using post hoc comparisons (single factor ANOVA; $F_{2,2208}=1427.5$, $p < 0.001$, multiple comparisons by Tukey's HSD).



SI Figure 5: The average fitness of resistant cells relative to their ancestor after 36 cycles of simulated evolution in a 12×16 lattice. Mean relative fitness of each treatment is shown and bars give the standard error of the mean. Letters distinguish treatments significantly different using post hoc comparisons (single factor ANOVA; $F_{2,3111}=31.24$, $p<0.001$, multiple comparisons by Tukey's HSD).

Number of Divisions.



SI Figure 6: The number of resistant cell divisions. Mean number of cell divisions in each treatment is shown and bars give the standard error of the mean. Significant differences among the treatments were not found (single factor ANOVA; $F_{2,12}=1.86$, $p=0.1976$).

Simulation Parameters.

Parameter	Description (units)	Value(s)
L	Length of the lattice (points)	16 or 100
W	Width of the lattice (points)	12 or 100
T	Duration of growth phase (hours)	12
C	Number of growth cycles (unitless)	36, 100, or 400
K	Number of mutant resistant strains (unitless)	7
m_Y	Maximum growth rate (abundance/hour)	0.61
$\{K_S, K_P, K_{R^0}, K_{R^1}, K_{R^2}, K_{R^3}, K_{R^4}, K_{R^5}, K_{R^6}, K_{R^7}\}$	Monod constant of ancestral strains and resistant mutants (abundance)	{0.165, 0.93, 0.341, 0.27, 0.28, 0.29, 0.30, 0.31, 0.32, 0.33}
$\{r^k(0), r^0(0)\}$	Initial competition amount (abundance)	{1/40, 1/40}
f	Dilution factor (unitless)	1/40
a	Probability of migration (unitless)	1/3
a_{crit}	Critical abundance for persistence (abundance)	0.00275
p	Mutation probability per transfer per well (unitless)	1/100
g	Fraction of resistant subpopulation converted to a random mutant given a mutation event (unitless)	1/2

Chapter 2: How the tortoise beats the hare: Slow and steady adaptation in structured populations suggests a rugged fitness landscape in *Escherichia coli*

Abstract:

Historical contingency in evolution requires genetic epistasis. Using Wright's metaphor of the adaptive landscape, epistasis can yield a multi-peaked or "rugged" topography. In an unstructured population, a lineage with selective access to multiple peaks is expected to rapidly fix on one, which may not be the highest peak. Contrarily, beneficial mutations in a population with spatially restricted migration take longer to fix, allowing distant parts of the population to explore the landscape semi-independently. Such a population can simultaneously discover multiple peaks and the genotype at the highest discovered peak is expected to fix eventually. Thus, structured populations sacrifice speed of adaptation for breadth of search. As in the Tortoise-Hare fable, the structured population (Tortoise) starts relatively slow, but eventually surpasses the unstructured population (Hare) in average fitness. In contrast, on single-peak landscapes (e.g., systems lacking epistasis), all uphill paths converge. Given such "smooth" topography, breadth of search is devalued, and a structured population only lags behind an unstructured population in average fitness (ultimately converging). Thus, the Tortoise-Hare pattern is an indicator of ruggedness. After verifying these predictions in simulated populations where ruggedness is manipulable, we then explore average fitness in metapopulations of *Escherichia coli*. Consistent with a rugged landscape topography, we find a Tortoise-

Hare pattern. Further, we find that structured populations accumulate more mutations, suggesting that distant peaks are higher. This approach can be used to unveil landscape topography in other systems, and we discuss its application for antibiotic resistance, engineering problems, and elements of Wright's Shifting Balance Process.

Introduction

The degree to which evolution is constrained by historical events experienced by an evolving lineage has been vigorously debated (64). Is it the case that distinct adaptive paths tend to converge (the evolutionary equivalent of all of Chaucer's roads leading to Rome), or do initial disparities lead to different end points (as where Frost's road less traveled makes all the difference)? Gould suggested that some evolutionary outcomes would be contingent on chance events, where an initial difference would lead to an altered result if "life's tape" was replayed (65). At the level of mutations, Gould's paradigm implies that different mutational steps lead to different ultimate evolutionary outcomes. In contrast, evolutionary convergence is predicted if every mutation has the same fitness effect despite the genotype in which it occurs. In such a case, adaptive evolution would be the accumulation of unconditionally beneficial mutations, and the endpoint of evolution would not depend on the order in which mutations are acquired. Counter to such mutational convergence, Gould's conclusions would require that the fitness effect of a mutation depends on genetic background. Does such genetic epistasis tend to make evolution contingent? Here, we use a classic metaphor to address this question.

Introduced by Sewall Wright (66), the adaptive landscape is a useful visual framework for discussing the factors that influence the contingency of evolutionary trajectories. One incarnation of Wright's landscape portrays the relationship between an organism's genotype and its fitness as a topographical map. Imagine

placing all possible genotypes of an asexual organism together on a plane, where the distance between two genotypes represents the number of mutations needed to generate one genotype from the other. Each genotype is assigned a height directly proportional to its fitness (the third dimension). An evolving population is represented as a cloud of points on the resulting landscape, where each member of the population is a point. Novel genotypes arise in the population via mutations, expanding the extent of the cloud. In contrast, natural selection diminishes the range of the cloud, shifting its weight uphill as less fit genotypes are culled. Thus the combination of mutation and selection leads to the population "climbing" adaptive hills to their "peak," which is a genotype from which all mutations are deleterious. If we assume strong selection and weak mutation (*SSWM*), the population-cloud exists as a single climbing point, where the rapid fixation of each rare beneficial mutation shifts the point uphill (64, 67). Overall, a population's evolutionary trajectory is taken to be sensitive to the gradients on this three-dimensional landscape.

As pointed out by many authors, including Wright (68–70), the actual geometry of the space of possible genotypes has extremely high dimensionality, which cannot be projected into two dimensional space in a way that preserves all distances. We explore an alteration of the classic representation (see (71–73) that ensures genotypes differing by a single mutation are equidistant (while the distance between genotypes differing by multiple mutations is distorted). This approach involves creating a network, in which nodes are genotypes and edges connect genotypes differing by a single substitution. This network is embedded in two

dimensions, where genotypes are grouped along the abscissa by their distance from a common genotype and along the ordinate by their fitness.

Using this representation, a mutation network without epistasis is shown in Figure 1a. This network would also be labeled as a “smooth landscape,” as the single peak is accessible from any other genotype. Evolution of populations on this landscape gives an example of mutational convergence. Under *SSWM* assumptions, we show that despite three different initial mutational steps, independent evolutionary trajectories converge at the peak (Fig. 1a) and fitness likewise converges (Fig. 1b). In contrast, Figure 1c shows a network with sign epistasis, where the sign of the fitness effect of a mutation depends on the background in which it occurs. Such sign epistasis is a necessary (but not sufficient) condition for the existence of multiple peaks. On this “rugged” landscape, the final genotype reached under three independent trajectories is contingent upon the initial mutation (Fig. 1c). In this case, fitness values of different populations can remain divergent over time if peaks are heterogeneous in height (Fig. 1d). Here we see that a population can become trapped at a suboptimal peak in the presence of epistasis.

Because Sewall Wright thought epistasis was pervasive (74), he was particularly concerned about confinement of populations at suboptimal peaks within rugged landscapes. He proposed the Shifting Balance Process (SBP) to explain how populations move from lower to higher peaks. Integral to the SBP is population structure. Wright envisioned a population that was distributed into semi-isolated,

sparsely-populated subpopulations (demes) in which genetic drift enabled some subpopulations to take downward steps by fixing deleterious mutations. In this way, a subset of the metapopulation is able to move from one peak's domain of attraction to another, thus crossing "adaptive valleys." Therefore, Wright's SBP depends on two critical assumptions: the presence of epistasis generating landscape ruggedness and the presence of population structure.

Ironically, by using the content of one assumption as an experimental variable, the other assumption can be tested: specifically, landscape ruggedness can be verified by manipulating population structure. Upon first glance, population structure would seem to hinder adaptation. For a population in which migration is not spatially restricted (unstructured population), a beneficial mutant that arises can rapidly fix in what is termed a selective sweep. On the other hand, a favored mutant arising in a population with restricted migration (structured population) advances more slowly in what might be termed a "selective creep." By reducing the rate of initial adaptation, the slow competitive displacement occurring within a structured population may also allow multiple semi-independent searches of the fitness landscape by geographically distant regions of the population. For a smooth landscape (e.g., Fig 1a), this enhanced exploration is superfluous as all roads lead to Rome (the single peak). Therefore, on smooth landscapes, structure only slows adaptation. However, on a rugged landscape, additional exploration may reveal alternate peaks. For instance, in Fig. 1c, while an unstructured population might exclusively follow one of the colored trajectories, a structured population may be

able to explore them all simultaneously. In this way, a structured population gets more plays of “life’s tape,” and can compare the outcomes. As discovered peaks may differ in height, on average such comparisons yield better endpoints. On a rugged landscape, fitnesses in populations differing in structure emulate the classic Tortoise-Hare fable. Specifically the unstructured population initially adapts faster (the Hare) but is overtaken by the structured population (the Tortoise), which is a poor starter but a strong finisher. Importantly, on a smooth landscape, the Tortoise never takes the lead, and the crossing of average fitness trajectories is not predicted. Thus, when manipulations to population structure produce a Tortoise-Hare pattern, we have a signature of ruggedness.

Before investigating this signature in a biological system in which landscape topography is cryptic, we confirm the above predictions using a computational system in which landscape topography is known and manipulable. Specifically, we control landscape ruggedness by employing Kaufmann’s NK model (75–77) and then track evolving metapopulations of bit strings, in which the pattern of migration between demes is manipulated. Following this simulation, we then turn to evolving metapopulations of *Escherichia coli* under a similar experimental manipulation of population structure. Discovery of a Tortoise-Hare pattern would be indicative of a rugged topography, in which evolution could be historically contingent.

Results and Discussion

Patterns of Average Fitness

In the NK model, simulated organisms are bit strings of length N , and the parameter K is the number of loci affecting the fitness contribution of each locus (see Methods). As K increases, the level of epistatic interaction increases, yielding more rugged landscapes; hereafter, we refer to K as a “ruggedness” parameter. We explore how ruggedness affects fitness trajectories in evolving metapopulations that differ in population structure. We consider either metapopulations with spatial restrictions to migration (hereafter, the Restricted treatment) or metapopulations where migration can occur between any two demes (Unrestricted treatment). For a smooth landscape topography ($K=0$, $N=15$), average fitness initially increases more rapidly in the Unrestricted treatment relative to the Restricted treatment; however, both trajectories converge over time (Fig. 2a). For a rugged landscape ($K=8$, $N=15$), fitness in the Restricted treatment once again lags behind fitness in the Unrestricted treatment at the outset. Instead of converging, however, the fitness trajectories cross, yielding a higher final fitness for the spatially restricted treatment (Fig. 2b). Indeed, we find significantly higher fitness in the Restricted treatment for $K>3$ at the end of our simulation (Fig. 2c; Mann-Whitney tests with Bonferroni corrections, $p<0.001$). The pattern in Figure 2b agrees with the Tortoise-Hare prediction, while the crossing does not occur for the Unrestricted treatment in Figure 2a. More generally, with sufficient ruggedness, a structured population can eventually outperform an unstructured population (Fig. 2c). This pattern is not limited to our

specific form of structure, as Bergman *et al.* (1995) (78) reported a similar result using an NK model where bit strings dispersed variable distances along a single dimension.

We next turned to examining fitness trajectories in evolving metapopulations of *Escherichia coli*. Similar to the NK model, we propagated the bacteria under two treatments differing in migration pattern: Restricted and Unrestricted (see Methods). Early and late in the evolutionary run, we sampled five random isolates from each metapopulation and determined their fitness relative to the common ancestor. Early in the experiment (at transfer 12), fitness in the Restricted treatment was significantly lower than the Unrestricted treatment (Fig. 3; Mann-Whitney test, $p=0.015$). However, at the end of the experiment (transfer 36), the opposite pattern was found, with fitness in the Restricted treatment surpassing the Unrestricted treatment (Fig. 3; Mann-Whitney test, $p=0.015$). This pattern is consistent with a rugged landscape topography.

Patterns of Evolutionary Distance

There are a few ways to account for the benefit that population structure confers on rugged landscapes (Figs. 2b and 3). First, a population may have access to multiple peaks that differ in height. A structured population can explore multiple domains in parallel, eventually comparing the results. Thus, it will tend to attain a higher endpoint; for the same reasons that the expectation of the maximum of a sample increases with sample size. This effect holds when all peaks are equidistant from the

ancestral population. A second possibility (not mutually exclusive with the first) is that peaks differ in both height and distance from the ancestral population. Suppose that initial mutations leading to the more distant peaks are less beneficial than mutations leading to the nearby peaks. In this case, intermediate genotypes approaching distant peaks risk being outcompeted in an unstructured population (consider the first mutant on the blue path to the more distant peak in Fig. 1c competing against the other first mutants). This is because the slower fixation of these intermediates allows for better competitors (from domains of nearer peaks) to arise. In contrast, these more distant peaks become accessible in a structured population due to reduced competitive displacement. If some of these distant peaks are also higher, then structured populations are predicted to both achieve better fitness and accumulate more mutations.

To explore the number of mutations accrued by evolving populations, we first return to the NK model. We define evolutionary distance to be the number of mutational differences between an evolved isolate and its ancestor. In the NK model, this is the Hamming distance (79). As ruggedness increases, the degree of population structure affects final evolutionary distance from the ancestor; we find a significantly higher distance in the Restricted treatment for $K > 3$ at the end of our simulation (Fig. 4; Mann-Whitney tests with Bonferroni corrections, $p < 0.001$). Thus, on a rugged landscape, a population with restricted migration moves both higher (Fig. 2c) and further (Fig. 4) than a less structured population.

Evolutionary distance was explored similarly in the *E. coli* metapopulations. In addition to assessing fitness on the isolates from the end of the experiment (transfer 36), we sequenced their genomes to determine their evolutionary distance from the common ancestor. The locations of all identified mutations in all sequenced isolates are shown in Figure 5. The number of mutations accumulated in each isolate is the evolutionary distance (the points to the right of the table in Fig. 5). We find that isolates have moved a significantly greater distance in the Restricted treatment (Mann-Whitney test, $p=0.045$). The increased distance traversed by digital and bacterial populations is consistent with rugged landscapes in which more distant peaks are being reached by structured populations.

The ability of the structured populations to move higher (in average fitness) and further (in evolutionary distance) is engendered by the capacity for parallel search. The presence of simultaneous selective creeps should increase the standing diversity within a structured population relative to an unstructured one. In line with this prediction, the metapopulations in the Restricted treatment had significantly higher genotypic diversity than the Unrestricted treatment (Mann-Whitney test on the nucleotide diversity index \mathbf{p} , $p=0.016$). While greater diversity in structured populations is expected despite the topography of the landscape (see Supplemental Figure 1), this observation does substantiate the basic logic of our argument.

Previous Empirical Work

An extreme form of population structure involves a set of completely isolated populations. If the landscape contains multiple peaks of different heights, these populations can diverge in fitness, genotype and phenotype. Korona *et al.* (80) and Melnyk & Kassen (81) measured phenotypic diversity within, and diversity in fitness among, a set of replicate evolving microbial populations. Under some growth conditions (growth on agar surfaces for *Rastonia eutropha* (80) and growth in minimal xylose for *Pseudomonas fluorescens* (81)) these authors found that both forms of diversity remained high at the end of the experiment, which they took as evidence for a rugged landscape. Under other growth conditions (growth in liquid for *R. eutropha* (80) and growth in minimal glucose for *P. fluorescens* (81)) they found lower diversity, consistent with a smoother landscape. Like the present study, these earlier studies used statistical patterns to infer topographical properties of landscapes.

This statistical approach was also used in a recent study by Kryazhimskiy *et al.* (82), which strongly parallels our experiment. They propagated metapopulations of asexual *Saccharomyces cerevisiae* and varied the rate (as opposed to pattern) of migration. In contrast to the results we report, they find yeast from treatments with higher rates of migration evolved higher fitness over the course of their experiment. This led the authors to conclude that epistasis was weak and the landscape was smooth. While it is absolutely plausible that their yeast and our bacteria differ in landscape topography, we discuss some alternative interpretations of their results in the Supplement.

Rather than inferring landscape topography from population-level statistical patterns, an alternative approach involves fully characterizing the fitnesses of genotypes in a small section of the landscape (83–91). This approach often involves considering two genotypes differing by M mutations, engineering all 2^M possible combinations, and assessing the fitness of each constructed genotype. Epistasis can be gauged directly by measuring the influence of genetic background on the fitness effect of a mutation. Some studies using this approach have uncovered instances of sign epistasis (83, 85, 90), while other studies have found only magnitude epistasis (84, 87). Some studies have reported multiple peaks (85, 90), while others have found only a single peak (83, 87). As with the inferential approach described above, this engineering approach has revealed a potential diversity of landscape topography. We suggest that combining the (top-down) inferential approach and (bottom-up) engineering approach is a promising direction for exploring the nature of adaptation.

Wright's Shifting Balance Process

During the Modern Synthesis, two divergent views on adaptation were born. One view considers adaptation as the sequential fixation of unconditionally beneficial mutations. This perspective (often associated with Fisher) does not highlight epistatic contingency, focusing instead on selection and mutation as the major processes of evolution. A second perspective (often linked to Wright) recognizes epistatic interaction as constraining adaptation. Based on his empirical

observations, Wright felt that such epistasis was pervasive, leading to multiple peaks in his adaptive landscape metaphor (64). In this view, evolving populations can become trapped in suboptimal positions. To illustrate how escape from the domains of suboptimal peaks was possible, Wright introduced the Shifting Balance Process (SBP). Importantly, he additionally assumed that adapting populations were spatially structured as a metapopulation of semi-isolated demes. The SBP is heuristically divided into three phases. In Phase I, the demes, which Wright posited were small, drift in genotype space, enabling movement into other domains. In Phase II, selection within demes produces a set of semi-independent hill-climbing episodes. In Phase III, the resulting genotypes are exchanged among demes and the best competitors spread through the metapopulation at large. Thus, Wright's view of adaptation, in contrast to Fisher's, invokes a complex combination of processes, specifically brought together to solve a problem generated by epistatic contingency.

While there have been various theoretical explorations of the plausibility of SBP (78, 92–95), Wright's ideas have been criticized because they demand a delicate balance of various evolutionary processes. For instance, populations need be small enough for effective drift (Phase I), but large enough for effective selection within demes (Phase II). Migration should be sufficiently restricted for drift and selection within demes (Phase I and II), but sufficiently unrestricted for effective exchange of genotypes among demes (Phase III). However, if we abandon Wright's goal of explaining how populations cross valleys, many of these conflicts vanish. Both Fisher and Wright acknowledged that environmental change could alter the

landscape, and, in the process, reposition the peaks. Imagine that a population experiences such a change and subsequently resides somewhere on the new landscape with access to multiple domains. In our experiment, our ancestor contained deleterious mutations and evolved in a stressful environment (see Methods), which potentially yielded access to multiple peaks. In this case, demes need not be small for the discovery of multiple peaks (and indeed, our experimental demes were large). With large subpopulations, selection within demes will proceed efficiently; however, limitations to migration between demes will still allow for parallel exploration of a rugged landscape. Thus, Phases II and III can jointly yield adapted populations even if Phase I is absent. If landscapes are indeed rugged, population structure can retain the critical role Wright foresaw, even if all the details of the SBP are not present.

Applications

One case where populations are potentially poised in multiple domains on a landscape involves the evolution of microbes exposed to antibiotics. When a bacterial population experiences a sufficiently high concentration of an antibiotic, susceptible genotypes are replaced by resistant mutants. When the drug is removed, these mutants tend to carry fitness costs relative to their susceptible progenitors. The cost can be alleviated by a mutation resulting in reversion to susceptibility or a mutation that compensates for the impairment without loss of resistance (37, 96). There is some evidence that reversion and compensation constitute distinct peaks in a rugged landscape (40, 97). Thus, we see that a changing environment (exposure

and removal of a drug) may position a microbial population at a landscape position where multiple peaks are accessible (71). It is at such a position that population structure may influence the evolutionary trajectory. Björkman *et al.* (97) and Nagaev *et al.* (40) evolved *Salmonella typhimurium* and *Staphylococcus aureus* resistant to fusidic acid either in well-mixed flasks or within murids (mice or rats). These authors found that the bacteria more often reverted when grown *in vivo* than *in vitro*. They explain these results by noting that the flask and murid environments differ markedly and may consequently place different selective pressures on revertants and compensated strains (indeed, they present data to this effect). In our terminology, the landscape in a flask and a mouse may be different. However, even if the landscape was identical (but rugged) in both, the results might not be unexpected because a murid environment is highly structured and a flask is not. Thus, if the “reversion peak” is higher than most to all of the “compensation peaks” (the authors present data consistent with this ordering) then evolution in a structured environment is predicted to revert at higher frequency. In this way, the structure that pathogenic bacteria experience (including in the bodies of human hosts) can potentially influence the course of antibiotic resistance evolution.

Not only can the ideas in this paper apply in a medical context, but also they may address practical engineering problems. Evolutionary principles have been utilized to find solutions to computational problems, a discipline known as evolutionary computation. In this field, putative solutions constitute a population, new solutions are generated by mutation and recombination, and better solutions can outcompete

their contemporaries. One defining feature of a difficult problem is the presence of multiple optima in the map from the specification of a solution (i.e., genotype) to its quality (i.e., fitness). As early as 1967, Bossert (98) suggested that dividing a population of solutions into subpopulations could yield better evolutionary outcomes. Subsequently, the inclusion of population subdivision in evolutionary algorithms has produced better solutions in a variety of applications, including analogue circuit design (99), financial trading models (100), and multi-objective scheduling (101). Besides the efficiency in networked computational resources that accompanies population subdivision, a deeper exploration of the landscape of solutions is predicted to occur when multiple domains can be semi-independently searched (102, 103).

Synthesis

The prevalence of contingency or convergence in evolution is affected by the underlying topography of adaptive landscapes. When no epistasis is present, a classic “smooth” landscape results, and ultimately evolution converges to the single peak. Contingency requires (sign) epistasis, specifically of a kind generating multi-peaked, or classic “rugged,” landscapes. (We again note that sign epistasis is not sufficient for multi-peaked landscapes and indeed can constrain evolution to a subset of paths to a single peak; see Weinreich *et al.* (83) for an example.) In the case of heterogeneous peak height, population structure may enable the simultaneous exploration of multiple domains and ultimately lead to the discovery of higher peaks than would be possible in an unstructured population. Thus, using

population structure as an experimental variable, we have a signal for this kind of ruggedness. In the case of our bacterial populations, we have presented a pattern consistent with ruggedness. Additionally, it appears that structured populations move further in the landscape, suggesting that the most accessible peaks may not be the highest. Ultimately it is an empirical issue whether other biological systems possess such ruggedness. However, statistical patterns in fitness and evolutionary distance may help to distinguish landscape topography, and thus shed light on the prospects for contingency. Specifically, a rugged landscape topography can be inferred by comparing structured (Tortoise) and unstructured (Hare) populations and assessing whether slow and steady adaptation “wins the race.”

Methods

Ancestral strain

The bacterial ancestor was derived from a K-12 strain of *E. coli* (BZB1011) by selecting for resistance to colicin E2, then colicin D, and then phage T6. Both resistance to colicin E2 and phage T6 are known to be individually costly (34, 104).

Experimental treatments

Each metapopulation was comprised of 96 subpopulations (the 96 wells of an 8×12 microtiter plate). The metapopulation was initiated with the ancestral strain in each well. These subpopulations grew for 12 hours in 200 mL of lysogeny broth (LB-Miller) supplemented with a sub-inhibitory concentration of tetracycline (0.25 mg/mL). After growth, each well in the metapopulation was diluted 40-fold into fresh growth medium using a 96 slot-pin multi-blot replicator (5 μL into 200 μL). Immediately following dilution, migrations among wells occurred. Migration was either restricted to occur between subpopulations adjacent to each other or was unrestricted. In both treatments, every well had 1/3 probability of experiencing an immigration event from one random well in its neighborhood. In the Restricted treatment, this neighborhood included the wells directly north, east, south or west of the focal well (using periodic boundary conditions to eliminate edge-effects). In the Unrestricted treatment, the neighborhood included all wells minus the focal well. All migration events were executed by a BioRobot 8000 liquid-handling robot (Qiagen), which transferred 5 μL from the source well within the plate from the previous transfer into the destination well within the plate from the current

transfer. Between transfers, plates were incubated (37 °C) and shaken (350 rpm using a microtiter shaker). Each metapopulation was propagated for a total of 36 transfers and each treatment contained five replicates.

Competition assay

We chose five random isolates from the last transfer of each metapopulation (here we denote any one of these strains as E). We marked our ancestor with resistance to phage T5 (we denote this marked ancestor as A^R). Before the competition, E and A were grown separately in 200 μL of growth medium for two 12-hour cycles (with 40-fold dilution at transfer). After this acclimation phase, we added 5 μL of E and 5 μL of A to a well containing 200 μL of fresh growth medium. The titer of each strain was assessed, through selective plating with and without phage T5, immediately after the competition was initiated and again after 12 hours. If $X(t)$ is the titer of strain X at time t , then the fitness of the evolved strain relative to its ancestor is given by:

$$w(E, A) = \frac{\ln(E(12)/E(0))}{\ln(A(12)/A(0))}.$$

Whole Genome Resequencing

Using the same isolates from the last transfer, we performed chromosomal DNA extractions using Qiagen Mini DNA Kits. Each sample was barcoded and multiplexed to 24 samples per lane with Illumina TruSeq. Whole genome resequencing (University of Washington High Throughput Genomics Unit) was performed with single-end 36-bp unpaired reads on Illumina HiSeq to an average of 30X coverage. Illumina reads were aligned for mutational discovery by Breseq 0.19 (J. E. Barrick, unpublished algorithm) against *E. coli* W311 [GenBank: AP009048]. Alignments were considered only if they covered 95% of a read. For every isolate, Sanger sequencing (Genewiz) of several loci (*fimE*, *marR*, *ompF*, and *stfR*) was used to confirm putative mutations.

NK Model

For the simulations, individuals were embedded within 96 demes in an 8 x 12 array. Each deme contained 1000 organisms. Each organism's genotype was a bit string (fixed length ordered list of 0's and 1's) of length $N=15$. The fitness of an organism was the sum of the fitness contributions of each of the 15 loci divided by the number of loci. The contribution of each locus was determined by its allelic state and the allelic states of the subsequent K loci (wrapping to the beginning of the bit string as needed). For each locus, 2^{K+1} random numbers (uniformly distributed between 0 and 1) described all possible fitness contributions of that locus (given any possible combination of alleles at relevant loci). Thus a mutation at a single locus affected the fitness contribution of the mutated locus and K other loci. Selection within a deme

involved the removal of a random organism, regardless of fitness, and its replacement by the birth of an organism from the same deme chosen by a fitness-weighted lottery. Upon birth, the offspring bit string differed from its parent at a random locus with probability 0.1 (the mutation rate). This Moran death-birth process was iterated 1000 times for each deme, followed by migration between demes. During each migration event, 25 individuals were chosen at random and removed from one deme (the destination), and then replaced by copies of 25 individuals chosen at random from the other deme (the source). Each deme experienced an immigration event with probability $1/3$. Migration was either restricted or unrestricted in precisely the same manner as the bacterial experiment above. Each replicate run of the Unrestricted treatment was paired with a replicate of the Restricted treatment, where each member of the pair shared the same NK landscape as well as the same ancestor (a random bit string used to populate the entire metapopulation). In the figures, one selection-migration episode is termed an “update.”

Chapter 2 Figures:

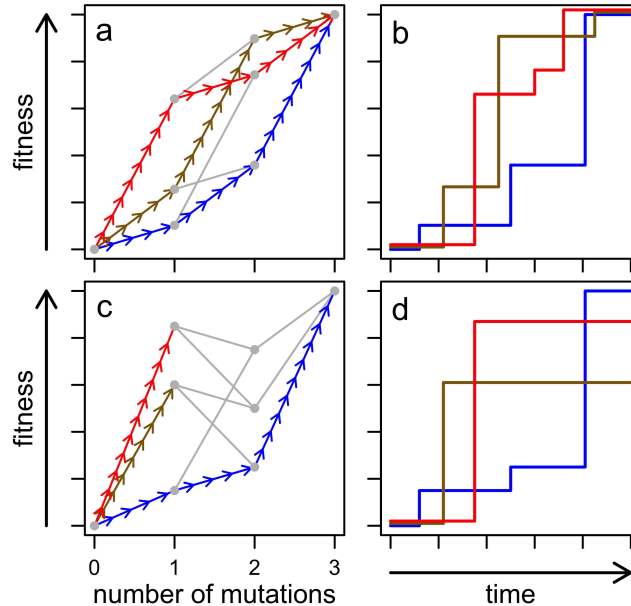


Figure 1. Adaptive paths in hypothetical landscapes. Here we consider a simple biallelic three locus system. (a) The adaptive landscape can be visualized by plotting genotype fitness as a function of the number of mutations on a wild type background. Each of the $2^3=8$ genotypes is given by a gray point and edges (arrows or gray lines) connect genotypes differing by a single mutation. An adaptive peak is a genotype from which all mutations are detrimental. A hypothetical landscape with a single peak (the triple mutant) is shown here. A selectively accessible path exists between two genotypes if a series of beneficial mutations connects the less fit genotype to the more fit genotype. On this “smooth” landscape, all of the $3!=6$ paths between the wild type (lacking mutations) and the triple mutant are selectively accessible; three of these paths are shown by the arrows in different colors. (b) Average fitness over time is shown for

three possible populations following the paths in part a. If we assume that selection is strong and mutation is weak, we can represent the fixation of each beneficial mutation as a step up in the fitness trajectory. All trajectories converge on the same final fitness value. (c) A hypothetical landscape with multiple peaks. Starting with the wild type, selection can take the population to different adaptive peaks on this “rugged” landscape, as illustrated by the different colored trajectories. (d) Average fitness over time is shown again for three possible populations following the paths in part c. The final fitness of different evolving populations can vary.




Figure 2. Fitness in the NK model. Metapopulations of bit strings of length $N=15$ evolved where either migration was restricted to occur between neighboring demes or migration was unrestricted (occurring between any two demes). (a) Average fitness in the metapopulation is shown over time on a smooth landscape ($K=0$) or (b) a rugged landscapes ($K=8$). (c) Average fitness at time point 1000 is shown as a function of the ruggedness parameter, K . Note the values at $K=0$ and $K=8$ in part c correspond to the values at time point 1000 in parts a and b, respectively. In all plots, points represent the mean of 50 replicates, shading gives the standard error of the mean and asterisks indicate significant differences.

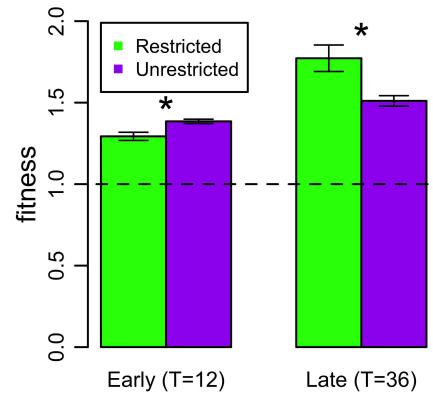


Figure 3. Bacterial fitness. Metapopulations of bacteria evolved where migration was spatially restricted or unrestricted. Average relative fitness of five isolates randomly sampled from the metapopulation is shown early in the experiment (at transfer T=12) and late in the experiment (at transfer T=36). As in Figure 2b, the ordering of fitnesses for the two treatments flips over time. Bars represent the mean of 5 replicate metapopulations, whiskers give the standard error, and asterisks denote significant differences.

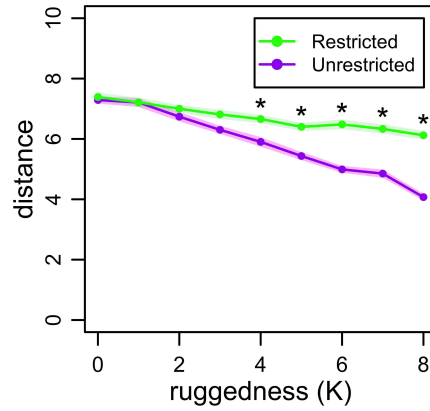


Figure 4. Distance in the NK model. Metapopulations of bit strings of length $N=15$ evolved where migration was spatially restricted or unrestricted. Evolutionary distance is the number of bits differing between an evolved isolate and its ancestor (the Hamming distance). Average distance at time point 1000 is shown as a function of the ruggedness parameter, K . Points represent the mean of 50 replicates, shading gives the standard error of the mean, and asterisks denote significant differences.




Figure 5. Bacterial distance. Metapopulations of bacteria evolved where migration was spatially restricted or unrestricted. At the end of the experiment, five isolates from each metapopulation were sequenced at the genome level. The top bar represents the genome of *Escherichia coli*. Genome regions with mutations are magnified for the table. Each isolate is a single row in the table and the location of each of its mutations is indicated by a white mark. The five isolates from each of the five replicate metapopulations are grouped by alternating shades of green (for the Restricted treatment) or purple (for the Unrestricted treatment). The horizontal distance of the point to the right of the table denotes the number of mutations in the isolate (its evolutionary distance). The horizontal distance of the green bar and the purple bar to the right of the table gives the average distance (the average of replicate averages) of

isolates from the Restricted and Unrestricted treatments, respectively. The asterisk denotes a significant difference between treatments.

Supplementary Information:

Diversity in digital and bacterial populations

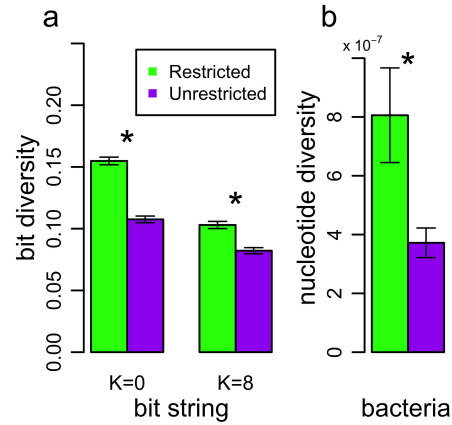
A structured population performs a broader search on the adaptive landscape, as the rate of competitive displacement is lower. Consequently, the standing genetic diversity of a structured population is expected to be greater than diversity in an unstructured population. In Supplemental Figure 1a, we see that this pattern does not depend on landscape ruggedness (Mann-Whitney tests, $p < 0.01$ for $K=0$ and $K=8$). Thus, despite landscape topography, we predict to find higher genetic diversity in a structured population, and this is what we find in our bacterial metapopulations (Mann-Whitney test, $p=0.015$; Supp. Fig. 1b).

Diversity methods

Consider a sample of G genotypes (bit strings or nucleotide sequences). We use the diversity index of Nei and Li (1979) (105):

$$\pi = \frac{2}{G^2} \sum_{i=2}^G \sum_{j=1}^{i-1} \pi_{ij},$$

where p_{ij} is average number of differences (in bits or bases) per site between genotype i and genotype j . We refer to p as bit diversity (in the NK model) or nucleotide diversity (for our bacterial system).



Supplemental Figure 1: Genetic diversity in the digital and bacterial populations. (a) Average bit diversity of a sample of eight evolved bit strings from time point 1000 in the NK model simulations. Whether the landscape is smooth (K=0) or rugged (K=8), diversity is significantly greater in the Restricted treatment than the Unrestricted treatment. Bars represent the mean of 40 replicates. (b) Average nucleotide diversity within bacterial metapopulations at the final transfer of the experiment (T=36). For each metapopulation, full genome data from each of its five isolates was used to compute the diversity index. Nucleotide diversity is significantly greater in the Restricted treatment than the Unrestricted treatment. Bars represent the mean of 5 replicates. In both parts of the figure, whiskers give the standard error and asterisks indicate significant differences.

How a rugged landscape can fail to give a Tortoise-Hare signal

On a rugged landscape, fitness in a structured population will increase more slowly than an unstructured population (the Tortoise initially lags behind, before overtaking, the Hare; see Fig. 2b). That is, for populations differing in structure evolving on a rugged landscape, early evolution will produce a pattern similar to that predicted under a smooth landscape (e.g., before time point 250, the pattern in Figure 2b would be hard to distinguish from the entire trajectory of Figure 2a). While the presence of a Tortoise-Hare pattern indicates ruggedness, its absence does not necessarily imply a smooth landscape.

For example, in the experiment of Kryazhimskiy *et al.* (2012) (82), the unstructured population ended the experiment with higher average fitness. This pattern is consistent with a smooth landscape, but is not inconsistent with a rugged one. As the authors themselves acknowledge, had their experiment run longer, they may have observed higher fitness under lower rates of migration (i.e., a fitness crossing).

Even when there is abundant time for evolution to take place, it is still possible that evolution on a rugged landscape will fail to yield the Tortoise-Hare pattern.

For instance, it is possible that the landscape is rugged, but peaks are of a homogeneous height. This could produce the fitness pattern shown in Figure 2a.

Finally, the landscape could be rugged with heterogeneous peak heights, but the ancestor could be positioned in the domain of a single peak. Consistent with this possibility, Kryazhimskiy *et al.* started their experiment with a lab-adapted strain of

yeast and evolved their populations under standard laboratory conditions. If their yeast had access to only a single domain in a rugged landscape, a Tortoise-Hare pattern would not be expected. In such a case Kryazhimskiy *et al.* would be justified in claiming that the *local* topography of such a landscape was smooth (and indeed, they restrict their claim of smoothness accordingly).


As outlined in our Methods, we introduced several deleterious mutations into our ancestor and evolved our populations under a stressful environment (in the presence of sub-lethal concentrations of the antibiotic tetracycline). Such manipulation was intended to displace our ancestral genotype from a peak, but it also may have placed it at a point where multiple domains were accessible. To address the effect of ancestor starting position, we describe additional NK simulations here. In addition to starting our ancestor at a random bit string, we consider three other starting positions: (i) valley, (ii) pre-adapted, and (iii) “silver-spoon.” For the valley simulations, we performed a “hill-plunge of steepest descent,” moving downhill from a random genotype until we hit a valley genotype (a genotype from which all mutations were beneficial), which served as the ancestor. For the pre-adapted simulations, we allowed a random ancestor to evolve briefly (in an unstructured population) to produce a “pre-adapted” ancestor. For the silver-spoon simulations, all genotypes were ranked for fitness and the genotype defining the 99th fitness percentile was chosen as the ancestor.

In the random and valley starting positions, the Tortoise-Hare pattern was observed and the Structured treatment ended at significantly higher average fitness than the Unrestricted treatment (Mann-Whitney tests, $p < 0.001$; Supp. Fig. 1b). However, in the

pre-adapted and silver-spoon starting positions, the Tortoise-Hare pattern was not seen and fitness was indistinguishable between the treatments in the long run (Mann-Whitney test, $p=0.36$ and $p=0.57$ respectively). These simulations demonstrate that the starting position of a population in a landscape will influence the statistical pattern of fitness of populations differing in structure.

Additional NK simulation methods

To study the effect of starting position in the adaptive trajectories in structured and unstructured populations, we examined starting the population with different types of ancestors. The “random” ancestor (used in the primary text) is simply a random bit string. To generate the “valley” ancestor, we start with a random bit string and substitutes the worst (lowest fitness) possible mutation until no deleterious mutations are possible, and that bit string is the ancestor for the evolutionary run. To produce the “pre-adapted” ancestor, we start with a random bit string and evolve a population initialized with this bit string under unrestricted migration for 50 updates. Then bit strings are sampled from the evolved population until one is found that has a higher fitness than the starting bit string. That adapted bit string is the ancestor. To determine the “silver-spoon” ancestor, all possible bit strings (2^{15}) are ranked according to fitness and the 99th percentile genotype is the ancestor.



Supplemental Figure 2: Metapopulations of bit strings of length $N=15$ evolved on a rugged landscape ($K=8$), where migration was either restricted or unrestricted. Average fitness in the metapopulation is shown at time point 1000 for a randomly chosen ancestor, an ancestor starting in a valley, an ancestor resulting from adaptation before the run, and an ancestor in the top percentile of fitness (the “silver-spoon” ancestor). Bars represent the mean of 40 replicates, whiskers give the standard error of the mean, and asterisks indicate significant differences.

Chapter 3: Improved adaptation in exogenously and endogenously changing environments

Introduction

On some level, all evolutionary studies involve exploration of how organisms adapt to their environment over some period of time. Adaptation by natural selection requires time in an environment to favor advantageous alleles. A reasonable expectation would be that longer exposure to an environment should yield increased adaptation, as there is more time for beneficial mutations to arise and spread. However, adaptation may be constrained by more than the time for new mutations to arrive. If the fitness effects of mutations depend on the genetic background, then a population may have difficulty arriving at the genotype of highest fitness, despite the length of time spent evolving.

To illustrate situations where adaptation is not limited solely by time, we turn to the metaphor of the “adaptive landscape.” Introduced by Sewall Wright (66), an adaptive landscape is the relationship between an organism’s genotype and its fitness. Imagine placing all individual genotypes of an asexual organism together on a plane where the distance between two genotypes represents the number of mutations needed to generate one genotype from another. Each genotype is assigned a height directly proportional to its fitness in a specified environment. The

evolving population consists of a cloud of points (one point for each member of a population). This cloud increases its range through mutation and shifts its weight uphill through selection (see Figure 1a). Thus the combination of mutation and selection leads to the population “climbing” fitness hills to their “peak,” which is a genotype from which all mutations are deleterious (downhill). If the landscape is rugged, i.e. having multiple peaks, merely climbing uphill doesn’t guarantee the population reaches the most fit genotype (75). The population may become “trapped” on a sub-optimal peak, a problem Wright addressed with his Shifting Balance Process.

The primary idea behind Shifting Balance Process is that small populations can escape sub-optimal adaptive peaks through genetic drift and begin climbing in a new place. One assumption made to simplify the model is that the environment remains constant during evolution (the peaks maintain their positions). Interestingly, a *changing* environment may provide an alternative mechanism to escape from sub-optimal peaks (66, 106–108). Different environments can produce a change in the mapping from genotype to fitness (e.g., some mutations detrimental in one environment may be beneficial in another). As an environment changes, former adaptive peaks may disappear and new peaks can appear where formerly there were none. In this reshaped adaptive landscape, even a large population, which evolved to a former peak, can be selected to move to a novel genotypic position. Upon returning to the original environment, the population can climb up to a new, perhaps higher, peak (see Figure 1b).

Changing environments can have profound effects on adaptation (109–112). Using a variety of different computational systems, Kashtan et al. (113) found that when the environment changed in specific ways (i.e., when goals from different environments shared sub-problems), the population was able to more rapidly evolve solutions to complex problems than in a constant environment. Alto et al. (114) found that alternating exposure of vesicular stomatitis virus to two different temperatures led to increased performance on both temperatures relative to virus evolved in a constant environment. In these studies, change was externally imposed on the evolving populations. Here we additionally investigate a qualitatively different form of environmental change—namely, where organisms are the agents of change in their environment.

We classify two different types of change in the environment: exogenous and endogenous. Exogenously driven change is brought about by factors outside the influence of the population (e.g., diurnal-nocturnal changes in light or seasonal changes in temperature). In contrast, the evolving population itself affects *endogenous* change. Actions of organisms in the population, including resource usage, waste production and habitat modification, affect the environment and alter the fitness landscape for themselves and future generations. The process whereby organisms modify their environments has been termed niche construction (115–118) or ecosystem engineering (119). Such niche construction can feed back to affect the evolution of the population. For instance, if organisms reduce the quality

of their environment by exploiting certain resources or become susceptible to infectious pathogens, they can depress the fitness of their own, and related, genotypes (see figure 1c and (120)). This change may additionally select for novel genotypes that differ from those currently constituting the population. In theory, both exogenous and endogenous change allows populations to leave (former) adaptive peaks by selection.

To distinguish the effects of exogenously and endogenously changing environments, we need a system that evolves in highly controlled settings. Ideally, this system would be simple (to deduce the fitness effects of individual mutations), fast (to allow for evolution across many generations), and tractable (to control the environment, and the effect organisms have on it). For these reasons we chose to perform our experiments with Avida, a digital platform for the study of evolution. Avida provides an ideal system to test the effects of a changing environment and has been used extensively to investigate a wide array of evolutionary questions (121). The organisms within Avida are simple (the mappings between genotype, phenotype, and fitness are easily determined) and fast (generations last less than a second). The environment for an evolving population can be measured and manipulated precisely. Most importantly, we can explore changing environments: either via exogenous change or by allowing the digital organisms themselves to influence the environment. For these reasons Avida serves as an ideal model system to explore the effects of environmental change on adaptive evolution.

The System

Avida is an evolution platform, wherein digital organisms (Avidians) can evolve in world with 3600 sites, where each site can hold at most one Avidian. Each Avidian has a genome composed of a sequence of simple computational instructions. For this experiment, the length of the genome was fixed at 100 instructions. When assembled in particular configurations, these commands perform functions like asexual replication or numerical computation. An Avidian's fitness (replication rate) can be improved by the "metabolizing" of resources in the environment, which can only be acquired by performing mathematical tasks determined by the experimenter. Many different globally available resources can be present in the environment, each paired with a particular task. During the replication of an Avidian, the mutation rate is the frequency at which substitutions of random instructions occur in the genome. Upon completion of replication, the offspring is placed in the world at a random site, supplanting any previous occupant. Each Avida run was seeded with an identical, self-replicating ancestor that initially could not perform any tasks. The unit of time is an 'update', the period for an organism to perform 30 instructions on average (one generation is approximately seven updates).

The environment within Avida is defined by the abundances of available resources and their associated computational tasks. If an organism successfully performs a task, it is rewarded with "merit" proportional to the abundance of the resource associated with the task. The more merit accumulated by an organism, the

faster the replication rate of its offspring. “Rigid” environments have unchanging resource concentrations that are not influenced by the tasks performed. The resources in a “Malleable” environment flow into the world at a fixed rate in a chemostat-like manner and are consumed by organisms when associated tasks are executed. The merit a task rewards is proportional to its abundance; hence, in a Malleable environment, the consumption of resources reduces their availability to other organisms.

In our experiments, every evolutionary run was broken into three periods, where the environment in each period was either rigid or malleable. The *Fixed* treatment uses the same rigid environment for each period. For all other treatments, the first and third periods are the same rigid environment (termed the “reference”) as the *Fixed* runs. However, the second period is a different environment (termed the “alternative”). The middle period in the *Flipped* treatment is a rigid environment with different resources (i.e., where different tasks are rewarded). The middle period of the *Negative Frequency Dependent (NFD)* treatment is a malleable environment, where the resources available in the reference environment become consumed when their task is performed. All populations were evolved for 100,000 updates (~12,000 generations).

Results

Evolution in Static and Dynamic Worlds

To assess the effect of a changing environment on evolution, a baseline for adaptation in an unchanging environment is needed. To obtain this baseline, we evolved populations in a single fixed environment (where rewards for resources did not change). At the conclusion of each evolutionary run in this *Fixed* treatment, the most abundant genotype was selected, and its line of descent to the ancestral genotype was determined. The fitness trajectory along such a line of descent is shown in Figure 2a for one *Fixed* population. For contrast, an example fitness trajectory from a population evolved in a changing environment (the *Flipped* treatment) is shown in Figure 2b. Here the shaded middle third of the evolutionary run represents exposure to an alternate environment (whereas the population evolves in the reference environment from Fig. 2a for the first and last third of the run). The blue trajectory during the middle third represents the fitness of the genotypes in the alternate environment (whereas the black trajectory gives fitness in the reference environment). In both static and dynamic worlds, nearly all mutations are beneficial with respect to the present environment. However, some mutations that are favored in the alternate environment (a lift in the blue line) would have been detrimental in the reference environment (a drop in the black line). The average of many runs of *Fixed* and *Flipped* treatments (Fig. 2c and 2d, respectively) demonstrate that fitness gains decrease over time to a plateau. However, populations exposed to an alternate environment (*Flipped* treatment) reach significantly higher fitness in the reference environment at the end of the run,

regardless of the specific resources being rewarded (Fig. 2e, Mann-Whitney test, $p=0.03$).

One possible advantage populations evolving in the *Flipped* treatment had relative to the *Fixed* treatment, is a greater availability of beneficial trajectories over the course of a run. If populations were exhausting potential beneficial mutations early in the run (leading to constraints in adaptive potential), organisms in the reference environment should have few possible beneficial mutations, but more possible beneficial mutations with respect to the alternate environment. To investigate this hypothesis, we extracted the organism in the lineage that existed at the conclusion of the first third of the run. Then we constructed every possible single mutation at every genome position and evaluated the fitness of each genotype in both the reference and alternate environment (Figure 3). After excluding mutations that were detrimental in both environments and nearly neutral in at least one environment, we found that mutations beneficial in the reference environment constituted a small minority (0.65%, Quadrants 1 and 4) of the possible mutations, while a much greater fraction of mutations were detrimental in the reference environment but beneficial in the alternate environment (99.35%, Quadrant 2). Thus, mutations beneficial in the alternate environment are often detrimental in the reference environment. This implies that the evolution in the alternate environment yields more selectively beneficial mutational steps and may lead to genotypic movement that could not have taken place in the reference environment. The

improved fitness outcome in the *Flipped* treatment over the *Fixed* treatment can be explained by a greater breadth of search of the adaptive landscape.

Evolution in Frequency Dependent Environments

Although populations can experience environmental change that is exogenous in origin (as in the *Flipped* treatment), they can also be the endogenous actors of the change as well. To model such a situation, we constructed an environment where resources were finite and consumable (whereas the reference environment had an infinite abundance of such resources). A small amount of each resource of the reference environment is continually flowing into the world, but when the resource is consumed (its associated task is performed) its availability is reduced and the fitness reward for further performance of its task decreases. In such an environment, when a phenotype increases in frequency, its relative fitness decreases as it consumes more of its resources, a situation termed negative frequency dependence (*NFD*). To examine the effect of endogenously driven change we mirrored the structure of the first set of runs. Specifically, we applied the *NFD* environment to the middle third (leaving the first and last third as the rigid reference environment). Populations evolved in the *NFD* treatment were better adapted to the reference environment than populations evolved in the *Fixed* treatment (Figure 4a and 4b, Mann-Whitney, $p=0.0075$).

From the results in the *Flipped* treatment, we know that environmental change can improve fitness. Was the effect of *NFD* merely a consequence of environmental

change that accompanies the resource consumption? To address this, we ran an additional treatment (*Paired Transplant*) where a population evolved in the environment generated by a separate *NFD* run. Specifically, we first evolved one population in the *NFD* treatment, where it affected its resources during the middle third of the run. We then evolved a second population (the focus of the *Paired Transplant* treatment) in the precise changing environment of the first. While this second population experienced fully the fluctuating resources of the first population, the latter population was completely unable to affect change in the environment itself. This isolates the effect of change alone (i.e., without the feedback between the population and environment). The mean final fitness of the *Paired Transplant* runs was found to be significantly lower than the *NFD* runs (Fig. 4c, Mann-Whitney, $p < 0.001$), indicating that environmental change alone does not account for the adaptive benefits conferred by negative frequency dependence. Thus the interchange between the environment and populations is necessary for the increased adaptation in the *NFD* treatment.

As feedback between an evolving population and its environment influences the degree of adaptation, we next examined whether the exact nature of feedback was important. We created another treatment where instead of depleting resources in the environment (as in *NFD*) organisms increased the abundance of a resource when its associated task was performed, an example of positive frequency dependence. This *Positive Frequency Dependence (PFD)* treatment is also characterized by feedback between the environment and the population during the

middle third of the run. Despite the presence of feedback, populations in the *PDF* treatment had significantly reduced mean final fitness relative to *NFD* and a trend toward reduced mean fitness relative to *Fixed* treatments (Mann-Whitney, $p < 0.001$ and $p = 0.08$ respectively), implying that the *negative* environmental feedback of *NFD* is necessary for enhanced adaptation.

One reason *NFD* may facilitate adaptation is that the population can become more diverse during the middle third of the run; that is, it occupies a broader section of the adaptive landscape. This increased distribution could yield improved outcome when the reference is revisited due to greater accessibility of adaptive peaks. As expected, during the middle third of runs, *NFD* treatments experienced an increase in genotype diversity while *PDF* runs decreased in diversity relative to the *Fixed* treatment (Fig 4j, Mann-Whitney, $p < .001$ and $p < .001$ respectively). With negative frequency dependence, prevalent genotypes reduce their own fitness, flattening the landscape and allowing for otherwise less fit genotypes to coexist. This leads to an increase in genotypic diversity.

Discussion

We observed that populations in an exogenously changing environment evolved to a higher fitness relative to populations in unchanging environments. We surmise that populations became constrained in genotype space during the first third of their evolutionary run, as only a few mutations in genotypes at this time

conferred benefits (Fig. 3). Changing the environment for the middle third of the run serves to liberate the population from the trap of a suboptimal peak. But this raises the question of why populations in general evolved to *better* positions in the reference landscape, as opposed to *different* positions. There is no immediate a priori reason why such populations wouldn't have moved into a section of the reference landscape with worse evolutionary prospects.

We postulate two reasons why populations spending time in alternative environments reach higher fitness genotypes after returning to their reference environment. If the fitness of a peak (its "height") is proportional to its basin of attraction (the number of genotypes that the peak can be reached by solely beneficial mutations), the additional movement in genotype space afforded by evolving in a different environment, should be beneficial. For instance, if a population initially evolved to a low peak, movement in the alternative environment will likely shift the population to a different/better basin of attraction of a higher peak. But if instead, the population initially evolved to a high peak, taking steps in the other environment will be less likely to cause the population to leave the larger basin, and thus it would likely return to the same high peak.

A second mechanism leading to better adaptation relies on association between fitnesses of the reference and alternative environments. If regions of higher fitness are shared between the environments (certain genotypes/traits may be favored in both), traversing the alternative environment may favor traits that are also beneficial in the reference environment. Prior work (122), finds that many of

the more complex traits favored in our reference environment (requiring a specific arrangement of instructions in the genome) could be co-opted by a few mutations to be beneficial in the alternative environment, and vice versa. However, the benefit of evolutionary exposure to a rigid, different environment will depend on the nature of the alternative landscape and the position of the population in genotype space. For instance, Kashtan et al. found that switching between environments, which shared common sub-problems, would yield better solutions than unrelated environments.

In light of these findings, we return to Wright's Shifting Balance Theory (SBT). Wright's SBT relies on drift and selection; however, these features are antagonistic, as circumstances that support drift hinder selection (e.g., small population sizes). Both Fisher and Wright realized that a change in the environment could move a population off a former peak. This requires environmental sign epistasis, where the fitness effect of a mutation is beneficial in one environment, but detrimental in another (86). Here we could imagine a recasting of SBT for a peak shift in a reference environment. A population is structured into demes, which need not be small. Suppose the whole population starts on some sub-optimal peak. If demes experience heterogeneous environments (temporally or spatially), they may be able to take different paths to different peaks. Migration between demes allows the highest fitness genotype to spread and fix globally. With such a model, dynamic environments may allow rapid evolution across rugged landscapes without the requirement for small subpopulation size.

In contrast to the previous descriptions of exogenously changed environments, malleable environments are shaped by organisms and, in turn, can affect their evolution. In our *Negative Frequency Dependent* treatment, as phenotypes utilizing novel resources increased in frequency, they became devalued. This “flattens” the landscape, as the rich (higher fitness genotypes) get poorer, which diversifies the population into different niches. This promotes the exploration of other regions of the genotype space, and indeed, in our *Negative Frequency Dependent* treatment, diversity was enhanced during the middle third. In contrast, in *Paired Transplant* runs, superior phenotypes found in this population remains superior, in this situation some rich stay rich, leading to reduced diversity. In the *Paired Transplant* of each *NFD* run, diversity does not increase, as novel phenotypes, which didn’t occur in the *NFD* run, can spread and fix. In *PNC*, an initial solution feeds back (enriches its environment) to make itself more superior. In this case the rich get richer and displace other potential phenotypes. Thus certain, malleable environments, such as *NFD*, can yield enhanced adaptation by favoring diversity and innovation.

Incorporating dynamic environments, especially malleable ones, to our models of evolutionary change may enhance our understanding of adaptation. Other natural systems that demonstrate negative frequency dependence such as resource consumption (123, 124) or host-pathogen coevolution (125–127) have stably diverse populations that may be better able to adapt to their environments. By incorporating common features of the natural world, namely, the dynamic aspect of environments and the feedback between populations and their environments, we

found that populations may be able to adapt faster. This result has implications for understanding evolution in nature, but may also suggest useful features to incorporate into evolutionary algorithms to solve engineering problems (128, 129). For instance, if the fitness function of an evolutionary algorithm discounted current high fitness solutions, alternative solutions can be more thoroughly explored, perhaps leading to a better overall outcome. Thus natural and artificial populations may yield adaptive benefits from exposure to exogenously or endogenously altered environments.

Chapter 3 Figures:




Figure 1: Adaptive Paths in Hypothetical Changing Landscapes. Here we consider two kinds of environmental change. The population is represented as one or more spheres that climb on the adaptive landscape. (a) A population initially stuck on a sub-optimal peak experiences an exogenous environmental shift, resulting in a different adaptive landscape (illustrated in green). On this new landscape, the selection allows the population to move to a new position in genotype space. Upon returning to the original environment, the population can (in this case) climb to a

new, higher peak. (b) A population initially stuck on a sub-optimal peak experiences an endogenously malleable environment. Through negative frequency dependence, the fitness of the genotype at the former peak is depressed (illustrated as a green depression in the landscape), allowing selection to favor a diversity of new genotypes (multiple spheres). When the population returns to the original environment, these diverse genotypes initiate trajectories to multiple fitness peaks.




Figure 2: Fitness in Static and Exogenously Changed Environments. (a) The fitness trajectory of the line of descent over the course of a single *Fixed* treatment run is shown. The black line denotes the fitness of the genotype in the reference environment. (b) A fitness trajectory from a single *Flipped* treatment run is shown, with the time in the alternate environment during the middle third shaded in blue. The blue line denotes the fitness of the line of descent in the alternate environment, whereas the black line gives fitness in the reference environment. Averaging 60 replicates, we next show mean fitness from the *Fixed* treatment (c) and the *Flipped* treatment (d). (e) The final mean fitness from *Fixed* and *Flipped* treatments, as measured in the reference environment. Ribbon and error whiskers denote standard error of the mean.

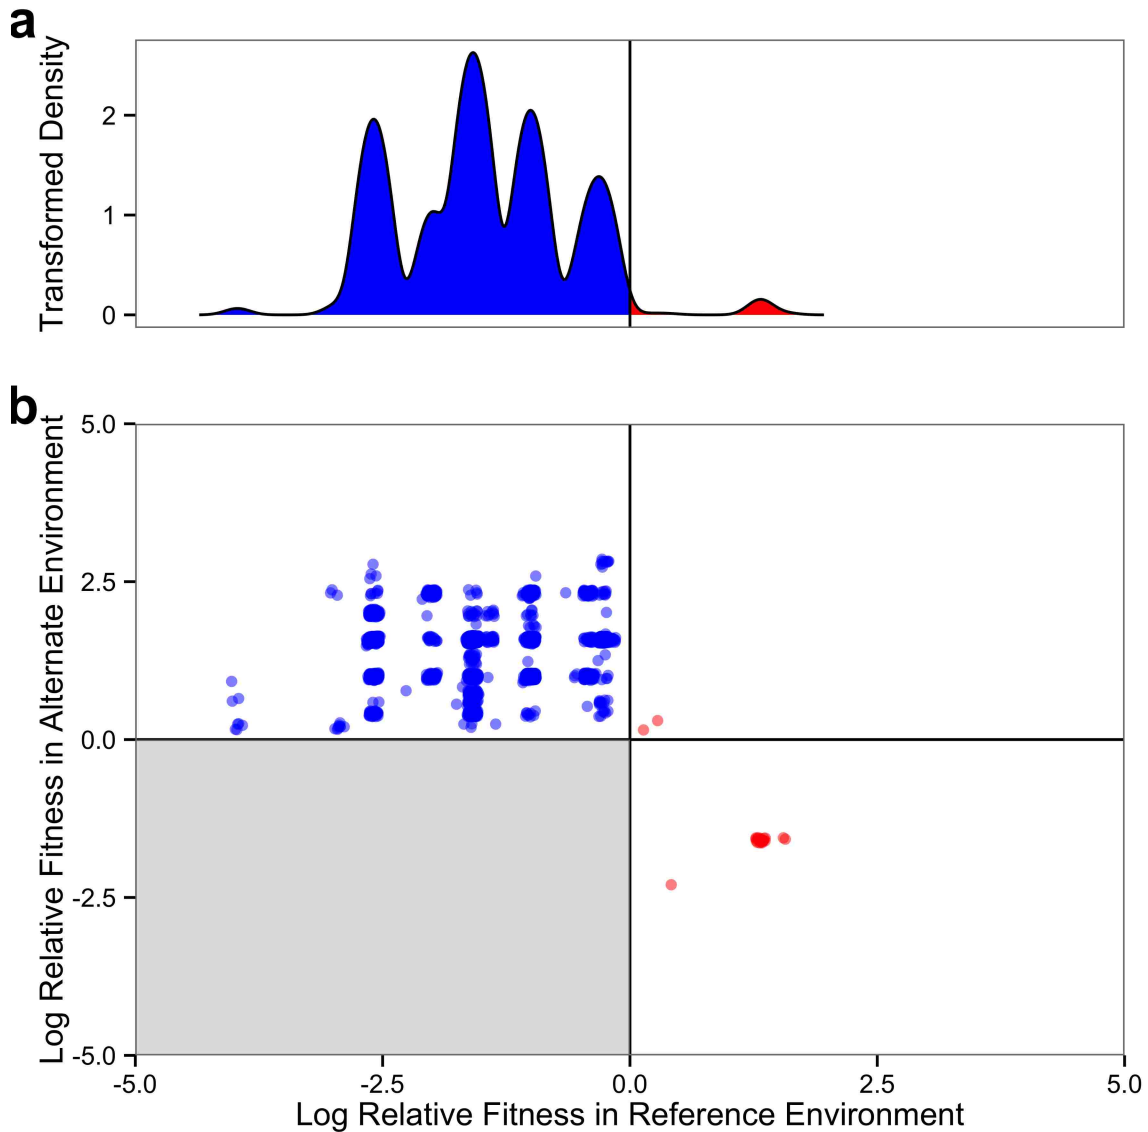


Figure 3: Mutational Distribution at the Environment Transition. We isolated a single genotype from the line of descent from a Fixed run that existed at update 33,000, which is right before an exogenous shift would have occurred in the Flipped treatment. All possible single mutations were introduced into this genotype and the fitness effects in both the reference and alternate environment were measured. In the lower plot, each point corresponds to a mutation and is positioned according to

its fitness consequences in each of the environments. Note, mutations neutral in both environments would fall on the origin. Red and blue points denote mutations beneficial and detrimental, respectively, in the reference environment. Mutations that are detrimental in both environments or have a fitness effect less than 5% in either environment are not shown. In the upper plot, the transformed density distribution ($\log(10 * \text{abundance} + 1)$) according to “ndr0” kernel smoothing) of non-excluded mutations according to fitness in the reference environment.

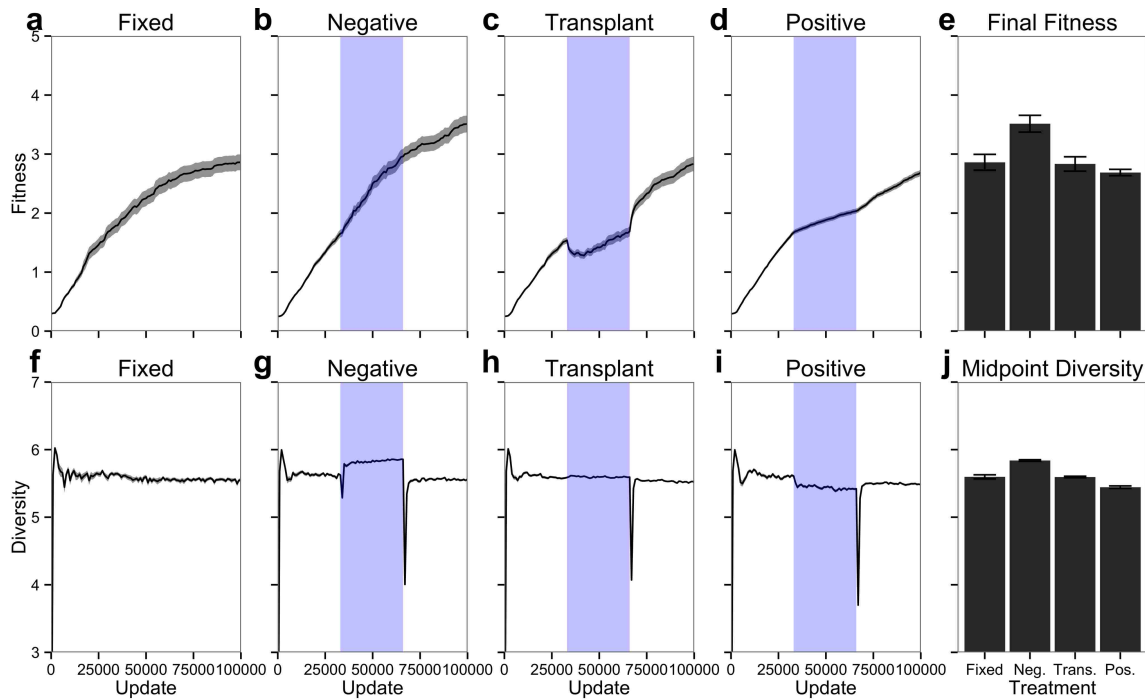


Figure 4: Fitness and Diversity in Endogenously Changed Environments.

Above, mean fitness of the line of descent from of 60 replicates of runs in the *Fixed*, *Negative Frequency Dependence*, *Paired Transplant*, and *Positive Frequency Dependence* treatments, and their final mean fitness as a bar plot. Below, the mean genotypic diversity (measured by Shannon Entropy) of the populations over time for the same treatments, and the mean diversity value at Update 50,000 (the midpoint of the run). Blue shading denotes exposure to a different environment(s) and ribbons and error whiskers denote standard error of the mean.

Supplemental Methods

Environments:

Reference environment the resources associated with the tasks Not, Nand, And, Nor, Xor and Equals are present in essentially infinite amounts, leading to no measurable depletion when the associated task is performed. In the same manner, the alternate environment (the middle third of the *Flipped* treatment) only contains the resources associated with the tasks: OrNot, Or, and AndNot. The reward for each task is only rewarded the first time an organism performs it and is equal to the complexity (number of nand instructions needed to compose) of the task. During the middle third of frequency dependent runs, the rewarded resources are the same as the reference, however, the resources have an inflow (100 units per update) and outflow rates (.01 proportion of concentration per update). The inflow and outflow rates are determined such that if an organism is the only one capable of performing a task, it will receive the same reward as the reference environment. The reward of each task is proportional to the resource concentration in a Michaelis-Menten manner. See supplemental configuration files for specifics. For the middle third of the *Negative Frequency Dependence* runs, performing a task consumed one unit of the resource, whereas in the *Positive Frequency Dependence* runs, performing a task increased the abundance of the resource associated with the task by 1. During the middle third of the Paired Transplant treatment, the resource abundance of an associated *Negative Frequency Dependence* was sampled every 1000 updates, and that environment was imposed on the evolving transplant run.

References

1. Nahum J, Harding B, Kerr B (2011) Evolution of restraint in a structured rock-paper-scissors community. *Proc Natl Acad Sci* 108:10831–10838.
2. Maynard Smith J (1964) Group Selection and Kin Selection. *Nature* 201:1145–1147.
3. Williams GC (1971) *Group Selection* (Beresford Book Serv).
4. Dawkins R (1976) *The selfish gene* (Oxford University Press, New York).
5. Okasha S Biological Altruism. Available at:
<http://plato.stanford.edu/entries/altruism-biological/> [Accessed December 11, 2010].
6. Pepper JW, Smuts BB (2002) A mechanism for the evolution of altruism among nonkin: positive assortment through environmental feedback. *Am Nat* 160:205–213.
7. Hamilton WD (1975) in *Biosocial anthropology*, pp 133–155.
8. Godfrey-Smith P, Kerr B (2009) Selection in ephemeral networks. *Am Nat* 174:906–911.
9. Fletcher JA, Doebeli M (2009) A simple and general explanation for the evolution of altruism. *Proc R Soc B Biol Sci* 276:13–19.

10. Queller DC (1992) Does population viscosity promote kin selection? *Trends Ecol Evol* 7:322–324.
11. Hamilton WD (1964) The genetical evolution of social behaviour. I. *J Theor Biol* 7:1–16.
12. Queller DC (1994) Genetic relatedness in viscous populations. *Evol Ecol* 8:70–73.
13. Wilson D, Pollock G, Dugatkin L (1992) Can altruism evolve in purely viscous populations? *Evol Ecol* 6:331–341.
14. Eshel I, Cavalli-Sforza LL (1982) Assortment of encounters and evolution of cooperativeness. *Proc Natl Acad Sci U S A* 79:1331–1335.
15. Matessi C, Jayakar SD (1976) Conditions for the evolution of altruism under darwinian selection. *Theor Popul Biol* 9:360–387.
16. Lehmann L, Keller L (2006) The evolution of cooperation and altruism – a general framework and a classification of models. *J Evol Biol* 19:1365–1376.
17. Toro M, Silio L (1986) Assortment of encounters in the two-strategy game. *J Theor Biol* 123:193–204.
18. Pepper JW (2007) Simple Models of Assortment through Environmental Feedback. *Artif Life* 13:1–9.

19. Frean M, Abraham ER (2001) Rock-scissors-paper and the survival of the weakest. *Proc R Soc B Biol Sci* 268:1323–1327.
20. Tainaka K (1993) Paradoxical effect in a three-candidate voter model. *Phys Lett* 176:303–306.
21. Tainaka K (1995) Indirect effect in cyclic voter models. *Phys Lett* 207:53–57.
22. Marsland T, Frank I eds. (2001) *Computers and Games* (Springer Berlin Heidelberg, Berlin, Heidelberg).
23. Kerr B, Riley MA, Feldman MW, Bohannan BJM (2002) Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature* 418:171–174.
24. Durrett R, Levin S (1997) Allelopathy in Spatially Distributed Populations. *J Theor Biol* 185:165–171.
25. Johnson CR, Seinen I (2002) Selection for restraint in competitive ability in spatial competition systems. *Proc R Soc Lond B Biol Sci* 269:655 –663.
26. Prado F, Kerr B (2008) The evolution of restraint in bacterial biofilms under nontransitive competition. *Evol Int J Org Evol* 62:538–548.
27. Jackson JBC, Buss L (1975) Allelopathy and spatial competition among coral reef invertebrates. *Proc Natl Acad Sci U S A* 72:5160 –5163.

28. Paquin CE, Adams J (1983) Relative fitness can decrease in evolving asexual populations of *S. cerevisiae*. *Nature* 306:368–370.
29. Sinervo B, Lively CM (1996) The rock-paper-scissors game and the evolution of alternative male strategies. *Nature* 380:240–243.
30. Clark AG, Dermitzakis ET, Civetta A (2000) Nontransitivity of sperm preference in *Drosophila*. *Evolution* 54:1030–1035.
31. Birkhead TR, Chaline N, Biggins JD, Burke T, Pizzari T (2004) Nontransitivity of paternity in a bird. *Evolution* 58:416–420.
32. Taylor DR, Aarssen LW (1990) Complex Competitive Relationships Among Genotypes of Three Perennial Grasses: Implications for Species Coexistence. *Am Nat* 136:305–327.
33. Feldgarden M, Riley MA (1998) High Levels of Colicin Resistance in *Escherichia coli*. *Evolution* 52:1270–1276.
34. Feldgarden M, Riley MA (1999) The Phenotypic and Fitness Effects of Colicin Resistance in *Escherichia coli* K-12. *Evolution* 53:1019–1027.
35. Kirkup BC, Riley MA (2004) Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo. *Nature* 428:412–414.
36. Schrag SJ, Perrot V, Levin BR (1997) Adaptation to the fitness costs of antibiotic resistance in *Escherichia coli*. *Proc R Soc B Biol Sci* 264:1287–1291.

37. Andersson DI, Levin BR (1999) The biological cost of antibiotic resistance. *Curr Opin Microbiol* 2:489–493.
38. Reynolds MG (2000) Compensatory Evolution in Rifampin-Resistant *Escherichia coli*. *Genetics* 156:1471–1481.
39. Andersson DI (2006) The biological cost of mutational antibiotic resistance: any practical conclusions? *Curr Opin Microbiol* 9:461–465.
40. Nagaev I, Bjorkman J, Andersson DI, Hughes D (2001) Biological cost and compensatory evolution in fusidic acid-resistant *Staphylococcus aureus*. *Mol Microbiol* 40:433–439.
41. Bolker BM, Pacala SW, Neuhauser C (2003) Spatial Dynamics in Model Plant Communities: What Do We Really Know? *Am Nat* 162:135–148.
42. Verhoef HA, Morin PJ (2010) *Community Ecology: Processes, Models, and Applications* (Oxford University Press US).
43. Chao L, Levin BR (1981) Structured habitats and the evolution of anticompetitor toxins in bacteria. *Proc Natl Acad Sci U S A* 78:6324–6328.
44. Van Baalen M, Rand DA (1998) The Unit of Selection in Viscous Populations and the Evolution of Altruism. *J Theor Biol* 193:631–648.
45. Kneitel JM, Chase JM (2004) Trade-offs in community ecology: linking spatial scales and species coexistence. *Ecol Lett* 7:69–80.

46. Morrison G, Barbosa P (1987) Spatial Heterogeneity, Population “Regulation” and Local Extinction in Simulated Host-Parasitoid Interactions. *Oecologia* 73:609–614.
47. Thrall PH, Burdon JJ (2002) Evolution of gene-for-gene systems in metapopulations: the effect of spatial scale of host and pathogen dispersal. *Plant Pathol* 51:169–184.
48. Boots M, Meador M (2007) Local interactions select for lower pathogen infectivity. *Science* 315:1284–1286.
49. Kerr B, Neuhauser C, Bohannan BJM, Dean AM (2006) Local migration promotes competitive restraint in a host–pathogen “tragedy of the commons.” *Nature* 442:75–78.
50. Korobeinikov A, Wake GC (1999) Global properties of the three-dimensional predator-prey Lotka-Volterra systems. *J Appl Math Decis Sci* 3:155–162.
51. Hardin G (1968) The tragedy of the commons. The population problem has no technical solution; it requires a fundamental extension in morality. *Science* 162:1243–1248.
52. Mitteldorf J, Croll DH, Ch S, Ravela U (2002) Multilevel selection and the evolution of predatory restraint. *Artif Life VIII Proc Eight Int Conf Artif Life* 8:146–152.

53. Killingback T, Bieri J, Flatt T (2006) Evolution in group-structured populations can resolve the tragedy of the commons. *Proc R Soc B Biol Sci* 273:1477–1481.
54. Eshelman CM et al. (2010) Unrestricted migration favours virulent pathogens in experimental metapopulations: evolutionary genetics of a rapacious life history. *Philos Trans R Soc B Biol Sci* 365:2503–2513.
55. Boerlijst MC, Hogeweg P (1991) Spiral wave structure in pre-biotic evolution: Hypercycles stable against parasites. *Phys Nonlinear Phenom* 48:17–28.
56. Frank SA (1994) Genetics of Mutualism: The Evolution of Altruism between Species. *J Theor Biol* 170:393–400.
57. Harcombe W (2010) Novel Cooperation Experimentally Evolved Between Species. *Evolution* 64:2166–2172.
58. Eberhard MJW (1975) The Evolution of Social Behavior by Kin Selection. *Q Rev Biol* 50:1–33.
59. Griffin AS, West SA (2002) Kin selection: fact and fiction. *Trends Ecol Evol* 17:15–21.
60. Pepper JW (2000) Relatedness in Trait Group Models of Social Evolution. *J Theor Biol* 206:355–368.
61. Wilson DS, Wilson EO (2007) Rethinking the theoretical foundation of sociobiology. *Q Rev Biol* 82:327–348.

62. Sober E, Wilson DS (1999) *Unto others: the evolution and psychology of unselfish behavior* (Harvard University Press, Cambridge, Mass.).
63. Kerr B (2009) in *Experimental Evolution: Concepts, Methods, and Applications of Selection Experiments* (University of California Press, Berkeley), pp 585–630..
1st Ed.
64. Orr HA (2005) The genetic theory of adaptation: a brief history. *Nat Rev Genet* 6:119–127.
65. Gould SJ (1990) *Wonderful life: the Burgess Shale and the nature of history* (Norton, New York).
66. Wright S (1932) in *Proc of the 6th International Congress of Genetics*, pp 356–366.
67. Gillespie JH (1984) Molecular Evolution Over the Mutational Landscape. *Evolution* 38:1116–1129.
68. Coyne JA, Barton NH, Turelli M (1997) Perspective: A Critique of Sewall Wright's Shifting Balance Theory of Evolution. *Evolution* 51:643–671.
69. Wright S (1988) Surfaces of Selective Value Revisited. *Am Nat* 131:115–123.
70. Gavrillets S (1997) Evolution and speciation on holey adaptive landscapes. *Trends Ecol Evol* 12:307–312.

71. Tanaka MM, Valckenborgh F (2011) Escaping an evolutionary lobster trap: drug resistance and compensatory mutation in a fluctuating environment. *Evol Int J Org Evol* 65:1376–1387.
72. Schenk MF, Szendro IG, Salverda MLM, Krug J, Visser JAGM de (2013) Patterns of Epistasis between Beneficial Mutations in an Antibiotic Resistance Gene. *Mol Biol Evol* 30:1779–1787.
73. Khan AI, Dinh DM, Schneider D, Lenski RE, Cooper TF (2011) Negative Epistasis Between Beneficial Mutations in an Evolving Bacterial Population. *Science* 332:1193–1196.
74. Provine WB (1989) *Sewall Wright and evolutionary biology* (University of Chicago Press, Chicago [Ill.]).
75. Kauffman S, Levin S (1987) Towards a general theory of adaptive walks on rugged landscapes. *J Theor Biol* 128:11–45.
76. Kauffman SA, Weinberger ED (1989) The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J Theor Biol* 141:211–245.
77. Østman B, Hintze A, Adami C (2012) Impact of epistasis and pleiotropy on evolutionary adaptation. *Proc R Soc B Biol Sci* 279:247–256.

78. Bergman A, Goldstein DB, Holsinger KE, Feldman MW (1995) Population structure, fitness surfaces, and linkage in the shifting balance process. *Genet Res* 66:85–92.
79. Hamming RW (1950) Error detecting and error correcting codes. *Bell Syst Tech J* 29:147–160.
80. Korona R, Nakatsu CH, Forney LJ, Lenski RE (1994) Evidence for multiple adaptive peaks from populations of bacteria evolving in a structured habitat. *Proc Natl Acad Sci* 91:9037 –9041.
81. MeInyk AH, Kassen R (2011) Adaptive Landscapes in Evolving Populations of *Pseudomonas Fluorescens*. *Evolution* 65:3048–3059.
82. Kryazhimskiy S, Rice DP, Desai MM (2012) Population Subdivision and Adaptation in Asexual Populations of *Saccharomyces Cerevisiae*. *Evolution* 66:1931–1941.
83. Weinreich DM, Delaney NF, DePristo MA, Hartl DL (2006) Darwinian Evolution Can Follow Only Very Few Mutational Paths to Fitter Proteins. *Science* 312:111 –114.
84. Lunzer M, Miller SP, Felsheim R, Dean AM (2005) The Biochemical Architecture of an Ancient Adaptive Landscape. *Science* 310:499–501.
85. Kouyos RD et al. (2012) Exploring the Complexity of the HIV-1 Fitness Landscape. *PLoS Genet* 8:e1002551.

86. Lindsey HA, Gallie J, Taylor S, Kerr B (2013) Evolutionary rescue from extinction is contingent on a lower rate of environmental change. *Nature* 494:463–467.
87. Chou H-H, Chiu H-C, Delaney NF, Segrè D, Marx CJ (2011) Diminishing Returns Epistasis Among Beneficial Mutations Decelerates Adaptation. *Science* 332:1190 –1192.
88. Poelwijk FJ, Kiviet DJ, Weinreich DM, Tans SJ (2007) Empirical fitness landscapes reveal accessible evolutionary paths. *Nature* 445:383–386.
89. MacLean RC, Perron GG, Gardner A (2010) Diminishing returns from beneficial mutations and pervasive epistasis shape the fitness landscape for rifampicin resistance in *Pseudomonas aeruginosa*. *Genetics* 186:1345–1354.
90. Visser JAGM de, Park S, Krug J (2009) Exploring the Effect of Sex on Empirical Fitness Landscapes. *Am Nat* 174:S15–S30.
91. Pepin KM, Wichman HA, Noor M (2007) Variable Epistatic Effects Between Mutations at Host Recognition Sites in ϕ X174 Bacteriophage. *Evolution* 61:1710–1724.
92. Bitbol A-F, Schwab DJ (2013) *Population subdivision with migration can facilitate evolution on rugged fitness landscapes* Available at: <http://arxiv.org/abs/1308.0278> [Accessed October 7, 2013].

93. Peck SL, Ellner SP, Gould F (1998) A Spatially Explicit Stochastic Model Demonstrates the Feasibility of Wright's Shifting Balance Theory. *Evolution* 52:1834–1839.
94. Weinreich DM, Watson RA, Chao L (2005) Perspective: Sign Epistasis and Genetic Constraint on Evolutionary Trajectories. *Evolution* 59:1165–1174.
95. Carneiro M, Hartl DL, Govindaraju DR (2010) Adaptive Landscapes and Protein Evolution. *Proc Natl Acad Sci U S A* 107:1747–1751.
96. Levin BR, Perrot V, Walker N (2000) Compensatory Mutations, Antibiotic Resistance and the Population Genetics of Adaptive Evolution in Bacteria. *Genetics* 154:985–997.
97. Björkman J, Andersson DI (2000) The cost of antibiotic resistance from a bacterial perspective. *Drug Resist Updat* 3:237–245.
98. Bossert W (1967) Mathematical optimization: are there abstract limits on natural selection? *Wistar Inst Symp Monogr* 5:35–46.
99. Koza JR, Bennett, F.H. I, Andre D, Keane MA, Dunlap F (1997) Automated synthesis of analog electrical circuits by means of genetic programming. *IEEE Trans Evol Comput* 1:109–128.
100. Oussaidène M, Chopard B, Pictet OV, Tomassini M (1997) Parallel genetic programming and its application to trading model induction. *Parallel Comput* 23:1183–1198.

101. Chang P-C, Chen S-H, Liu C-H (2007) Sub-population genetic algorithm with mining gene structures for multiobjective flowshop scheduling problems. *Expert Syst Appl* 33:762–771.
102. Alba E, Tomassini M (2002) Parallelism and evolutionary algorithms. *IEEE Trans Evol Comput* 6:443–462.
103. Skolicki Z (2005) in *In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2005* (ACM Press), pp 1295–1302.
104. Bremer E, Middendorf A, Martinussen J, Valentin-Hansen P (1990) Analysis of the tsx gene, which encodes a nucleoside-specific channel-forming protein (Tsx) in the outer membrane of Escherichia coli. *Gene* 96:59–65.
105. Nei M, Li WH (1979) Mathematical model for studying genetic variation in terms of restriction endonucleases. *Proc Natl Acad Sci U S A* 76:5269–5273.
106. Fisher RA (1930) *The Genetical Theory Of Natural Selection* (At The Clarendon Press) Available at: <http://archive.org/details/geneticaltheory031631mbp> [Accessed November 1, 2013].
107. Whitlock MC (1997) Founder Effects and Peak Shifts Without Genetic Drift: Adaptive Peak Shifts Occur Easily When Environments Fluctuate Slightly. *Evolution* 51:1044–1048.
108. Collins S, de Meaux J, Acquisti C (2007) Adaptive Walks Toward a Moving Optimum. *Genetics* 176:1089–1099.

109. Waxman D, Peck JR (1999) Sex and Adaptation in a Changing Environment. *Genetics* 153:1041–1053.
110. Meyers L.A., Bull J.J. (2002) Fighting change with change: adaptive variation in an uncertain world. *Trends Ecol Evol* 17:551–557.
111. Ancel Meyers L, Ancel FD, Lachmann M (2005) Evolution of Genetic Potential. *PLoS Comput Biol* 1:e32.
112. Parter M, Kashtan N, Alon U (2007) Environmental variability and modularity of bacterial metabolic networks. *BMC Evol Biol* 7:169.
113. Kashtan N, Noor E, Alon U (2007) Varying environments can speed up evolution. *Proc Natl Acad Sci* 104:13711–13716.
114. Alto BW, Wasik BR, Morales NM, Turner PE (2013) Stochastic Temperatures Impede Rna Virus Adaptation. *Evolution* 67:969–979.
115. Laland KN, Odling-Smee FJ, Feldman MW (1996) The evolutionary consequences of niche construction: a theoretical investigation using two-locus theory. *J Evol Biol* 9:293–316.
116. Heino M, Metz JAJ, Kaitala V (1998) The enigma of frequency-dependent selection. *Trends Ecol Evol* 13:367–370.
117. Odling-Smee FJ, Laland KN, Feldman MW (2003) *Niche construction: the neglected process in evolution* (Princeton University Press, Princeton).

118. Sterelny K (2005) Made By Each Other: Organisms and Their Environment. *Biol Philos* 20:21–36.
119. Jones CG, Lawton JH, Shachak M (1994) Organisms as Ecosystem Engineers. *Oikos* 69:373–386.
120. Laland KN, Odling-Smee FJ, Feldman MW (1999) Evolutionary consequences of niche construction and their implications for ecology. *Proc Natl Acad Sci* 96:10242–10247.
121. Ofria C, Wilke CO (2004) Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artif Life* 10:191–229.
122. Lenski RE, Ofria C, Pennock RT, Adami C (2003) The evolutionary origin of complex features. *Nature* 423:139–144.
123. Ross-Gillespie A, Gardner A, West SA, Griffin AS (2007) Frequency Dependence and Cooperation: Theory and a Test with Bacteria. *Am Nat* 170:331–342.
124. Svanbäck R, Bolnick DI (2007) Intraspecific competition drives increased resource use diversity within a natural population. *Proc R Soc B Biol Sci* 274:839–844.
125. Carius HJ, Little TJ, Ebert D (2001) Genetic Variation in a Host-Parasite Association: Potential for Coevolution and Frequency-Dependent Selection. *Evolution* 55:1136–1145.

126. Koskella B, Lively CM (2009) Evidence for Negative Frequency-Dependent Selection During Experimental Coevolution of a Freshwater Snail and a Sterilizing Trematode. *Evolution* 63:2213–2221.
127. Thrall PH et al. (2012) Rapid genetic change underpins antagonistic coevolution in a natural host-pathogen metapopulation. *Ecol Lett* 15:425–435.
128. Bäck T, Schwefel H-P (1993) An Overview of Evolutionary Algorithms for Parameter Optimization. *Evol Comput* 1:1–23.
129. Sauer U (2001) Evolutionary engineering of industrially important microbial phenotypes. *Adv Biochem Eng Biotechnol* 73:129–169.