Predicting the First Steps of Evolution in Randomly Assembled Communities

John McEnany[1] and Benjamin H. Good[2,3,4,†]

[1]*Biophysics Program, Stanford University, Stanford, CA 94305, USA*
[2]*Department of Applied Physics, Stanford University, Stanford, CA 94305, USA*
[3]*Department of Biology, Stanford University, Stanford, CA 94305, USA*
[4]*Chan Zuckerberg Biohub – San Francisco, San Francisco, CA 94158, USA*

[†]Correspondence should be addressed to: B.H.G. (bhgood@stanford.edu)

Abstract

Microbial communities can self-assemble into highly diverse states with predictable statistical properties. However, these initial states can be disrupted by rapid evolution of the resident strains. When a new mutation arises, it competes for resources with its parent strain and with the other species in the community. This interplay between ecology and evolution is difficult to capture with existing community assembly theory. Here, we introduce a mathematical framework for predicting the first steps of evolution in large randomly assembled communities that compete for substitutable resources. We show how the fitness effects of new mutations and the probability that they coexist with their parent depends on the size of the community, the saturation of its niches, and the metabolic overlap between its members. We find that successful mutations are often able to coexist with their parent strains, even in saturated communities with low niche availability. At the same time, these invading mutants often cause extinctions of metabolically distant species. Our results suggest that even small amounts of evolution can produce distinct genetic signatures in natural microbial communities.

Main Text

Microbes often live in ecologically complex communities containing hundreds of coexisting species (1–4). As residents of these communities compete with each other, they can evolve over time by acquiring mutations (5–7). These evolutionary changes can alter the ecological interactions between species, driving shifts in community composition (8–10). Conversely, the community also creates the context in which organisms evolve, by influencing the structure of the adaptive landscape (11–15). Understanding how the community influences these evolutionary paths (and vice versa) is a crucial step toward predicting and controlling microbial ecosystems.

Longstanding conceptual models suggest that community structure can impact evolution in different ways. Some models predict that the rate of evolution should decline in larger communities, as more niches are filled by other species (7, 16–18). Others have proposed that diverse communities could create new opportunities for local adaptation, by sup-

pressing key competitors or creating new niches through cross-feeding (11–14, 19, 20). These conceptual models also make different predictions about how adaptive mutations will impact their community when they invade. Some mutations will replace their parent strain, while others can encroach on other species (9, 21) or diversify into coexisting lineages (5, 6, 22). Each of these behaviors has been observed empirically, yet it remains challenging to predict which should dominate in a given community.

The source of this challenge lies in the niches inhabited by different species, and how mutations alter or move between them. While much progress has been made in small communities where the relevant niches can be explicitly defined (21–24), it is difficult to extend this approach to larger communities like the gut microbiome, where species can compete for many different combinations of resources. In this high-diversity limit, even basic questions about the effects of community structure remain unresolved: how does the benefit of a mutation depend on the diversity of the community and the metabolic overlap between its members? Do mutations primarily compete with their parent strain, or do they continue to stably diversify, as suggested by recent evidence from the gut microbiome (5, 6, 25) and some *in vitro* communities (26, 27)?

Resource competition models provide a mechanistic framework to address these questions (28–32). In these models, niches are not defined in advance, but emerge organically through differences in resource consumption (33). An extensive body of work has used this framework to investigate the process of community assembly, where species compete to colonize a new environment (28, 34–37). These model communities can recapitulate some large-scale features of natural (38–40) and experimental ecosystems (39–42). By contrast, the evolutionary dynamics that emerge from resource competition are difficult to model with traditional community assembly theory. While some studies have started to explore these effects, previous work has mostly focused on small communities (30, 31, 43) or the long-term states attained over infinite evolutionary time (30, 44). Both approaches are poorly suited for understanding how a focal species evolves in different community backgrounds, which is the scenario most accessible in experiments. To address this gap, we develop a theoretical framework for predicting the initial steps of evolution in an assembled community with many coexisting species. By extending random matrix approaches from community assembly theory, we derive analytical predictions that describe how the fitness benefits and fates of new mutations scale with the diversity and metabolic overlap of the surrounding community, enabling quantitative tests of the conceptual models above.

Modeling first-step mutations in randomly assembled communities

To study the interplay between ecological competition and new mutations in a mathematically tractable setting, we turn to a simple resource competition model (28–32, 36, 45) where microbes compete for $\mathcal{R} \gg 1$ substitutable resources that are continuously supplied by the environment (Fig. 1A). Each strain μ in the community is characterized by a resource utilization vector $\vec{r}_\mu = (r_{\mu,1}, \ldots, r_{\mu,\mathcal{R}})$, which describes how well it can grow on each of the supplied resources. While our model allows for simple forms of cross-feeding

(see below), we neglect additional factors like regulation (46), resource sequestration (47), or antagonistic interactions (48), allowing us to focus on the basic evolutionary pressures imposed by resource competition alone.

In the simplest version of our model, the concentrations of the microbial strains ($\vec{n}$) and abiotic resources ($\vec{c}$) can be described by the coupled system of equations,

$$\frac{\partial n_\mu}{\partial t} = \sum_{i=1}^{\mathcal{R}} u_i(\vec{c}) r_{\mu,i} n_\mu - \delta \cdot n_\mu, \tag{1a}$$

$$\frac{\partial c_i}{\partial t} = K_i - \sum_{\mu=1}^{\mathcal{S}} u_i(\vec{c}) r_{\mu,i} n_\mu - \delta \cdot c_i, \tag{1b}$$

where δ is the dilution rate, K_i is the external supply rate of resource i, and $u_i(\vec{c}) \cdot r_{\mu,i} \cdot n_\mu$ is the total uptake rate of resource i by species μ (SI Section 1.1). When the dilution rate is sufficiently low, or biomass sufficiently high, we can coarse-grain over time to obtain an effective model for the relative abundances of the strains (f_μ),

$$\frac{\partial f_\mu}{\partial t} = f_\mu \left[\sum_{i=1}^{\mathcal{R}} r_{\mu,i} h_i(\vec{f}\,) - 1 \right], \tag{2}$$

where time is now measured in generations and $h_i(\vec{f}\,) \equiv \left(K_i / \sum_j K_j\right) \cdot \left(\sum_\mu r_{\mu,i} f_\mu\right)^{-1}$ denotes the local availability of resource i (Fig. S1; SI Section 1.2). Simple forms of cross-feeding (39, 49) can also be included in this model by replacing the external supply rates $\vec{K}$ with an effective value that accounts for internal metabolic conversion (SI Section 1.3).

Previous work has used this model to study the process of community assembly (28, 36, 37), where $\mathcal{S}$ initial strains arrive in a new environment and form an ecologically stable community containing $\mathcal{S}^* \leq \mathcal{R}$ survivors (Fig. 1B, left). The steady-state values of $\mathcal{S}^*$ and $\vec{h}$ depend on the environmental supply rates $\vec{K}$ and the resource preferences $\vec{r}_\mu$ of the $\mathcal{S}$ initial strains. While it is difficult to measure these kinetic parameters directly, past research has shown that emergent features of large ecosystems can often be mimicked by randomly drawing the uptake rates of each strain from a common statistical distribution (28, 29, 34–36).

For concreteness, we will initially focus on the binary resource usage model from Ref. (36), in which each strain has an independent probability $\mathcal{R}_0/\mathcal{R}$ of utilizing each resource, and an overall uptake budget $X_\mu \equiv \log \sum_i r_{\mu,i}$ that is independently drawn from a Gaussian distribution (SI Section 1.4). However, our results will apply for a variety of different sampling schemes, which we verify in Figures S3, S8, and S9 and SI Section 4.

Individual realizations of this model produce assembled communities with similar numbers of surviving species, which can be predicted using methods from the physics of disordered systems (28, 29, 36; Fig. 2B; SI Section 3). In our analysis below, it will be convenient to treat the expected number of surviving species as an input parameter, and classify the

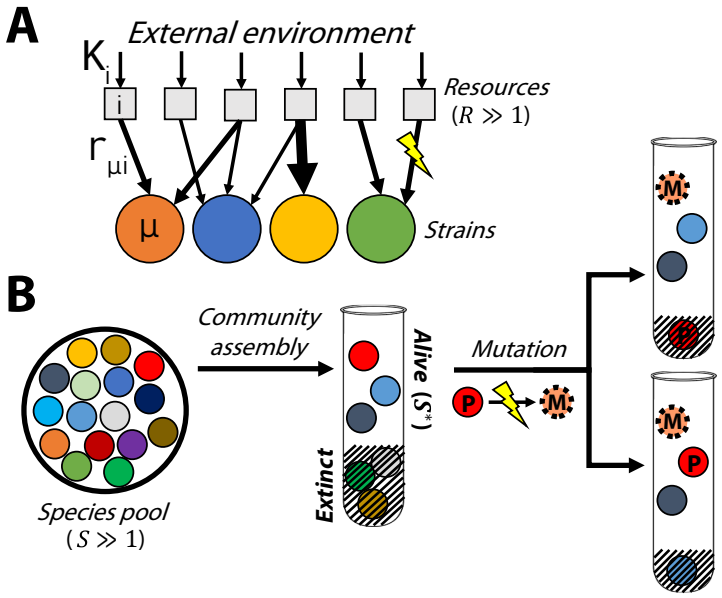

Figure 1: Modeling the first steps of evolution in a randomly assembled community that competes for substitutable resources. (A) Microbial strains compete for $\mathcal{R}$ resources that are continuously supplied by the environment at rates K_i. Each strain μ has a characteristic set of uptake rates $r_{\mu,i}$ (arrows), which can be altered by further mutations. **(B)** A local pool of $\mathcal{S}$ initial species, whose phenotypes are randomly drawn from a common statistical distribution, self-assembles into an ecological equilibrium with $\mathcal{S}^* \leq \mathcal{S}$ surviving species (left). A new mutation (M) arises in one of the surviving species (P); if the mutation provides a fitness benefit, its descendants can either replace the parent strain (top right) or stably coexist with the parent, potentially driving another species to extinction (bottom right).

assembled communities as a function of their niche saturation, $\mathcal{S}^*/\mathcal{R}$ (SI Section 3.2). This re-parameterization allows for easier generalization to other community assembly models, since communities with similar levels of niche saturation will often have similar statistical properties, even if they resulted from different sampling distributions (29, 35, 36).

To account for evolution, we model the first mutational steps that occur in a randomly assembled community (Fig. 1B). This scenario might apply to the initial phases of *in vitro* passaging experiments (12, 26), or recently colonized gut microbiomes (50, 51). Through much of our analysis, we will focus on a particularly simple class of "knock-out" mutations, where the mutant loses its ability to consume one of its resources ($r_i \to 0$). We assume that the cell can compensate for this deletion by increasing the uptake of other resources in its repertoire, but this compensation may not be perfect, corresponding to a shift in the effective uptake budget ($X \to X + \Delta X$). We also consider "knock-in" mutations, where a strain gains the ability to consume a resource (e.g. through horizontal gene transfer; 25), as well as more general changes that influence the uptake rates of multiple resources simultaneously ($\vec{r} \to \vec{r} + \Delta\vec{r}$). Following previous work (30, 32), it will be convenient to classify these mutations by their impact on the strain's total resource uptake rate $X \equiv \log \sum_i r_i$ and its normalized resource consumption strategy $\vec{\alpha} \equiv \vec{r}/\sum_i r_i$, which describes the relative effort devoted to importing each of the resources. A general mutation will therefore involve a shift in one or both of these parameters ($X \to X + \Delta X$ and $\vec{\alpha} \to \vec{\alpha} + \Delta\vec{\alpha}$).

While the phenotypes of these mutants are simple, their fitness consequences will depend on the local environment, which is shaped by the other resident strains. Using the assembled communities as a backdrop allows us to quantify how this evolutionary landscape varies with the size and composition of the surrounding community.

Distribution of fitness effects of mutations in newly assembled communities

The rate of evolution in a newly assembled community will depend on the supply of beneficial mutations. This landscape is often summarized by the local distribution of fitness effects (DFE), denoted by $\rho_\mu(s)$, which gives the relative probability that a mutation in focal strain μ will have an invasion fitness s. The shape and scale of $\rho_\mu(s)$ determine the availability of beneficial mutations, and therefore the rate of evolutionary change. Our resource competition framework allows us to ask how these features depend on the composition of the larger community.

For a community at ecological equilibrium, a mutation that arises in a resident strain μ and changes its resource uptake phenotype to a new value $X_\mu + \Delta X$ and $\vec{\alpha}_\mu + \Delta\vec{\alpha}$ will have an invasion fitness,

$$s_{\text{inv}} \approx \Delta X + \sum_i \Delta\alpha_i g_i \,, \tag{3}$$

where $g_i \equiv h_i/\overline{h} - 1$ is the excess availability of resource i relative to the ecosystem average, $\overline{h} \equiv \frac{1}{\mathcal{R}} \sum_i h_i$ (30) (SI Section 1.5). Equation (3) shows that the invasion fitness of a mutation that only affects the overall uptake budget of a strain ($\Delta\vec{\alpha} = \vec{0}$) is independent of the

surrounding community. In contrast, the benefits of mutations that change the resource consumption strategy of a strain will depend on the interactions between species, which are mediated by the values of the resource availabilities, g_i. For example, if a resource has a lower relative availability ($g_i < 0$), then it is not worth devoting a portion of its uptake budget to consume it, and a knock-out mutation for that resource should be beneficial.

Replica-theoretic calculations similar to Ref. (36) allow us to predict the distribution of the g_i for a typical assembled community when $\mathcal{R} \gtrsim \mathcal{R}_0 \gg 1$ (SI Section 3.2). In large communities, the excess resource availabilities are well-approximated by a collection of Gaussian random variables,

$$g_i \sim \left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right) \left[\frac{K_i}{\overline{K}} - 1 + Z_i \cdot C \cdot (\mathcal{S}^*/\mathcal{R})^{-1/2} \cdot \sqrt{\frac{1 - \mathcal{R}_0/\mathcal{R}}{\mathcal{R}_0}}\right], \tag{4}$$

where $\overline{K} \equiv \frac{1}{\mathcal{R}} \sum_i K_i$ is the average resource supply rate, C is an $\mathcal{O}(1)$ factor that depends on the sampling depth $\mathcal{S}/\mathcal{S}^*$, and the Z_i are uncorrelated standard Gaussians. Analogous expressions for other sampling distributions can be obtained by replacing the $\mathcal{R}_0$-dependent factor with the corresponding spread in $\vec{\alpha}$ (e.g. SI Section 4.4). This mean-field solution requires that the resources are supplied at comparable rates ($|K_i - \overline{K}| \ll \overline{K}$) and each resource is utilized by a substantial number of surviving strains ($\mathcal{R}_0\mathcal{S}^* \gg \mathcal{R}$; SI Section 3.5; Fig. S3D). When these conditions are satisfied, we can use Eqs. (3) and (4) to calculate the distribution of fitness effects $\rho_\mu(s)$ by aggregating over mutations with different values of $\Delta\vec{\alpha}$ and ΔX.

To understand how the community influences the DFE of a focal species, it is helpful to begin by considering the simplest case, where the resource supply is uniform ($K_i = \overline{K}$) and mutations have no direct costs or benefits ($\Delta X = 0$). In this case, Eqs. (3) and (4) imply that $\rho_\mu(s)$ will also follow a Gaussian distribution with mean zero and standard deviation

$$\sigma_{\text{inv}} \sim \|\Delta\vec{\alpha}\| \cdot (1 - \mathcal{S}^*/\mathcal{R}) \cdot (\mathcal{S}^*/\mathcal{R})^{-1/2} \cdot \sqrt{\frac{1 - \mathcal{R}_0/\mathcal{R}}{\mathcal{R}_0}}, \tag{5}$$

where $\|\Delta\vec{\alpha}\| \equiv \sqrt{\sum_i \Delta\alpha_i^2}$ is the magnitude of the phenotypic change produced by the mutation. Examples for knock-out and knock-in mutations (where $\|\Delta\vec{\alpha}\| \approx 1/\mathcal{R}_0$) are shown in Fig. 2C-D. Since Eq. (5) does not explicitly depend on the parent strain μ, it implies that the DFEs should be statistically similar for all strains in the community (Fig. S2). Furthermore, since the details of the mutations only enter through their overall magnitude $\|\Delta\vec{\alpha}\|$, this implies that knock-out and knock-in mutations — as well as multi-resource mutations with the same value of $\|\Delta\vec{\alpha}\|$ — will have statistically similar DFEs (Fig. 2C).

Equation (5) shows that the community influences the DFE of a focal strain primarily through the degree of niche saturation, $\mathcal{S}^*/\mathcal{R}$. The magnitude of the typical fitness effect approaches zero as $\mathcal{S}^*/\mathcal{R}$ increases (Fig. 2D), which is consistent with the idea that the rate of evolution will be slower in communities where more niches are already filled, since the establishment probability of a beneficial mutant is proportional to its fitness effect (52).

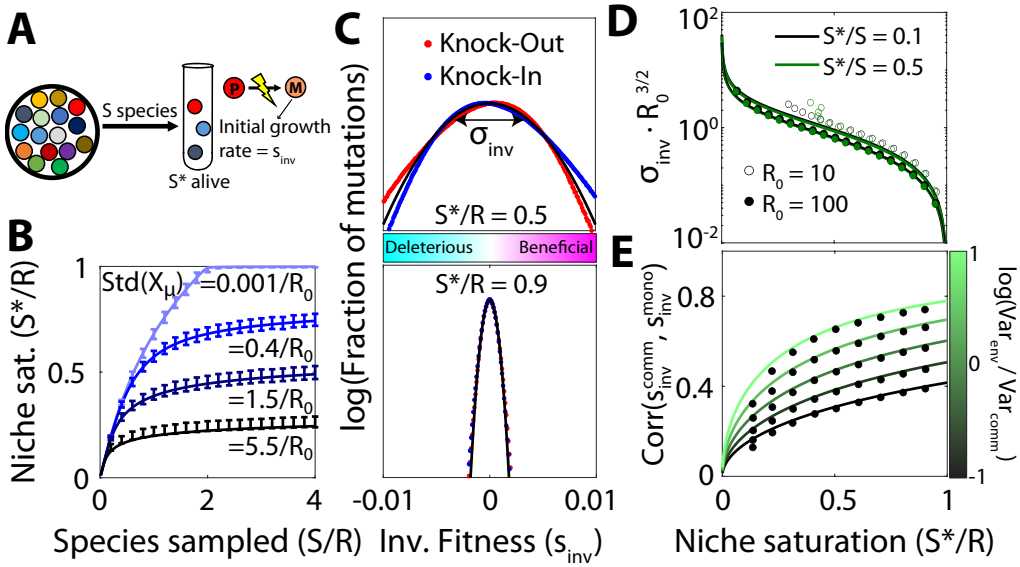

Figure 2: Distribution of fitness effects of first-step mutations as a function of community complexity. (A) As in Fig. 1, a community is assembled from $\mathcal{S}$ initial species, leaving $\mathcal{S}^*$ alive at equilibrium. The surviving species produce mutations, whose invasion fitness s_{inv} is equal to their initial relative growth rate. **(B)** Number of surviving species $\mathcal{S}^*$ as a function of the sampling depth $\mathcal{S}$ and the standard deviation of the total uptake budget among sampled species, $\text{Std}(X_\mu)$. Curves show theory predictions from SI Section 3.1, while points show means and standard deviations over 10^3 simulation runs with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, and uniform resource supply. **(C)** Distribution of fitness effects of knock-out (and knock-in) mutations with $\Delta X = 0$ in communities with two different levels of niche saturation ($\mathcal{S}^*/\mathcal{R}$). Black curve shows the theoretical predictions from Eq. (5), while dots represent a histogram over all possible strategy mutations in 10^3 simulation runs using the same parameters as panel B, with $\mathcal{S}^*/\mathcal{S} = 0.1$. **(D)** The width σ_{inv} of the distribution of fitness effects in panel C as a function of niche saturation, for various values of sampling permissivity $\mathcal{S}^*/\mathcal{S}$ and per-species resource usage $\mathcal{R}_0$. Curves show the theoretical predictions from Eq. (5), while the dots show the average over 10^3 simulation runs. **(E)** Pearson correlation between the fitness effect of a mutation in the community and its fitness effect in monoculture, for different values of niche saturation and scaled variation in resource supply rates, $\text{Var}_{\text{env}}/\text{Var}_{\text{comm}} \equiv \left[\text{Var}(K)/\overline{K}^2\right] \cdot \mathcal{R}_0/(1 - \mathcal{R}_0/\mathcal{R})$. Curves show the theoretical predictions from SI Section 4.1, while points show the average over all mutations in 10^3 simulated communities with parameters the same as panel C.

However, since the mean of the DFE is still centered at $s = 0$, the overall fraction of beneficial mutations remains constant as the niche saturation increases (Fig. 2C), even though a smaller number are likely to survive genetic drift. This implies that surviving organisms are not necessarily at a local evolutionary optimum, where any change to their resource consumption strategy tends to be deleterious.

However, this symmetry between the frequency of beneficial and deleterious mutations critically depends on the assumption of perfect trade-offs ($\Delta X = 0$), which will not always hold in practice. For example, a beneficial mutation could halt the production of an enzyme which is used for metabolizing a low-availability resource, while leaving the expression of other enzymes in that now-defunct pathway intact – resulting in a net cost to pure fitness. If all mutations carried such a direct cost ($\Delta X < 0$), then Eq. (3) implies that the entire DFE would shift to the left by a constant amount $-|\Delta X|$. If this shift is larger than the typical spread of the DFE ($|\Delta X| \gg \sigma_{\text{inv}}$), then the beneficial tail of $\rho_\mu(s)$ will approach an exponential shape, whose height and width will both strongly decline with the degree of niche saturation $\mathcal{S}^*/\mathcal{R}$. Thus, the availability of beneficial mutations can sensitively depend on the genetic architecture of the strain's resource consumption rates.

Similar results apply when the resources are supplied at different rates ($K_i \neq \overline{K}$; SI Section 4.1; Fig. S7). In this case, the fitness effect of a mutation will depend on both the community context and the external environment, as encoded by the resource availabilities in Eq. (4). The overall magnitude of the environmental contribution declines as niche saturation $\mathcal{S}^*/\mathcal{R}$ increases, consistent with previous work showing that larger communities self-organize to "shield" their member species from the external environment (30, 32, 36). Interestingly, however, Eq. (4) shows that the *relative* impact of $\vec{K}$ on the fitness effect a mutation actually increases with the degree of niche saturation, so the external environment can still exert an influence on the overall direction of the fitness landscape. This effect is strikingly illustrated when we compare the fitness effect of each mutation with its expected value in the absence of the community (Fig. 2E). We find that the correlations between the two DFEs can be substantial when the coefficient of variation in K_i exceeds a critical value, $\sqrt{\mathcal{R}/\mathcal{S}^*\mathcal{R}_0} \lesssim 1$. This indicates that environmental shielding is often incomplete: the net direction of natural selection can be preserved across community backgrounds, even when inter-specific competition exerts a strong effect on the local resource availabilities.

Ecological diversification in large communities

While the invasion fitness of a mutation describes its initial growth rate, a successful variant will eventually reach a size where it starts to impact the other members of the community. Some of these mutants will eventually replace their parent strain, due to the principle of competitive exclusion (28). Others can stably coexist with their parents by exploiting a different ecological niche (30). How does the frequency of these *in situ* ecological diversification events depend on the composition of the surrounding community?

We can analyze the probability of mutant-parent coexistence in our model by recasting

it as the outcome of two correlated assembly processes. First, an initial ecosystem E_0 is formed through our standard community assembly process (Fig. 1B, left). Then, a second ecosystem is formed when one of the surviving species in E_0 produces a beneficial mutant M, and the combined community $E_0 + M$ is allowed to reach its new ecological equilibrium (Fig. 1B, right). The latter event requires that the parent strain has a positive relative abundance in the first ecosystem ($f_P^{E_0} > 0$), and that the mutant survives in the second ecosystem ($f_M^{E_0+M} > 0$). The probability that the mutant coexists with its parent can then be expressed as a conditional probability,

$$\mathbb{P}_{\text{coex}} = \mathbb{P}[f_P^{E_0+M} > 0 \,\big|\, f_P^{E_0} > 0,\, f_M^{E_0+M} > 0], \tag{6}$$

which averages over the random resource uptake rates in the initial community, as well as the random effect of the adaptive mutation.

The correlated assembly process in Eq. (6) is challenging to analyze with existing methods (37), since the mutant and parent phenotypes are closely related to each other. Fortunately, we will show that we can often approximate the coexistence probability by considering a third community assembly process, in which the mutant and parent are introduced simultaneously with the other species. This "simultaneous assembly" approximation differs from the two-ecosystem model in Eq. (6), since the final community can contain "rescued" species that would not have survived in E_0 before the mutant strain invaded. However, simulations indicate that these differences result in only small corrections to the coexistence probability across a broad range of parameters (SI Section 2.1, Fig. S4), so this is often a reasonable approximation.

When the simultaneous assembly approximation holds, we can evaluate the coexistence probability by extending the replica-theoretic calculations used for Eq. (4) (SI Section 3.3). We find that coexistence can be expressed as the probability that the invasion fitness of the mutant lies below a critical fitness threshold,

$$s_{\text{coex}}(f_P) \equiv \sigma_{\text{inv}} \cdot (f_P \mathcal{S}^*) \cdot (\mathcal{S}^*/\mathcal{R})^{-1/2} \cdot \frac{\sqrt{2}\|\Delta\vec{\alpha}\|}{\|\Delta\vec{\alpha}_{\text{comm}}\|}\,. \tag{7}$$

where f_P is the relative abundance of the parent strain, $\|\Delta\vec{\alpha}\|^2 \equiv \sum_i \Delta\alpha_i^2$ is the phenotypic effect size of the mutation, and $\|\Delta\vec{\alpha}_{\text{comm}}\|^2 \approx 2(1 - \mathcal{R}_0/\mathcal{R})/\mathcal{R}_0$ is the equivalent spread between random pairs of strains in the community. Averaging over the abundance of the parent strain (SI Section 3.4) yields a corresponding expression for the coexistence probability as an integral over the DFE in Fig. 2C,

$$\mathbb{P}_{\text{coex}} \approx \frac{\int_0^\infty s\rho(s)e^{-s/\bar{s}_{\text{coex}}}}{\int_0^\infty s\rho(s)ds} \tag{8}$$

where $\bar{s}_{\text{coex}}$ is the coexistence threshold for a typical genetic background ($f_P \sim 1/\mathcal{S}^*$). For mutations with modest phenotypic effects, such as a single resource knock-out ($\|\Delta\vec{\alpha}\| \sim 1/\mathcal{R}_0$), the magnitude of $\bar{s}_{\text{coex}}$ will be much smaller than the typical spread in the DFE in Fig. 2C. This implies that only mutations with anomalously low invasion fitnesses will have an appreciable chance of coexisting with their parent (Fig. 3A).

For a mutation with perfect trade-offs ($\Delta X = 0$), the integral in Eq. (8) reduces to a simple form,

$$\mathbb{P}_{\text{coex}} \approx \frac{1}{\mathcal{S}^*/\mathcal{R}} \cdot \frac{2\|\Delta\vec{\alpha}\|^2}{\|\Delta\vec{\alpha}_{\text{comm}}\|^2}\,, \tag{9}$$

which depends on the niche saturation $\mathcal{S}^*/\mathcal{R}$ and the phenotypic effect size of the new mutation.

Equation (9) shows that the coexistence probability is largest for small values of $\mathcal{S}^*/\mathcal{R}$ (Fig. 3A), consistent with the idea that diversification is easier when there are many open niches left to exploit, as in an adaptive radiation (20). Interestingly, however, we find that the coexistence probability does not vanish as the community approaches saturation ($\mathcal{S}^*/\mathcal{R} \to 1$), but instead plateaus at a nonzero value. This is true even in fully saturated communities ($\mathcal{S}^* = \mathcal{R}$), where other species must be driven to extinction when the successful mutant invades.

In this regime, the coexistence probability is most strongly determined by the phenotypic effect size of the mutation. A mutation that changes the resource consumption strategy by an infinitesimal amount ($\|\Delta\vec{\alpha}\| \to 0$) can never coexist with its parent, while a mutation with $\|\Delta\vec{\alpha}\| \sim \|\Delta\vec{\alpha}_{\text{comm}}\|$ has a much higher probability of coexistence. Single-resource knockouts fall between these two extremes, with the coexistence probability scaling inversely with $\mathcal{R}_0$ (the average number of resources utilized per strain), rather than the total number of resources consumed by the community. This suggests that *in situ* ecological diversification can continue to occur even in large communities containing many species and resources (Fig. 3B and Fig. S5).

We see that simulation results generally support the theoretical predictions in Eq. (9), but start to exceed this value for communities very close to full saturation. The source of this deviation coincides with the breakdown of our "simultaneous assembly" approximation above, which allowed the invading mutant to "rescue" some species that had previously gone extinct. The deviations disappear if extinct species are allowed to re-invade (Fig. S4), suggesting that competition from rescued species plays a significant role in our theory for extremely high levels of niche saturation. This rescuing effect could be relevant in some natural settings, e.g. if the species pool is maintained in a separate spatial niche with frequent migration back into the community. Regardless, our results in Eq. (9) provide a lower bound on the coexistence probability in large ecosystems, suggesting that diversification could be even more common under some conditions.

In reality, most mutations that change an organism's resource consumption strategy are unlikely to be perfect trade-offs. More generally, we find that mutations that carry a direct cost or benefit ΔX change the coexistence probability in Eq. (9) by the factor

$$\frac{\mathbb{P}_{\text{coex}}(\Delta X)}{\mathbb{P}_{\text{coex}}(\Delta X = 0)} \sim \begin{cases} \frac{\sigma_{\text{inv}}}{\sqrt{2\pi}\Delta X} e^{-\Delta X^2/2\sigma_{\text{inv}}^2} & \text{if } \Delta X \gg \sigma_{\text{inv}}, \\ 1 & \text{if } |\Delta X| \ll \sigma_{\text{inv}}, \\ \Delta X^2/\sigma_{\text{inv}}^2 & \text{if } \Delta X < 0, |\Delta X| \gg \sigma_{\text{inv}}. \end{cases} \tag{10}$$

which shows that direct costs or benefits begin to exert an effect when ΔX becomes compa-

rable to σ_{inv}. However, the direction of this effect is somewhat counterintuitive: strategy mutations that carry a direct fitness cost ($\Delta X < 0$) are more likely to coexist with their parent, provided that their overall invasion fitness is still positive. Conversely, mutations with strong direct benefits ($\Delta X > 0$) are more likely to drive their parent strain to extinction (Fig. 3B). These differences arise because nonzero values of ΔX change the typical invasion fitnesses of successful mutants. The coexistence probability in Eq. (8) is highly sensitive to these shifts, as mutants with an invasion fitness below $\bar{s}_{\text{coex}}$ are much more likely to coexist with their parent (Fig. 3A,D). This is a reflection of the fact that diversification is driven by a narrow range of mutants with fitness high enough to establish but small enough to coexist.

We can gain some intuition for this effect by considering a loss-of-function mutation for a resource that is currently overutilized by the population ($g_i < 0$). Equation (2) shows that the growth rate of this variant (relative to its ancestor) is given by $\sim g_i(t)$. As the mutant grows in abundance, it begins to replace its parent, which tends to reduce the overall utilization of the target resource ($\partial_t g_i > 0$). If g_i reaches zero before the parent strain goes extinct, then the mutant will coexist with its parent, having lost its initial growth advantage. This coexistence is most likely to happen if the parent strain has high abundance, or if $|g_i|$ (and therefore s_{inv}) is initially small (Fig. S6).

The same argument applies for more complex mutations which affect multiple resources (Fig. S5). It also explains why mutations with small phenotypic effects (e.g. a pure fitness mutation with no strategy change) cannot coexist with their parent. If the mutant and parent have near-identical resource consumption strategies, the mutant's invasion produces very little negative feedback on its growth rate relative to its parent, making it likely the parent strain will be driven to extinction (Fig. S5B). Similar logic applies for extensions of our basic model, including non-uniform resource supply rates and alternative community assembly schemes (Figures S3, S8, and S9; SI Section 4). Together, these results suggest that *in situ* diversification could be common even in large and saturated communities, particularly for mutants on abundant backgrounds with lower-than-expected invasion fitnesses.

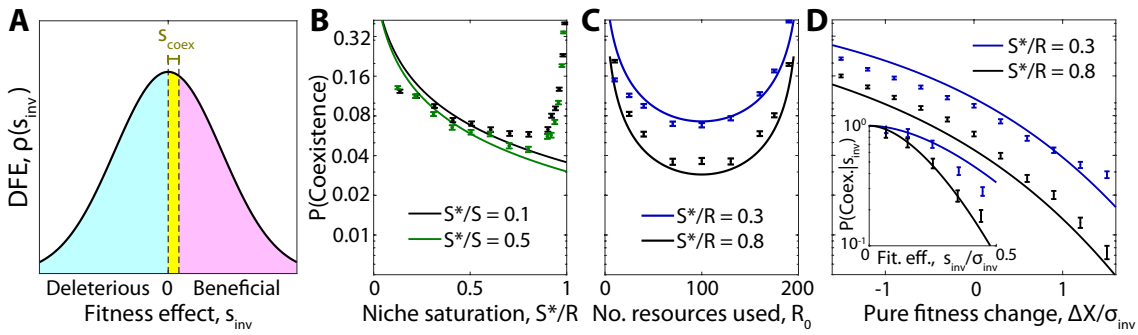

Figure 3: Ecological diversification in large communities. (A) Schematic showing the subset of mutations that are able to stably coexist with their parent strain. Coexistence occurs in the yellow region where the invasion fitness is below a critical threshold s_{coex}, which depends on the abundance of the parent and the phenotypic effect size of the mutation. **(B-D)** Probability that a successful knock-out mutant coexists with its parent strain as a function of (B) the niche saturation $\mathcal{S}^*/\mathcal{R}$, (C) the typical number of resources used per species $\mathcal{R}_0$, and (D) the change in overall uptake budget of the mutant ΔX. Inset shows the dependence on the total invasion fitness s_{inv}. In all panels, curves show the theory predictions from SI Sections 3.3 and 3.4, while points show means and standard errors over 10^4 simulation runs with base parameters $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and uniform resource supply.

Successful mutations drive extinctions in other niches

In addition to displacing their parents, successful mutants can also drive other species in the community to extinction. This is particularly evident in saturated communities ($\mathcal{S}^* = \mathcal{R}$), where competitive exclusion implies that the invasion of a new strain must be accompanied by extinction of at least one other. Figure 4B shows that the number of extinctions steadily rises with the degree of niche saturation, exceeding 1% of the initial community ($\sim$1-10 species) for many combinations of parameters. Moreover, these extinctions are not completely independent of each other, since they are somewhat overdispersed compared to a simple Poisson expectation (Fig. 4B, inset). Previous work has suggested that even distantly related strains could exhibit correlated dynamics if their resource consumption strategies are anomalously similar to each other (26). Could such hidden metabolic similarity be responsible for driving the extinction events above?

We tested this idea by examining the number of resources that were jointly consumed by the invading and displaced strains (a proxy for their overall metabolic similarity). Interestingly, we found that the number of resources shared with the mutant was not substantially higher for the displaced species, and was comparable to a randomly drawn species from the larger community (Fig. 4C). This suggests that the extinction events in Fig. 4 cannot be explained by traditional measures of niche overlap (33). Rather, successful mutants can displace species outside of their apparent niche, even when they stably coexist with a more metabolically similar parent.

Since extinctions were not well-predicted by their overall metabolic similarity to the mutant, we conjectured that these displaced species may possess other features that render them vulnerable to extinction. For example, low-abundance species may be less well-adapted to the current environment, and thus sensitive to perturbations like the invasion of a mutant. Consistent with this hypothesis, we found that the extinction probability for very low-abundance species is much higher than the community average, approaching $\sim$50% at the highest levels of niche saturation (Fig. 4D). Furthermore, although the displaced species are metabolically distant from the invading mutant, we find they are more likely to share the mutant's strategy for the resource targeted by the mutation. Displaced species are more likely to use the resource gained by a successful "knock-in" mutation, and are correspondingly less likely to use the resource lost in a successful knock-out strain (Fig. 4E). These results illustrate that in high-dimensional ecosystems, the invasion of a new mutant can have a small impact on a diverse range of metabolic strategies. If a resident strain is maladapted enough to already be on the edge of extinction, the invasion of a mutant can be enough of a perturbation to displace it from the community. Thus, the abundance of an organism might often be a better predictor of its fate than its apparent metabolic niche.

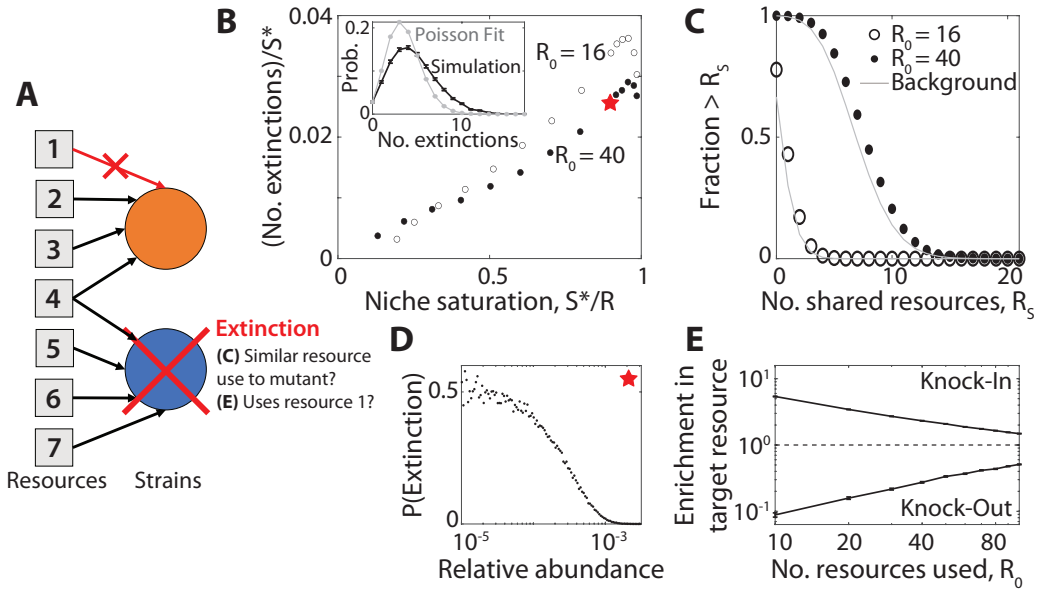

Figure 4: Successful mutations drive extinctions of metabolically distant species. (**A**) Schematic showing the extinction of an unrelated species (blue) after a beneficial knockout mutation (orange) invades. In this example, the displaced species and the mutant share one common resource, but not the one targeted by the knock-out mutation. (**B**) Average number of extinctions after the invasion of a successful mutant, as a function of niche saturation $\mathcal{S}^*/\mathcal{R}$; parent strains are excluded from the extinction tally. Inset: full distribution of extinctions for the starred parameters, compared to a Poisson distribution with the same fraction of zero counts. Points denote the averages over 10^4 simulation runs with the same base parameters as Fig. 3. (**C**) Distribution of the number of resources jointly utilized by the displaced species and the invading mutant (parent strains excluded). Points denote the results of simulations with $\mathcal{S}^*/\mathcal{R} = 0.9$. Gray curves show the analogous background distribution between the mutant and all other species in the community, regardless of whether they become extinct. (**D**) Probability of extinction as a function of initial relative abundance for the starred point in panel B. (**E**) The fold change in probability that the displaced species uses the same resource targeted by the mutant (i.e., the resource being knocked in or out), relative to the background distribution of resource use in the population. Points denote means and standard errors for knock-in and knock-out mutations as a function of $\mathcal{R}_0$ for $\mathcal{S}^*/\mathcal{R} = 0.8$, with the the remaining parameters the same as panel B.

Robustness of ecosystem to subsequent mutations

Our analysis above focused on the first wave of mutations arising in a newly assembled community, where the initial states could be predicted using existing community assembly theory (28, 29, 35, 36). However, subsequent waves of evolution could eventually drive the ecosystem away from this well-characterized initial state (30). To assess the robustness of our results under the acquisition of further mutations, we simulated successive waves of mutations using a generalization of the approach in Fig. 3. We considered the simplest case where resident strains could generate knock-out mutations on any of the resources they currently utilized. We also assumed that the dynamics were mutation-limited (30, 31), so that the community relaxes to a well-defined steady state between each successive mutation. We continued this process until one of the surviving strains had accumulated 10 mutations in total, which typically corresponded to 100-200 successful mutations in the larger community.

We first asked how these subsequent waves of mutations altered the genetic structure of their community. While the total number of surviving strains decreased slightly over time (eventually stabilizing at an intermediate value), the number of strains related through *in situ* diversification events increased approximately linearly over the same time window (Fig. 5B). This indicates that the ecological diversification events in Fig. 3 continue to occur at a high rate even after additional mutations have accumulated, more quickly than individual branches of diversified lineages go extinct. Nonetheless, the abundance trajectories in Figure 5A indicate that these extinction events among close relatives are not uncommon. For example, Lineage 1 seeds eight ecological diversification events over the course of the simulation, but only three of these strains are alive at the end of the simulation. On average, we find that newly diversified lineages coexist for $\sim$40 mutational steps before one of them goes extinct (Fig. 5C). Nonetheless, longer coexistence is possible: one of the diversification events in Fig. 5A is maintained for $>$150 mutational steps, enough time for multiple additional mutations to accumulate within each lineage. Thus, coexisting mutant-parent pairs are often maintained through further evolutionary perturbations, even as the total number of species stabilizes over time.

Our first-step analysis showed that mutations in more abundant strains were more likely to coexist with their parent. This prediction is borne out by our multi-step simulations as well: while some highly abundant lineages seeded many *in situ* ecological diversification events (e.g. Lineage 1 in Fig. 5A), the typical surviving lineage experienced no more than one (e.g. Lineage 2), and many species went extinct without diversifying at all (e.g. Lineage 3). Consistent with our earlier predictions, lineages at low abundance were more prone to extinction when a new mutant invaded, while coexistence events tended to happen to high-abundance lineages (Figure 5A, right). The combination of these factors creates a "rich get richer" effect where high-abundance species diversify at the expense of driving low-abundance species to extinction (30).

Finally, we asked how the accumulation of mutations shifts the landscape of beneficial mutations from the initial "assembled" state in Figs. 2 and 3. We found that the distribution of fitness effects does not dramatically differ between the start and the end of the

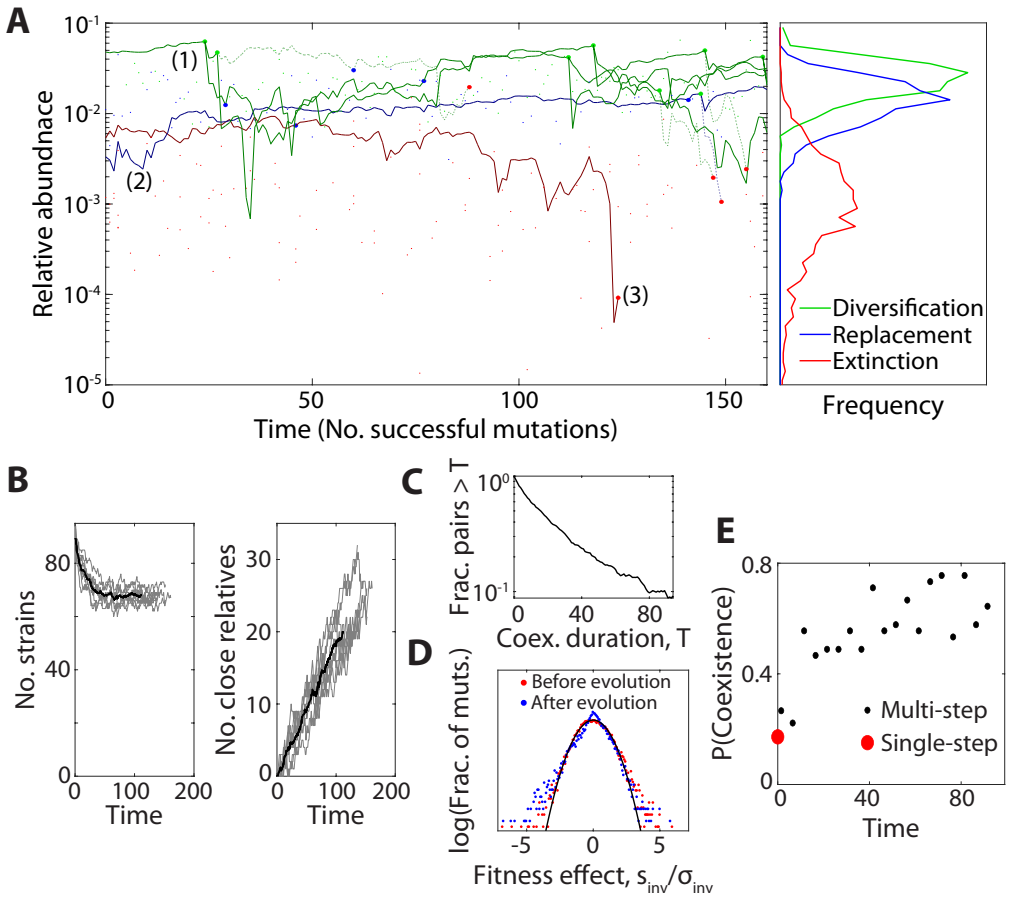

Figure 5: Mutation and diversification over longer evolutionary timescales. (**A**) Left: an example simulation showing the step-wise accumulation of ∼100 adaptive knock-out mutations in a community with $\mathcal{R} = 100$, $\mathcal{R}_0 = 20$, $\mathcal{S}^*/\mathcal{R} = 0.9$, $\mathcal{S}^*/\mathcal{S} = 0.1$. Solid lines denote the abundance trajectories of 3 example lineages. Large points indicate extinction events in these lineages (red), diversification events (green), and mutation events that displace their parent strain (blue) in the highlighted lineages, while smaller points indicate analogous events for other species in the community. Dashed lines illustrate offshoots of the highlighted lineages that eventually went extinct. Right: Relative abundances of strains when they experienced mutation, diversification and extinction events, respectively. Lines denote histograms aggregated over 9 simulation runs. (**B**) Left: Total number of surviving strains over time. Grey lines denote replicate simulations for the same parameters in panel A, while their average is shown in black. Right: Total number of strains related to another surviving strain through one or more *in situ* diversification events. (**C**) The probability of mutant-parent coexistence being maintained (i.e., both lineages surviving) as a function of time since initial divergence, over ten simulations. (**D**) The distribution of fitness effects at the start of the simulation (red) and after 90 accumulated mutations (blue), compared to the first-step predictions from Fig. 2 (black). (**E**) The probability that a mutation event leads to stable diversification as a function of time. Black points denote binned values aggregated over nine replicate simulations, while red point denotes the analogous result for 10^4 first-step simulations.

simulation, although its shape changes slightly from a Gaussian distribution toward a two-tailed exponential (Fig. 5D). In contrast, the impacts of successful mutations exhibit larger changes: the proportion of mutation events which result in mutant-parent coexistence increases substantially with evolutionary time, rising from about 20% to 80% over the course of the simulation (Fig. 5E). This shift cannot be explained by changes in the overall number of species or the availability of beneficial mutations, but instead deviates entirely from the replica-theoretic predictions that describe our initial state. Thus, while many our qualitative conclusions continue to hold on longer evolutionary timescales, the accumulation of mutations can lead to quantitative differences from our first-step analysis above. In both cases, the evolved communities exhibit distinct genetic signatures compared to their purely assembled counterparts, which could motivate future tests of *in situ* evolution.

Discussion

Large ecological communities apply complex evolutionary pressures to their resident species, leading to controversy about how these organisms' evolutionary trajectories are affected by the background community (17). Here, we address this challenge by developing a theoretical framework for predicting the first steps of evolution in large, randomly assembled communities that compete for substitutable resources. These "mean-field" results hold when surviving strains consume many resources, and each resource is consumed by many surviving strains. Under these conditions, our results provide a mechanistic approach for understanding how the fitness benefits and fates of new mutations should scale with the diversity and metabolic overlap of the surrounding community.

Our results show that the supply of beneficial mutations does not necessarily run out in larger communities – as expected in the simplest models of niche filling – but rather that the benefits of these mutations will systematically decline with the degree of niche saturation ($\mathcal{S}^*/\mathcal{R}$). We also find that the fitness effects of mutations can be broadly correlated with the external environment, even in large communities where internal resource concentrations are broadly shielded from external environmental shifts (32, 36). These distributions of fitness effects can be measured in modern experiments, e.g. by performing barcoded fitness assays in communities of different size (12, 50).

Our finding that successful mutants often coexist with their parents is reminiscent of empirical observations from the gut microbiome, where recently diverged strains differing by only a handful of mutations appear to stably coexist within their host (5, 6). While spatial structure could also contribute to this coexistence (27, 53), our model demonstrates that *in situ* diversification can continue to occur even in a well-mixed environment when most niches are already filled. Similar behavior has also been observed in generalized Lotka-Volterra models with spatiotemporal chaos (54). This suggests that ongoing diversification may be a generic feature of large microbial communities, providing an alternative mechanism for the "diversity begets diversity" effect (17, 55, 56) that does not rely on explicit cross-feeding interactions.

Beyond diversification, Figure 5 demonstrates that successful mutants can also drive dis-

tantly related species to extinction. This behavior is consistent with recent empirical observations in the human gut microbiome (8) and strain-swapping experiments in synthetic gut communities (57). It could also contribute to the extinction events observed in community passaging experiments (58). The fact that the displaced species are metabolically diverged from the invading mutants creates obstacles for inferring these interactions from metabolomics measurements (59) or metabolic reconstructions (60). Our findings suggest that future efforts should instead focus on how the invading mutant impacts low-abundance species more generally, and whether they utilize the specific resource(s) targeted by the mutation. These results align with recent work emphasizing the importance of collective interactions in shaping microbial community dynamics (61).

Here, we focused on the simplest possible resource competition model, neglecting important factors like spatial structure (62), metabolic regulation (46), and more general forms of cross-feeding (49, 63), which are all thought to play key roles in natural microbial communities. We also considered the simplest regime of evolutionary dynamics that neglect competition between simultaneously occurring mutations. These clonal interference effects can enhance diversity by allowing strains to temporarily evade competitive exclusion (31). On the other hand, competition between lineages also selects for more strongly beneficial mutations (64), which we predict are less likely to coexist with their parent. Our results provide a baseline for incorporating these effects, which will be crucial for understanding how large microbial communities will evolve.

References

[1] Huttenhower C, Gevers D, Knight R, Abubucker S, Badger JH, Chinwalla AT, et al. Structure, function and diversity of the healthy human microbiome. Nature. 2012;486(7402):207–214. doi:10.1038/nature11234.

[2] Yang J, Pu J, Lu S, Bai X, Wu Y, Jin D, et al. Species-Level Analysis of Human Gut Microbiota With Metataxonomics. Frontiers in Microbiology. 2020;11. doi:10.3389/fmicb.2020.02029.

[3] Bai Y, Müller DB, Srinivas G, Garrido-Oter R, Potthoff E, Rott M, et al. Functional overlap of the Arabidopsis leaf and root microbiota. Nature. 2015;528(7582):364–369. doi:10.1038/nature16192.

[4] Sunagawa S, Coelho LP, Chaffron S, Kultima JR, Labadie K, Salazar G, et al. Structure and function of the global ocean microbiome. Science. 2015;348(6237). doi:10.1126/science.1261359.

[5] Roodgar M, Good BH, Garud NR, Martis S, Avula M, Zhou W, et al. Longitudinal linked-read sequencing reveals ecological and evolutionary responses of a human gut microbiome during antibiotic treatment. Genome Res. 2021;31(8):1433–1446.

[6] Zhao S, Lieberman TD, Poyet M, Kauffman KM, Gibbons SM, Groussin M, et al. Adaptive Evolution within Gut Microbiomes of Healthy People. Cell Host & Microbe. 2019;25(5):656–667.e8. doi:https://doi.org/10.1016/j.chom.2019.03.007.

[7] Scheuerl T, Hopkins M, Nowell RW, Rivett DW, Barraclough TG, Bell T. Bacterial adaptation is constrained in complex communities. Nature Communications. 2020;11(1):754. doi:10.1038/s41467-020-14570-z.

[8] Good BH, Rosenfeld LB. Eco-evolutionary feedbacks in the human gut microbiome. Nature Communications. 2023;14(1). doi:10.1038/s41467-023-42769-3.

[9] Meroz N, Tovi N, Sorokin Y, Friedman J. Community composition of microbial microcosms follows simple assembly rules at evolutionary timescales. Nature Communications. 2021;12(1). doi:10.1038/s41467-021-23247-0.

[10] Barber JN, Sezmis AL, Woods LC, Anderson TD, Voss JM, McDonald MJ. The evolution of coexistence from competition in experimental co-cultures of Escherichia coli and Saccharomyces cerevisiae. The ISME Journal. 2020;15(3):746–761. doi:10.1038/s41396-020-00810-z.

[11] Evans R, Beckerman AP, Wright RCT, McQueen-Mason S, Bruce NC, Brockhurst MA. Eco-evolutionary Dynamics Set the Tempo and Trajectory of Metabolic Evolution in Multispecies Communities. Current Biology. 2020;30(24):4984–4988.e4. doi:https://doi.org/10.1016/j.cub.2020.09.028.

[12] Venkataram S, Kuo HY, Hom EFY, Kryazhimskiy S. Mutualism-enhancing mutations dominate early adaptation in a two-species microbial community. Nature Ecology & Evolution. 2023;7(1):143–154. doi:10.1038/s41559-022-01923-8.

[13] Lawrence D, Fiegna F, Behrends V, Bundy JG, Phillimore AB, Bell T, et al. Species interactions alter evolutionary responses to a novel environment. PLoS Biol. 2012;10(5):e1001330.

[14] Barroso-Batista J, Pedro MF, Sales-Dias J, Pinto CJG, Thompson JA, Pereira H, et al. Specific Eco-evolutionary Contexts in the Mouse Gut Reveal Escherichia coli Metabolic Versatility. Current Biology. 2020;30(6):1049–1062.e7. doi:10.1016/j.cub.2020.01.050.

[15] Tawk C, Lim B, Bencivenga-Barry NA, Lees HJ, Ramos RJF, Cross J, et al. Infection leaves a genetic and functional mark on the gut population of a commensal bacterium. Cell Host & Microbe. 2023;31(5):811–826.e6. doi:10.1016/j.chom.2023.04.005.

[16] Schluter D, Pennell MW. Speciation gradients and the distribution of biodiversity. Nature. 2017;546(7656):48–55. doi:10.1038/nature22897.

[17] Madi N, Chen D, Wolff R, Shapiro BJ, Garud NR. Community diversity is associated with intra-species genetic diversity and gene loss in the human gut microbiome. eLife. 2023;12:e78530. doi:10.7554/eLife.78530.

[18] Hall JPJ, Harrison E, Brockhurst MA. Competitive species interactions constrain abiotic adaptation in a bacterial soil community. Evolution Letters. 2018;2(6):580–589. doi:10.1002/evl3.83.

[19] Calcagno V, Jarne P, Loreau M, Mouquet N, David P. Diversity spurs diversification in ecological communities. Nature Communications. 2017;8(1):15810. doi:10.1038/ncomms15810.

[20] Meyer JR, Kassen R. The effects of competition and predation on diversification in a model of adaptive radiation. Nature. 2007;446(7134):432–435. doi:10.1038/nature05599.

[21] Debray R, Conover A, Zhang X, Dewald-Wang EA, Koskella B. Within-host adaptation alters priority effects within the tomato phyllosphere microbiome. Nature Ecology & Evolution. 2023;7(5):725–731. doi:10.1038/s41559-023-02040-w.

[22] Knope ML, Forde SE, Fukami T. Evolutionary History, Immigration History, and the Extent of Diversification in Community Assembly. Frontiers in Microbiology. 2012;2. doi:10.3389/fmicb.2011.00273.

[23] Nadeau CP, Farkas TE, Makkay AM, Papke RT, Urban MC. Adaptation reduces competitive dominance and alters community assembly. Proceedings of the Royal Society B: Biological Sciences. 2021;288(1945):20203133. doi:10.1098/rspb.2020.3133.

[24] Gül E, Abi Younes A, Huuskonen J, Diawara C, Nguyen BD, Maurer L, et al. Differences in carbon metabolic capacity fuel co-existence and plasmid transfer between Salmonella strains in the mouse gut. Cell Host & Microbe. 2023;doi:https://doi.org/10.1016/j.chom.2023.05.029.

[25] Frazão N, Seixas E, Barreto HC, Mischler M, Güleresi D, Gordo I. Massive lateral gene transfer under strain coexistence in the gut. bioRxiv. 2023;doi:10.1101/2023.09.25.559333.

[26] Goyal A, Bittleston LS, Leventhal GE, Lu L, Cordero OX. Interactions between strains govern the eco-evolutionary dynamics of microbial communities. eLife. 2022;11:e74987. doi:10.7554/eLife.74987.

[27] Jin X, Yu FB, Yan J, Weakley AM, Dubinkina V, Meng X, et al. Culturing of a complex gut microbial community in mucin-hydrogel carriers reveals strain- and gene-associated spatial organization. Nature Communications. 2023;14(1). doi:10.1038/s41467-023-39121-0.

[28] Advani M, Bunin G, Mehta P. Statistical physics of community ecology: a cavity solution to MacArthur's consumer resource model. J Stat Mech. 2018;2018(3):033406.

[29] Cui W, Marsland R, Mehta P. Effect of Resource Dynamics on Species Packing in Diverse Ecosystems. Phys Rev Lett. 2020;125:048101. doi:10.1103/PhysRevLett.125.048101.

[30] Good B, Martis S, Hallatschek O. Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources. PNAS. 2018;115(44). doi:10.1073/pnas.1807530115.

[31] Amicone M, Gordo I. Molecular signatures of resource competition: Clonal interference favors ecological diversification and can lead to incipient speciation. Evolution. 2021;75(11):2641–2657. doi:https://doi.org/10.1111/evo.14315.

[32] Posfai A, Taillefumier T, Wingreen NS. Metabolic Trade-Offs Promote Diversity in a Model Ecosystem. Phys Rev Lett. 2017;118:028103. doi:10.1103/PhysRevLett.118.028103.

[33] Letten A, Ke PJ, Fukami T. Linking modern coexistence theory and contemporary niche theory. Ecological Monographs. 2016;87. doi:10.1002/ecm.1242.

[34] Barbier M, Arnoldi JF, Bunin G, Loreau M. Generic assembly patterns in complex ecological communities. Proc Natl Acad Sci U S A. 2018;115(9):2156–2161.

[35] Cui W, Marsland R, Mehta P. Diverse communities behave like typical random ecosystems. Phys Rev E. 2021;104(3-1):034416.

[36] Tikhonov M, Monasson R. Collective Phase in Resource Competition in a Highly Diverse Ecosystem. Physical Review Letters. 2017;118(4). doi:10.1103/physrevlett.118.048103.

[37] Tikhonov M, Monasson R. Innovation rather than improvement: A solvable high-dimensional model highlights the limitations of scalar fitness. J Stat Phys. 2018;172(1):74–104.

[38] Barbier M, de Mazancourt C, Loreau M, Bunin G. Fingerprints of high-dimensional coexistence in complex ecosystems. Phys Rev X. 2021;11(1).

[39] Marsland R, Cui W, Mehta P. A minimal model for microbial biodiversity can reproduce experimentally observed ecological patterns. Sci Rep. 2020;10(1):3308.

[40] Ho PY, Good BH, Huang KC. Competition for fluctuating resources reproduces statistics of species abundance over time across wide-ranging microbiotas. Elife. 2022;11.

[41] Goldford JE, Lu N, Bajić D, Estrela S, Tikhonov M, Sanchez-Gorostiaga A, et al. Emergent simplicity in microbial community assembly. Science. 2018;361(6401):469–474.

[42] Dal Bello M, Lee H, Goyal A, Gore J. Resource-diversity relationships in bacterial communities reflect the network structure of microbial metabolism. Nat Ecol Evol. 2021;5(10):1424–1434.

[43] Doebeli M. A model for the evolutionary dynamics of cross-feeding polymorphisms in microorganisms. Population Ecology. 2002;44(2):59–70. doi:10.1007/s101440200008.

[44] Caetano R, Ispolatov Y, Doebeli M. Evolution of diversity in metabolic strategies. eLife. 2021;10:e67764. doi:10.7554/eLife.67764.

[45] MacArthur R. Species Packing, and What Competition Minimizes. Proceedings of the National Academy of Sciences. 1969;64(4):1369–1371. doi:10.1073/pnas.64.4.1369.

[46] Pacciani-Mori L, Giometto A, Suweis S, Maritan A. Dynamic metabolic adaptation can promote species coexistence in competitive microbial communities. PLOS Computational Biology. 2020;16(5):e1007896. doi:10.1371/journal.pcbi.1007896.

[47] Liu Y, Hu J, Lee H, Gore J. Complex ecosystems lose stability when resource consumption is out of niche. biorxiv. 2023;doi:10.1101/2023.11.30.569477.

[48] Niehus R, Oliveira NM, Li A, Fletcher AG, Foster KR. The evolution of strategy in bacterial warfare via the regulation of bacteriocins and antibiotics. eLife. 2021;10:e69756. doi:10.7554/eLife.69756.

[49] Mehta P, Marsland R. Cross-feeding shapes both competition and cooperation in microbial ecosystems. arXiv. 2021;doi:10.48550/ARXIV.2110.04965.

[50] Wong DPGH, Good BH. Quantifying the adaptive landscape of commensal gut bacteria using high-resolution lineage tracking. bioRxiv. 2023;doi:10.1101/2022.05.13.491573.

[51] Yilmaz B, Mooser C, Keller I, Li H, Zimmermann J, Bosshard L, et al. Long-term evolution and short-term adaptation of microbiota strains and sub-strains in mice. Cell Host & Microbe. 2021;29(4):650–663.e9. doi:10.1016/j.chom.2021.02.001.

[52] Gillespie J. Population Genetics: A Concise Guide. Johns Hopkins University Press; 2004.

[53] Mark Welch JL, Hasegawa Y, McNulty NP, Gordon JI, Borisy GG. Spatial organization of a model 15-member human gut microbiota established in gnotobiotic mice. Proceedings of the National Academy of Sciences. 2017;114(43). doi:10.1073/pnas.1711596114.

[54] Mahadevan A, Pearce MT, Fisher DS. Spatiotemporal ecological chaos enables gradual evolutionary diversification without niches or tradeoffs. eLife. 2023;12. doi:10.7554/elife.82734.

[55] Estrela S, Diaz-Colunga J, Vila JCC, Sanchez-Gorostiaga A, Sanchez A. Diversity begets diversity under microbial niche construction. bioRxiv. 2022;doi:10.1101/2022.02.13.480281.

[56] San Roman M, Wagner A. Diversity begets diversity during community assembly until ecological limits impose a diversity ceiling. Molecular Ecology. 2021;30(22):5874–5887. doi:https://doi.org/10.1111/mec.16161.

[57] Wang M, Osborn LJ, Jain S, Meng X, Weakley A, Yan J, et al. Strain dropouts reveal interactions that govern the metabolic output of the gut microbiome. Cell. 2023;186(13):2839–2852.e21. doi:10.1016/j.cell.2023.05.037.

[58] Morella NM, Weng FCH, Joubert PM, Metcalf CJE, Lindow S, Koskella B. Successive passaging of a plant-associated microbiome reveals robust habitat and host genotype-dependent selection. Proc Natl Acad Sci U S A. 2020;117(2):1148–1159.

[59] Ho PY, Nguyen TH, Sanchez JM, DeFelice BC, Huang KC. Resource competition predicts assembly of in vitro gut bacterial communities. biorxiv. 2022;doi:10.1101/2022.05.30.494065.

[60] Levy R, Borenstein E. Metabolic modeling of species interaction in the human microbiome elucidates community-level assembly rules. Proceedings of the National Academy of Sciences. 2013;110(31):12804–12809. doi:10.1073/pnas.1300926110.

[61] Chang CY, Bajić D, Vila JCC, Estrela S, Sanchez A. Emergent coexistence in multispecies microbial communities. Science. 2023;381(6655):343–348. doi:10.1126/science.adg0727.

[62] Ghosh OM, Good BH. Emergent evolutionary forces in spatial models of luminal growth and their application to the human gut microbiota. Proceedings of the National Academy of Sciences. 2022;119(28). doi:10.1073/pnas.2114931119.

[63] Goyal A, Maslov S. Diversity, Stability, and Reproducibility in Stochastically Assembled Microbial Ecosystems. Physical Review Letters. 2018;120(15). doi:10.1103/physrevlett.120.158102.

[64] Good BH, Rouzine IM, Balick DJ, Hallatschek O, Desai MM. Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. Proceedings of the National Academy of Sciences. 2012;109(13):4950–4955. doi:10.1073/pnas.1119910109.

[65] Doebeli M. Adaptive Dynamics: A Framework for Modeling the Long-Term Evolutionary Dynamics of Quantitative Traits. In: The Adaptive Landscape in Evolutionary Biology. Oxford University Press; 2013.Available from: `https://doi.org/10.1093/acprof:oso/9780199595372.003.0014`.

Acknowledgments

We thank Daniel Fisher for useful discussions, and Sophie Walton, James Ferrare, Zhiru Liu, Daniel Wong, and Avaneesh Narla for comments and feedback on the manuscript. This work was supported in part by the Alfred P. Sloan Foundation (FG-2021-15708), NIH NIGMS Grant No. R35GM146949, and a Terman Fellowship from Stanford University. B.H.G. is a Chan Zuckerberg Biohub – San Francisco Investigator.

Author contributions: Conceptualization: J.M. and B.H.G.; theory and methods development: J.M. and B.H.G..; analysis: J.M. and B.H.G.; writing: J.M. and B.H.G.

Competing interests: None declared.

Code Availability: Source code for community assembly simulations, numerical calcula-

tions and figure generation are available at Github (`https://github.com/jdmcenany/First_Step_Muts`).

Supplementary Information

Contents

Page

1 Model Outline **13**

1.1 General Ecological Dynamics 13

1.2 Coarse-Grained Ecological Dynamics 14

1.3 Incorporating Simple Forms of Cross-feeding 16

1.4 Defining the Local Species Pool 17

1.5 Ecological Equilibrium as Optimization 17

1.6 Analytic Approximation of Ecological Equilibrium 18

2 Replica-Theoretic Analysis of First-Step Invasion **19**

2.1 Single-Ecosystem Approximation of First-Step Invasion 19

2.2 Partition Function and the Replica Trick 21

2.3 Introducing First-Step Mutations 24

2.4 Saddle Point Approximations 25

2.5 Joint Distribution of Δ_P and Δ_M 27

2.6 Parent-Mutant Coexistence Probability 30

3 Scaling Analysis of First-Step Mutations **31**

3.1 Sampling Depth and Fitness Gauge 31

3.2 Niche Saturation and σ_{inv} 32

3.3 Scaling Analysis of $\mathbb{P}_{\text{coex}}$ 34

3.4 Dependence of $\mathbb{P}_{\text{coex}}$ on Invasion Fitness and Abundance 36

3.5 Self-Consistency of Replica Approximations 37

4 Extensions of Model Assumptions **39**

4.1 Non-Uniform Resource Supply 39

4.2 Variation in Number of Metabolized Resources 40

4.3 Metabolic Trade-offs in Sampling 42

4.4 Continuous Resource Usage 42

4.5 Specialist Community Assembly 43

5 Simulations and Numerics **44**

5.1 Community Simulations 44

5.2 Theory Numerics 45

List of Figures

S1 Explicit community dynamics with dilution 3

S2 Distribution of fitness effects for a single organism 4

S3 Alternative consumer resource models 5

S4 Simulations with returning species 6

S5 Mutant-parent coexistence for alternate strategy mutations 7

S6 Dependence of mutant-parent coexistence on invasion fitness and background abundance 8

S7 Distribution of fitness effects and mutant-parent coexistence probabilities under heterogeneous resource supply. 9

S8 Fitness effects and mutant-parent coexistence for different types of metabolic trade-offs 10

S9 Fitness effects and mutant-parent coexistence with a range of specialist and generalist resource strategies 11

S10 Checks for numerical stability 12

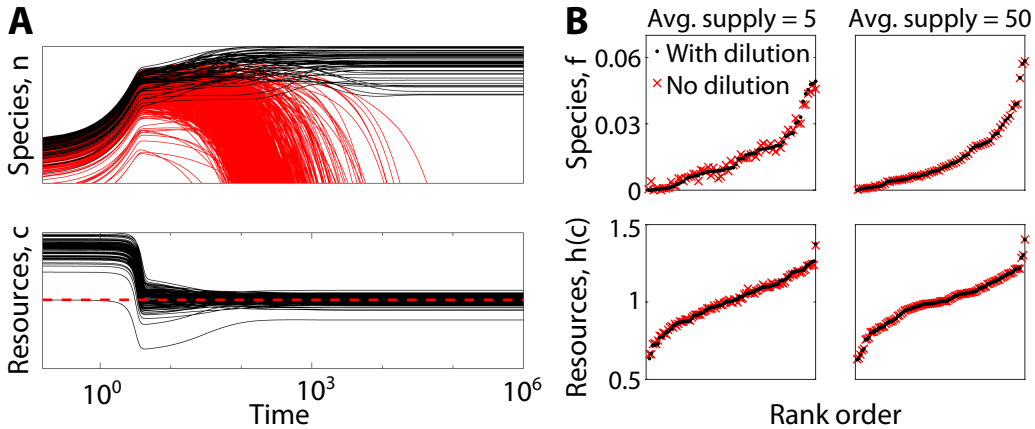

Figure S1: (A) Absolute abundance of species (top) and resource concentrations (bottom) over time on a log-log scale, for an example community simulated by explicitly integrating Eq. (S1a). Red curves are species which go extinct at equilibrium; the red dashed line in the lower plot shows the expected average value of c at equilibrium. **(B)** Equilibrium species relative abundances and resource availabilities $h(c)$, obtained from explicit simulation (black dots) and finding the equilibrium through optimization, ignoring the dilution term (red crosses). Predictions are shown for two values of average resource supply $\overline{K}$. Simulations were performed with $\mathcal{R} = 100$, $\mathcal{R}_0 = 20$, $\mathcal{S}/\mathcal{R} = 10$, $\mathrm{Std}(X_\mu) = 0.5/\mathcal{R}_0$, and $\mathrm{Var}(K)/\overline{K}^2 = 0.1$. The uptake function $u(c)$ was taken as a Monod function, $u_i(c_i) = u_{\max} c_i/(c_{1/2} + c_i)$, with $u_{\max} = 4$ and $c_{1/2} = 1$.

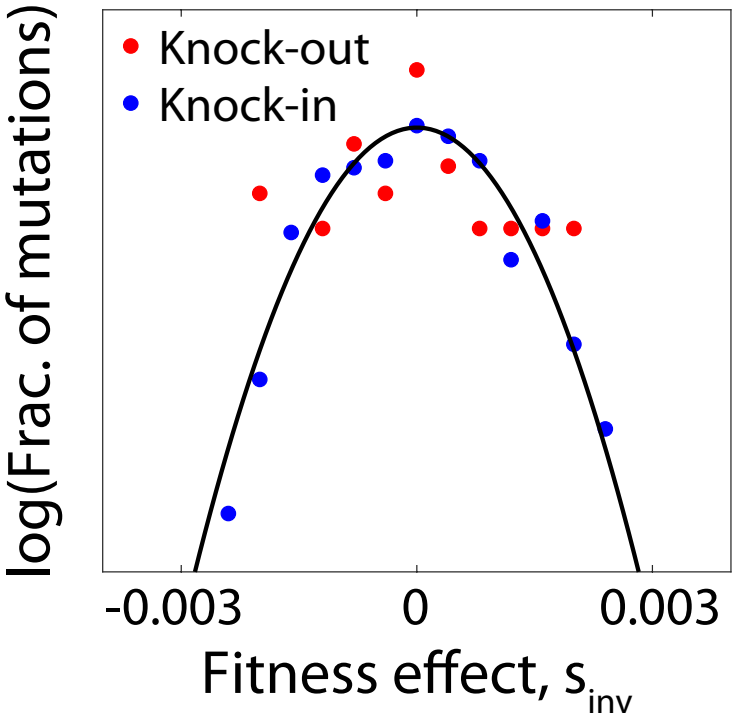

Figure S2: Distribution of invasion fitnesses for knock-out and knock-in strategy mutations of a single organism within a single sampled community. Black curve shows Gaussian theory prediction, while dots are histogram values over all possible strategy mutations. Simulation was run for $\mathcal{S}^*/\mathcal{R} = 0.8, \mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and uniform resource supply.

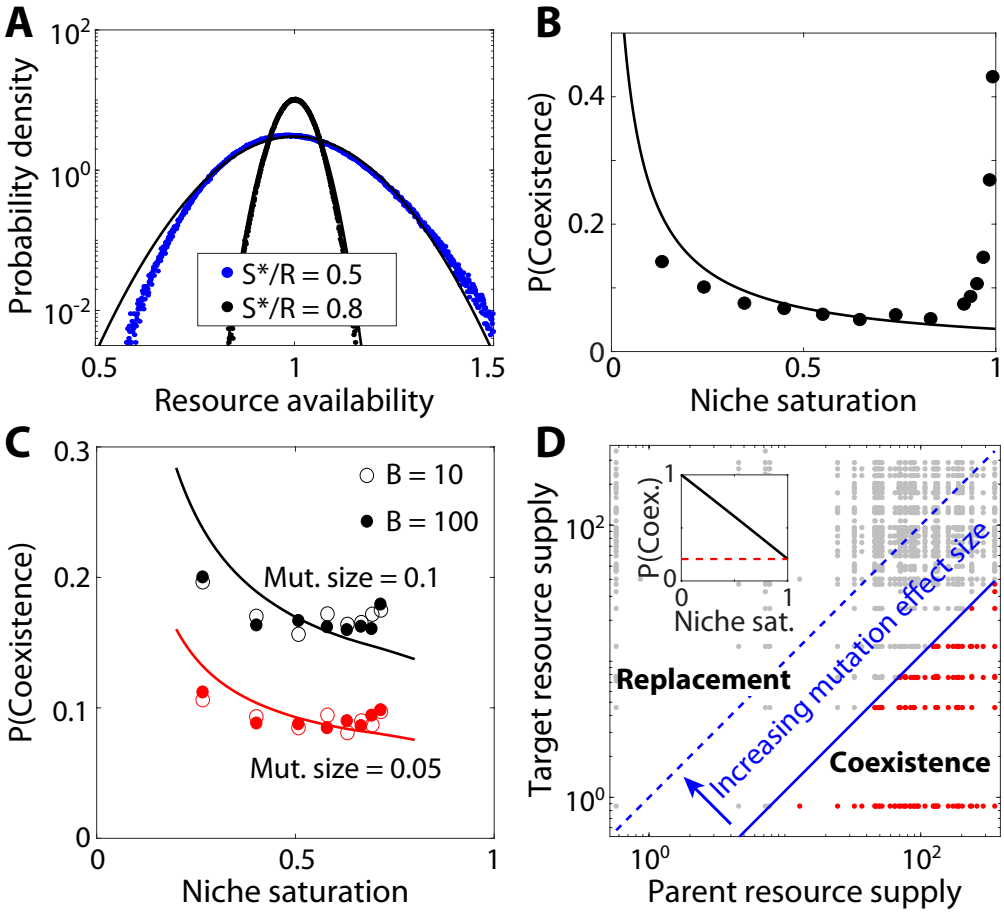

Figure S3: Our main conclusions about mutation in assembled communities can be generalized to other consumer-resource sampling structures. **(A)** Distribution of resource availabilities h_i for an ecosystem assembled with strategy vectors drawn from the Dirichlet distribution as described in SI Section 4.4. Black curves show theoretical prediction from the binary resource use case studied in this paper, while points show simulated results. Simulations were run with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$ and $\mathcal{S}^*/\mathcal{S} = 0.1$ for uniform resource supply. **(B)** Probability of mutant-parent coexistence for knock-out strategy mutations in this Dirichlet assembled community, with the same parameters as in (A). Black curve shows the theoretical prediction for binary resource use. **(C)** Mutant-parent coexistence probability for a community where all species consume all resources at similar levels (SI Section 4.4), for different mutation effect sizes Γ. Curves show theoretical prediction from the binary model, with an effective value of $\mathcal{R}_0$ corresponding to the shape parameter B of the Dirichlet distribution. Other parameters were $\mathrm{Std}(X_\mu) = 1/B\mathcal{R}$ and $\mathcal{R} = 200$. **(D)** Conditions for mutant-parent coexistence in a specialist community (SI Section 4.5). Each point corresponds to a beneficial mutation in a 50-member community, with red points showing mutant-parent coexistence. Axes indicate the supply of the resource the parent consumes, and the supply of the "target" resource the mutant gains access to. The blue line is the predicted boundary between coexistence and replacement for mutations with effect size $\gamma_{\mathrm{spec}} = 0.1$. Inset: Mutant-parent coexistence probability as a function of niche saturation; the red dashed line corresponds to a saturated community, as shown in the rest of (D).

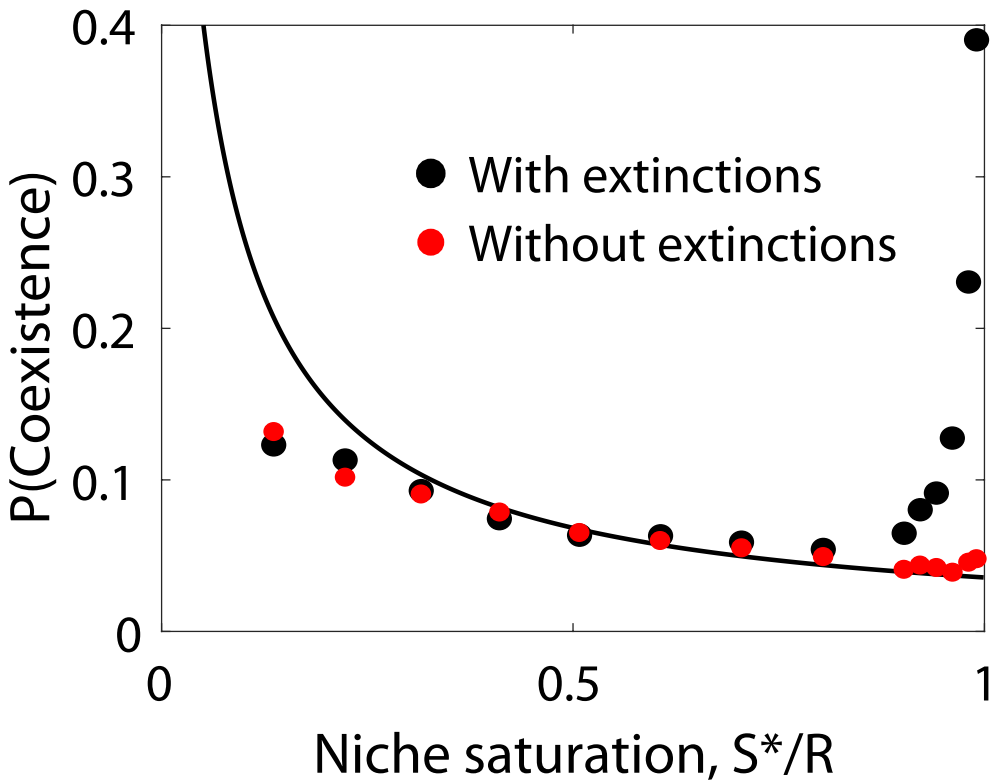

Figure S4: Probability of coexistence as a function of niche saturation, with simulations using the "simultaneous assembly" approximation where mutant invasion is modeled using a single assembly process (SI Section 2.1). Red points show the simultaneous assembly approximation, which allows reappearance of extinct species. Black points show simulations where extinct species were disallowed from returning (as in the rest of our study); this change results in coexistence probability increasing at high saturation. Simulations were performed with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, and $\mathcal{S}^*/\mathcal{S} = 0.1$ using knock-out mutations.

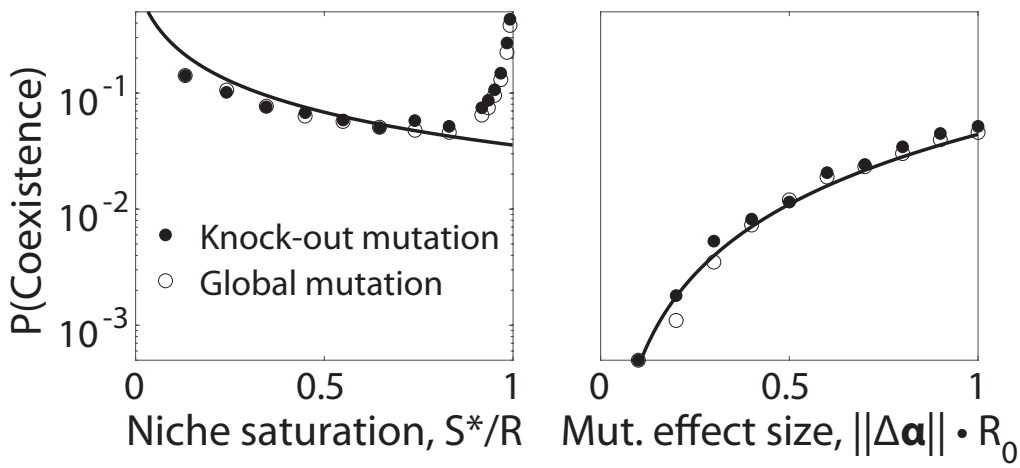

Figure S5: Coexistence probability for different strategy mutations. Left: Mutant-parent coexistence probability as a function of niche saturation $\mathcal{S}^*/\mathcal{R}$ for full knock-out mutations and "global" strategy mutations with $\|\Delta\vec{\alpha}\| = 1/\mathcal{R}_0$ (which have similarly-sized changes to all nonzero entries of $\vec{\alpha}$, as described in SI Section 4.4). Right: Mutant-parent coexistence probability as a function of mutation effect size $\|\Delta\vec{\alpha}\|$ for $\mathcal{S}^*/\mathcal{R} = 0.8$, for partial knock-out and global strategy mutations. Black curves show theoretical prediction. All simulations run with $\mathcal{R} = 100$, $\mathcal{R}_0 = 40$, and $\mathcal{S}^*/\mathcal{S} = 0.1$ using Dirichlet-distributed resource usage vectors.

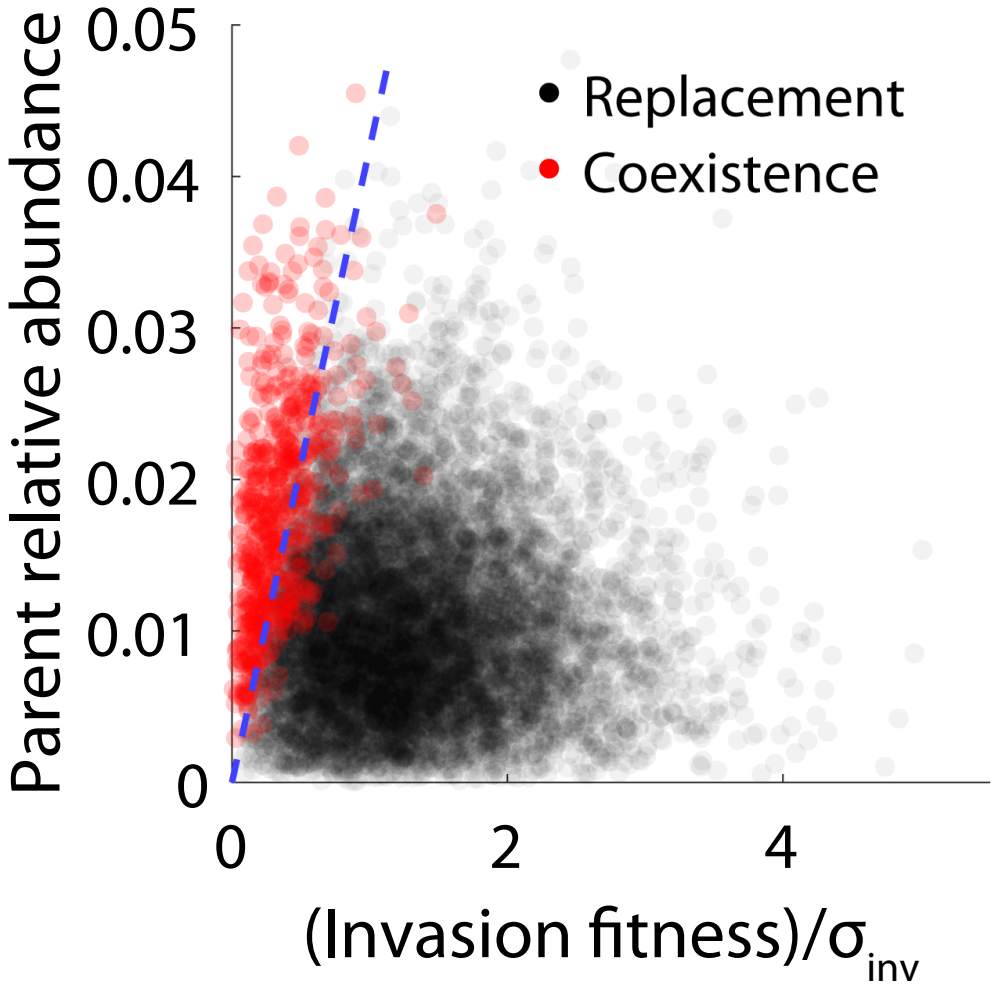

Figure S6: Scatterplot showing the invasion fitness of knock-out strategy mutations against the relative abundance of the parent before mutant invasion, with colors indicating whether the mutants coexist with or replace their parent at ecological equilibrium. The dashed blue line shows the theoretical prediction that should divide replacement events from coexistence events. Simulations run for $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and $\mathcal{S}^*/\mathcal{R} = 0.8$.

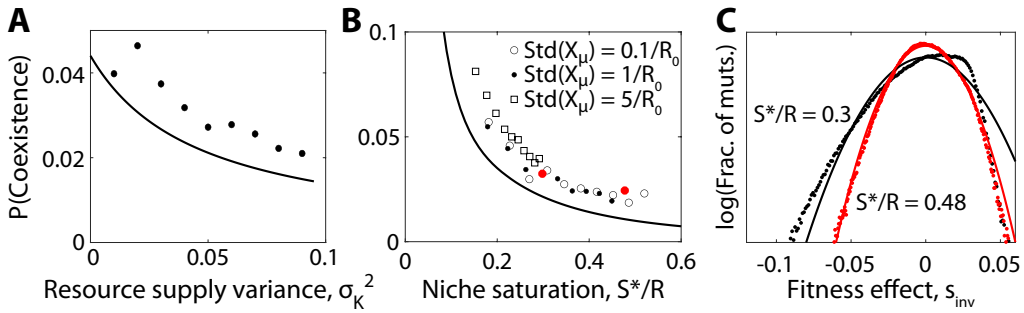

Figure S7: The distribution of fitness effects and mutant-parent coexistence probability in the presence of heterogeneous resource supply rates. **(A)** Mutant-parent coexistence probability for knockout mutants as a function of scaled variance in resource supply $\sigma_K^2 \equiv \mathrm{Var}(K)/\overline{K}^2$, with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and $\mathcal{S}^*/\mathcal{R} = 0.8$. **(B)** Coexistence probability of parents and knock-out mutants for communities assembled with exponentially distributed resource supply (corresponding to $\sigma_K^2 = 1$), and different levels of uptake budget variation $\mathrm{Std}(X_\mu)$ among sampled organisms. Simulations were run with $\mathcal{R} = 150$ and $\mathcal{R}_0 = 30$. **(C)** Distribution of fitness effects for knock-out mutants for communities corresponding to the red points in (B). In all plots, points show simulation results while curves show theory predictions.

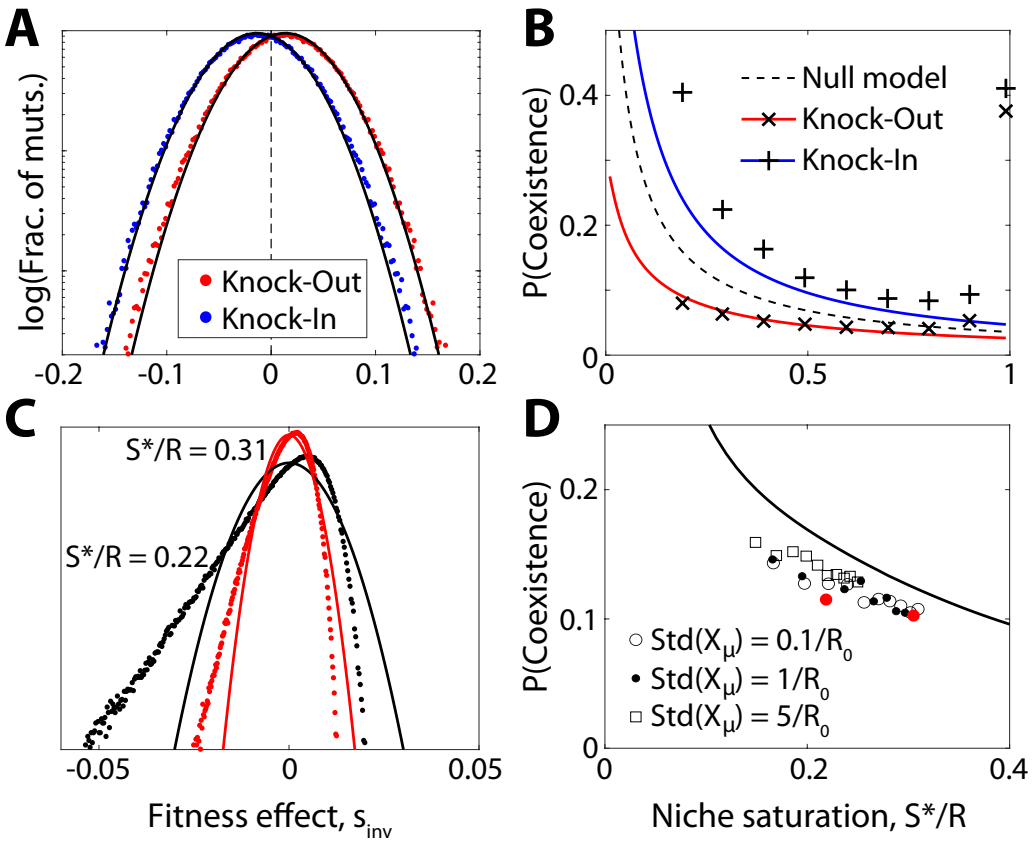

Figure S8: **(A)** Distribution of invasion fitnesses for the alternative model in Ref. (36), where the cost of using a resource contributes to the death rate of an organism, rather than occupying a portion of its overall resource uptake budget. Black curves show theoretical prediction for knock-in and knock-out mutations, while points show simulation results. Simulations run with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and $\mathcal{S}^*/\mathcal{R} = 0.8$. **(B)** Mutant-parent coexistence probability for the Ref. (36) CRM with the same parameters as (A), as a function of niche saturation, $\mathcal{S}^*/\mathcal{R}$. The average fitness benefit (cost) for knock-out (knock-in) mutations effectively gives strategy mutations a nonzero change in pure fitness, which lowers (raises) the theoretically predicted coexistence probability relative to the neutral case. Red and blue curves show these adjusted predictions, while crosses and pluses show simulated results. **(C)** Distribution of fitness effects for knock-out mutations in a community where sampled species lack metabolic trade-offs, as described in SI Section 4.3. Simulations run with $\mathcal{R} = 150$ and $\mathcal{R}_0 = 30$. **(D)** Mutant-parent coexistence probability for the same ecosystem as a function of niche saturation, for different levels of variation in uptake budget $\mathrm{Std}(X_\mu)$. Red points correspond to the parameter values in (C).

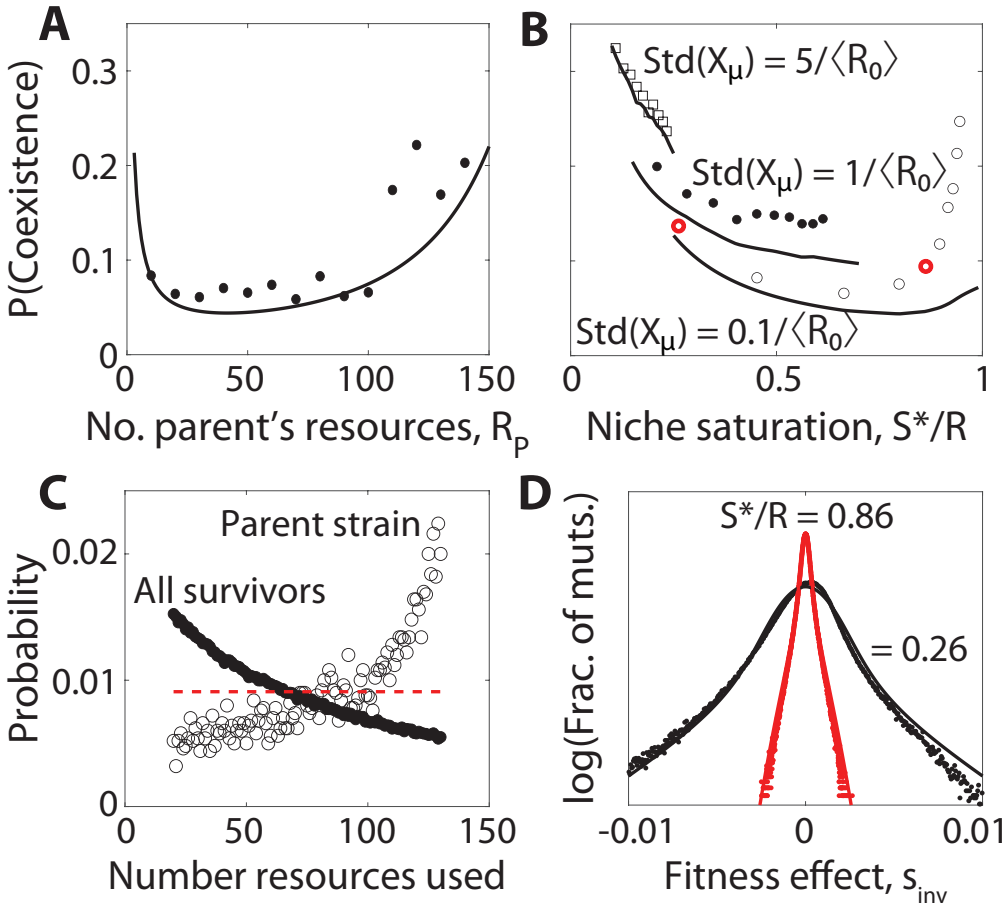

Figure S9: Extensions to communities with wider variation in the total number of resources consumed per species. **(A)** Probability that the mutant coexists with its parent strain as a function of the number of resources $\mathcal{R}_P$ used by the parent, in a background community with $\mathcal{R}_0 = 40$. Points denote simulation results for $\mathcal{R} = 200$, $\mathcal{S}^*/\mathcal{S} = 0.1$, and $\mathcal{S}^*/\mathcal{R} = 0.8$. **(B)** Mutant-parent coexistence probability when the number of resources used by each strain is uniformly distributed between 30 and 120 (SI Section 4.2). Points denote simulation results for knock-out mutations with $\mathcal{R} = 150$, and different levels of variation in uptake budget $\mathrm{Std}(X_\mu)$. **(C)** Distribution of the actual number of resources used by surviving species in the community corresponding to the rightmost red circle in (B), including all survivors (closed circles) and the subset of species which produced a successful knock-out mutation (open circles). The red dashed line shows the corresponding uniform distribution for sampled species. **(D)** Distribution of knock-out fitness effects for the community in (B), with parameter values corresponding to the red circles. All points show simulation results, while curves show theory predictions, using the distribution in (C) as an input.

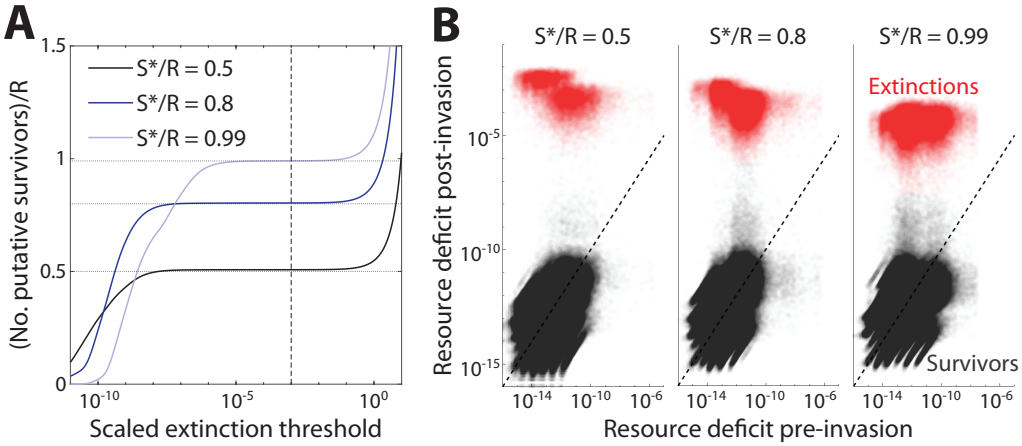

Figure S10: Checks to ensure that we have consistently identified extinctions numerically. **(A)** Plot of the number of putatively identified survivors for an assembled community, as a function of the numerical extinction threshold. If the resource deficit for an organism is above this threshold times $\mathrm{Std}(h_i)$, the organism is labeled extinct. The plateau over multiple orders of magnitude indicates a numerically stable definition of survival vs. extinction; the dashed line shows the threshold used throughout our analysis. **(B)** Scatterplot of numerical resource deficit of organisms before and after invasion of a knock-out strategy mutant, with color indicating whether the resource deficit after invasion passes the extinction threshold, (i.e., whether the species has gone extinct). The fact that most red and gray points are clearly separated shows that most species identified as extinct are insensitive to our precise choice of threshold. Simulations were run with $\mathcal{R} = 200$, $\mathcal{R}_0 = 40$, $\mathcal{S}^*/\mathcal{S} = 0.1$.

1 Model Outline

In this section, we outline the setup that allows us to derive our main theoretical results. First, we describe the consumer resource model we study, and how its dynamics can be coarse-grained to a limit where species consume most of the available resources. The task of finding an ecological equilibrium in this coarse-grained model can be described as an optimization problem. Then, we define a convenient fitness gauge which brings our model in line with previous work.

1.1 General Ecological Dynamics

As summarized in the Model section, we are interested in the following system: consider an ecosystem formed by combining $\mathcal{S}$ species, identically and independently drawn from some underlying distribution, competing for $\mathcal{R}$ resources. We allow this ecosystem to come to ecological equilibrium, at which point $\mathcal{S}^* \leq \mathcal{S}$ species will survive at nonzero abundance. Then, suppose that one randomly selected surviving species undergoes a mutation to produce a new strain, which can differ from its parent in its ability to consume resources. If the mutant has positive invasion fitness, it will rise to nonzero abundance in the population, and the abundances of other species may change as well. We are interested in analyzing the properties of this new ecosystem with a single mutant, particularly the probability that the mutant and its parent strain are able to coexist at ecological equilibrium.

First, we define our community assembly process and consumer resource model, which generally follows previous work (28–32, 36, 45). We consider a well-mixed community formed from the combination of $\mathcal{S}$ strains, indexed by $\mu \in \{1, ..., \mathcal{S}\}$, and $\mathcal{R}$ resources, indexed by $i \in \{1, ..., \mathcal{R}\}$. Each organism is associated with a resource uptake vector, $\vec{r}_\mu = (r_{\mu,1}, \ldots, r_{\mu,\mathcal{R}})$, which could represent the expression of enzymes used to metabolize each resource. The consumer-resource dynamics for the strain abundances (n_μ) and the local resource concentrations (c_i) are given by Eq. (1) in the Main Text:

$$\frac{\partial n_\mu}{\partial t} = \sum_{i=1}^{\mathcal{R}} Y_{\mu,i} u_i(\vec{c}) r_{\mu,i} n_\mu - \delta \cdot n_\mu \tag{S1a}$$

$$\frac{\partial c_i}{\partial t} = K_i - \sum_{\mu=1}^{\mathcal{S}} u_i(\vec{c}) r_{\mu,i} n_\mu - \delta \cdot c_i, \tag{S1b}$$

where δ is the dilution rate, K_i is the input flux of each resource, $u_i(\vec{c}) r_{\mu,i} n_\mu$ is the total uptake rate of resource i by strain μ, and $Y_{\mu i}$ is the yield of resource i when converted into biomass by strain μ. For simplicity, we will assume that this yield matrix can be factored into a rank-one form $Y_{\mu,i} = a_i / b_\mu$. While this assumption can approximately account for effects like different stoichiometric yields and cell sizes, it neglects idiosyncratic effects such as resource sequestration (47), where some organisms remove resources from the environment without corresponding growth in abundance. This simplification allows us to choose "biomass-equivalent" units such that $Y_{\mu,i} = 1$, yielding a conservation of mass equation

corresponding to a symmetry between resource consumption and biomass growth:

$$\frac{\partial}{\partial t}\left[\sum_i c_i + \sum_\mu n_\mu\right] = \sum_i K_i - \delta\left[\sum_i c_i + \sum_\mu n_\mu\right]. \tag{S2}$$

The total mass rapidly approaches the $\mathcal{S}+\mathcal{R}-1$ dimensional manifold,

$$\sum_i c_i(t) + \sum_\mu n_\mu(t) = \frac{\sum_i K_i}{\delta}, \tag{S3}$$

on a timescale proportional to the inverse dilution rate δ^{-1} (which is equivalent to the average generation time at steady state). Since we are interested in the dynamics on timescales that are much longer than a single generation, we assume that Eq. (S3) always applies.

1.2 Coarse-Grained Ecological Dynamics

The dynamics in Eq. (S1) can be further simplified if we assume that the resource uptake rates are sufficiently fast, and the total biomass of organisms sufficiently large, that most supplied resources are consumed by organisms rather than diluted out. When this is the case, we can coarse-grain over time such that the resource concentrations are instantaneously set by the abundance of each strain, rather than being independent dynamical variables ("integrating out the resources"). While this procedure follows previous work (30, 32), we will restate this approximation and its self-consistency conditions for completeness.

As shown by Equation (S1), it is reasonable to approximate that most resources are converted to biomass rather than diluted out when the local concentration of each resource is much smaller than $c_i^{\max} \equiv K_i/\delta$, which is the equilibrium value that would be obtained in the absence of any microbes. In other words, the community must be able to deplete each resource far below the level that would be expected from the environmental supply rates alone. When this condition is satisfied, the total biomass in Eq. (S3) will be dominated by the microbial contribution,

$$\sum_\mu n_\mu(t) \approx \frac{\sum_i K_i}{\delta}, \tag{S4}$$

and the uptake rates of each resource will approach a local steady state set by the current strain abundances:

$$u_i(\vec{c}) \approx \frac{K_i}{\sum_\mu n_\mu r_{\mu,i}} \approx \left(\frac{K_i}{\sum_j K_j}\right)\frac{\delta}{\sum_\mu f_\mu r_{\mu,i}}, \tag{S5}$$

where we have defined the relative abundance variables,

$$f_\mu \equiv \frac{n_\mu}{\sum_\nu n_\nu} \approx \frac{\delta \cdot n_\mu}{\sum_i K_i}. \tag{S6}$$

Substituting these expressions into Eq. (S1a) yields the coarse-grained model in Eq. (2) in the Main Text:

$$\frac{\partial f_\mu}{\partial \tau} = f_\mu \left[\sum_{i=1}^{\mathcal{R}} r_{\mu,i} h_i(\vec{f}) - 1 \right] , \tag{S7a}$$

where time is measured in generations ($\tau \equiv \delta \cdot t$) and

$$h_i(\vec{f}) \equiv \frac{u_i(\vec{c})}{\delta} \approx \left(\frac{K_i}{\sum_j K_j} \right) \left[\sum_\mu r_{\mu,i} f_\mu \right]^{-1} \tag{S7b}$$

can be interpreted as the local availability of resource i.

Self-consistency of coarse-grained dynamics. These coarse-grained dynamics will be self-consistent when the relative abundances that are attained in Eq. (S7) satisfy

$$u_i^{-1} \left(\frac{K_i}{\sum_j K_j} \frac{\delta}{\sum_\mu r_{\mu,i} f_\mu} \right) \ll \frac{K_i}{\delta} , \tag{S8}$$

where $u_i^{-1}(\cdot)$ is the inverse of the uptake function $u_i(\vec{c})$. Since the relative abundances only depend on the relative supply rates, $K_i / \sum_j K_j$, this condition can always be satisfied if δ is sufficiently small, or if the total biomass is sufficiently large. For example, if the uptake rates can be described by a Monod function,

$$u_i(\vec{c}) = \frac{u_{\max,i} \cdot c_i}{c_i^* + c_i} , \tag{S9}$$

then the condition in Eq. (S8) reduces to

$$\frac{K_i}{c_i^* \delta} \left[\frac{u_{\max,i}}{\delta} \frac{\sum_j K_j}{K_i} \sum_\mu r_{\mu,i} f_\mu - 1 \right] \gg 1 . \tag{S10}$$

This will be valid for small enough values of δ, or large enough values of $\sum_j K_j / K_i$.

Rescaled variables. In our analysis below, it will be convenient to reparameterize the model in Eq. (S7) by decomposing the resource uptake vector $\vec{r}_\mu$ into an overall "uptake budget" $X_\mu \equiv \log \sum_i r_{\mu,i}$, and a normalized "strategy vector" $\vec{\alpha}_\mu \equiv \vec{r}_\mu / \sum_i r_{\mu,i}$, which contains $\mathcal{R}$ non-negative entries that sum to one. The original uptake rates can then be expressed as

$$r_{\mu,i} = \alpha_{\mu,i} e^{X_\mu} . \tag{S11}$$

The uptake budget dictates the overall capacity of the organism to consume resources, setting its maximum growth rate under ideal growth conditions; the strategy vector describes how the organism allocates this uptake budget between the different substitutable resources. This alternative parametrization is equivalent to the original basis $\vec{r}_\mu$, but will lead to simpler expressions in many of our analytical calculations below.

To streamline notation, it will also be useful to define the relative resource supply rates,

$$\kappa_i \equiv \frac{K_i}{\frac{1}{\mathcal{R}}\sum_j K_j}\,, \tag{S12}$$

which describe how the supply of a given resource deviates from the average value, $\overline{K} = \frac{1}{\mathcal{R}}\sum_i K_i$. The relative supply rates are normalized such that $\frac{1}{\mathcal{R}}\sum_i \kappa_i = 1$. We will also refer to the variance of κ_i, $\sigma_K^2 \equiv \mathrm{Var}(\kappa) = \mathrm{Var}(K)/\overline{K}^2$.

1.3 Incorporating Simple Forms of Cross-feeding

While the derivation of Eq. (S7) assumed that the resources were externally supplied, similar dynamics can continue to hold for simplified forms of cross-feeding, provided that κ_i is replaced by an effective supply rate that accounts for internal metabolic conversion. To see this, we consider the "universal" cross-feeding model from Ref. (49), where the metabolism of each resource is associated with a corresponding "leakage rate" ℓ_i, which describes the proportion of that resource that is excreted back into the environment in the form of other resources, rather than being converted into biomass. This conversion process is described by a stoichiometric matrix $\mathbf{D}$, whose entries $D_{i,j}$ represent the fraction of resource j that is excreted in the form of resource i. This model assumes that every organism capable of metabolizing a resource produces the same byproducts in the same proportions; species can still differ in their overall excretion profiles, but these are completely determined by the corresponding uptake rates $\vec{r}_\mu$. These cross-feeding dynamics can be described by a generalization of Eq. (S1),

$$\frac{\partial n_\mu}{\partial t} = n_\mu\left[\sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} e^{X_\mu}(1-\ell_i)h_i(\vec{c}) - 1\right], \tag{S13a}$$

$$\frac{\partial c_i}{\partial t} = K_i - \delta_i c_i + \sum_{\mu=1}^{\mathcal{S}} n_\mu e^{X_\mu}\left[-h_i(\vec{c})\alpha_{\mu,i} + \sum_{j=1}^{R}\ell_j D_{i,j} h_j(\vec{c})\alpha_{\mu,j}\right], \tag{S13b}$$

where the first term in the brackets of Eq. (S13b) represents the consumption of that resource, while the second term represents the production via metabolism of other resources. If we can continue to make the approximation that the dilution of resources is negligible compared to their consumption by other organisms, then the steady-state resource availabilities $h_i(\vec{c})$ will obey a generalization of Eq. (S5),

$$h_i \sum_\mu n_\mu e^{X_\mu}\alpha_{\mu,i} = \sum_j (\mathbf{A}^{-1})_{i,j} K_j, \tag{S14}$$

where $\mathbf{A}$ is an $\mathcal{R}$-by-$\mathcal{R}$ matrix whose entries are given by $A_{i,j} = \delta_{i,j} - \ell_j D_{i,j}$ (note that $\delta_{i,j}$ is the Kronecker delta). This shows that we can recover the same dynamics in Eq. (S7) by defining the effective supply rates,

$$\kappa_i^{\mathrm{eff}} = (1-\ell_i)\sum_j (\mathbf{A}^{-1})_{i,j}\kappa_j = \frac{\sum_j (1-\ell_i)(\mathbf{A}^{-1})_{i,j}K_j}{\sum_{k,\ell}(1-\ell_k)(\mathbf{A}^{-1})_{k,\ell}K_\ell/\mathcal{R}}. \tag{S15}$$

which depends on the external supply rates and the stoichiometry of metabolic conversions, but is otherwise independent of the species that are present in the community.

1.4 Defining the Local Species Pool

All that remains is to choose the resource consumption parameters of the $\mathcal{S}$ initial colonizers. As described in the text, we assume that the phenotypes of each species are drawn from a common statistical distribution. For our theoretical analysis, we follow the procedure of Ref. (36), where the uptake budgets X_μ are drawn from a normal distribution with standard deviation $\epsilon/\mathcal{R}_0$, and the resource consumption strategies $\vec{\alpha}_\mu$ are normalized binary vectors where each entry has an independent probability $\mathcal{R}_0/\mathcal{R}$ of being nonzero. The independence of X_μ and $\vec{\alpha}_\mu$ results in a soft metabolic tradeoff, where strains that consume more resources are usually less effective at using any individual one. To test the importance of this and other assumptions, we also perform simulations for several alternative distributions (such as Dirichlet-distributed resource consumption strategies) to verify that our qualitative results are robust (SI Section 4; Figures S3, S8, and S9).

1.5 Ecological Equilibrium as Optimization

For any fixed set of species, the ecological dynamics in Eq. (S7) can be shown to possess a convex Lyapunov function,

$$\Lambda(\vec{f}) \equiv \frac{1}{\mathcal{R}} \sum_i \kappa_i \log h_i(\vec{f}) - \sum_\mu f_\mu \,, \tag{S16}$$

and a single stable equilibrium (see, e.g. Refs. 30, 32, 45). At this equilibrium, none of the surviving strains can increase in abundance, while none of the extinct strains could invade the ecosystem if they were re-introduced at low abundance. Thus, at ecological equilibrium, the bracketed term in Eq. (S7) (the "resource surplus") must be non-positive for all μ, and is equal to zero if and only if species μ survives at nonzero abundance.

Conveniently, the task of finding the unique stable equilibrium in this model can be recast as a constrained optimization problem (30, 36, 37). Specifically, we treat the resource availabilities h_i (rather than the relative abundances f_μ) as free variables to optimize over, and consider the task

$$\begin{aligned} &\text{maximize} \sum_i \kappa_i \log h_i \\ &\text{subject to} \sum_i \alpha_{\mu,i} h_i - e^{-X_\mu} \leq 0 \; \forall \mu. \end{aligned} \tag{S17}$$

The constraint is the same as that discussed above, requiring that the resource surplus be non-positive for each organism. This model possesses an important symmetry which will be useful when making approximations later: it is invariant under a rescaling of the total uptake budget, which we can view as a change in "pure fitness" decoupled from the environment. In particular, if we change variables into translated fitness $\tilde{X}_\mu = X_\mu + \chi$, and

rescaled $\tilde{h}_i = e^{-\chi} h_i$, then the optimization problem becomes

$$\begin{aligned} &\text{maximize} \sum_i \kappa_i \left(\log \tilde{h}_i + \chi \right) \\ &\text{subject to } e^{\chi} \left(\sum_i \tilde{h}_i \alpha_{\mu i} - e^{-\tilde{X}_\mu} \right) \leq 0 \; \forall \mu. \end{aligned} \tag{S18}$$

Since χ is just an arbitrary constant, it is clear that these optimization problems are the same, up to the rescaling of the h_i. Suppose that h_i^* solves the first optimization problem. Then, $\tilde{h}_i^* = e^{-\chi} h_i^*$ will solve the gauge-shifted optimization problem

$$\begin{aligned} &\text{maximize} \sum_i \kappa_i \log \tilde{h}_i \\ &\text{subject to} \sum_i \tilde{h}_i \alpha_{\mu i} - e^{-\tilde{X}_\mu} \leq 0 \; \forall \mu. \end{aligned} \tag{S19}$$

To convince ourselves that these two problems really have the same solution, we return to the equation defining h_i in terms of the relative equilibrium abundances f_μ:

$$h_i = \frac{\kappa_i / \mathcal{R}}{\sum_i \alpha_{\mu i} e^{X_\mu} f_\mu}. \tag{S20}$$

Translating the uptake budgets by χ rescales h_i by $e^{-\chi}$, while leaving f_μ unchanged. Therefore, both optimization problems correspond to the same f_μ, which is the actual observable we are interested in. In our numerical simulations, we solve this optimization problem directly, with an arbitrary fitness gauge chosen for numerical convenience.

Once the h_i have been determined, the strain relative abundances f_μ can be calculated numerically. First, the living strains with $f_\mu > 0$ are identified as the strains for which the inequality constraint in Eq. (S17) achieves equality. Given these strains, the definition of h_i in Eq. (S7b) can be inverted to obtain a system of linear equations for f_μ:

$$\sum_\mu r_{\mu,i} f_\mu = \frac{\kappa_i}{\mathcal{R} h_i}. \tag{S21}$$

Because f_μ for extinct strains must be zero, there are only $\mathcal{S}^* \leq \mathcal{R}$ remaining abundances to be solved for. A non-negative linear least squares solver can be used to find the unique set of f_μ that satisfy all $\mathcal{R}$ equations.

1.6 Analytic Approximation of Ecological Equilibrium

By using this gauge symmetry and a small number of additional approximations, we can rewrite this optimization problem in a more analytically tractable form. Indeed, we find that we can put our model in a similar form to that analyzed in Refs. (36) and (37), by choosing a gauge such that $\langle h_i \rangle = 1$. Let's define a constant h_0 based on the equilibrium solution in our (arbitrary) original gauge,

$$h_0 = \frac{1}{\mathcal{R}} \sum_i h_i^*, \tag{S22}$$

and consider what happens when we set $\chi = \log h_0$. The rescaled solution is

$$\tilde{h}_i^* = \frac{h_i^*}{h_0}, \tag{S23}$$

which by definition has mean 1. It is now natural to define a new variable, $g_i = 1 - \tilde{h}_i$, which has mean zero. The optimization problem is now

$$\begin{aligned} &\text{optimize} \sum_i \kappa_i \log(1 - g_i) \\ &\text{subject to } 1 - e^{-\tilde{X}_\mu} - \sum_i g_i \alpha_{\mu i} \leq 0 \; \forall \mu. \end{aligned} \tag{S24}$$

In large ecosystems, the g_i's will often be small compared to one, allowing us to Taylor expand the optimization function to second order. (We will verify the range of parameters where this holds in SI Section 3.5.) When $g_i \lesssim 1$, the third term on the LHS of the constraint equation must also be less than one. This implies that for any surviving species (where the constraint equation is exactly satisfied) the fitness term $1 - e^{-\tilde{X}_\mu}$ must also be small compared to one. In other words, choosing a gauge where $\langle h_i \rangle = 1$ is the same as choosing a gauge where the pure fitnesses of surviving organisms are centered at zero, allowing us to Taylor expand the $\tilde{X}_\mu$ as well. This relationship arises from the fact that strains with very large differences in total uptake budget are unlikely to coexist. Altogether, the approximate optimization problem is

$$\begin{aligned} &\text{optimize} \sum_i \kappa_i \left(g_i + \frac{g_i^2}{2} \right) \\ &\text{subject to } \tilde{X}_\mu - \sum_i g_i \alpha_{\mu i} \leq 0 \; \forall \mu. \end{aligned} \tag{S25}$$

For organisms that go extinct, the Taylor expansion of the $e^{-\tilde{X}_\mu}$ terms may not be very accurate, since $\tilde{X}_\mu$ can be arbitrarily negative. However, since $1 - \tilde{X}_\mu \leq e^{-\tilde{X}_\mu}$, this approximation will not cause us to erroneously label any extinct species as alive – and for the species which are alive, we will check that the approximation is sufficiently accurate to avoid worry. While it is not yet clear why this approximation is useful, we will see below that remapping the consumer resource problem in this way makes it amenable to replica-theoretic analysis. Note that for numerical simulations, we do not make these approximations, instead solving the original optimization problem in Eq. (S17) – as such, our simulation results implicitly test these assumptions across a range of parameters. We discuss the regime of validity of these assumptions and their interpretations in SI Section 3.5.

2 Replica-Theoretic Analysis of First-Step Invasion

2.1 Single-Ecosystem Approximation of First-Step Invasion

Having defined our species pool and consumer resource model, we now consider adding first-step mutations into this community. As discussed in the Model section, mutant-

parent coexistence can be described as two correlated community assembly processes. Here, we argue that we can reasonably approximate the conditions for mutant-parent coexistence while considering only one community assembly process, which simplifies our theoretical work. First, we describe the two-ecosystem framework. An ecosystem E_1 is formed through standard community assembly: $\mathcal{S}$ species are placed into competition, and $\mathcal{S}^*$ remain alive at equilibrium. Then, one of the surviving species produces a beneficial mutant offspring. This effectively creates a new community assembly problem for ecosystem E_2, where the initial $\mathcal{S}^*$ species plus one mutant are allowed to reach ecological equilibrium. The mutant-parent coexistence probability can therefore be written as

$$\mathbb{P}_{\text{coex}} = \mathbb{P}\left[\text{parent alive in } E_2 | \text{mutant alive in } E_2, \text{parent alive in } E_1\right], \tag{S26}$$

where the conditional probability requires that (1) the mutant be beneficial and (2) the parent be alive when it produces a mutant. Standard manipulation of probabilities allows us to rewrite this as

$$\mathbb{P}_{\text{coex}} = \frac{\mathbb{P}\left[\text{P and M in } E_2\right] - \mathbb{P}\left[\text{P and M in } E_2,\ \text{P not in } E_1\right]}{\mathbb{P}\left[\text{M in } E_2\right] - \mathbb{P}\left[\text{M in } E_2,\ \text{P not in } E_1\right]}. \tag{S27}$$

First, we argue that $\mathbb{P}\left[\text{P and M in } E_2,\ \text{P not in } E_1\right] \approx 0$. This probability describes a situation where the parent is unable to survive in the initial ecosystem, but its hypothetical mutant offspring is. Moreover, the invasion of the mutant offspring adjusts the ecosystem in a way where the original parent strain is also able to simultaneously invade. Intuitively, both of these requirements are unlikely: it will be difficult for a small mutation to allow an otherwise unfavored species to invade, and we would naively expect that the presence of the mutant would make it more difficult for the parent species to survive – not easier – due to their shared competition for resources. Making this approximation, we find

$$\mathbb{P}_{\text{coex}} \approx \frac{\mathbb{P}\left[\text{P and M in } E_2\right]}{\mathbb{P}\left[\text{M in } E_2\right] - \mathbb{P}\left[\text{M in } E_2,\ \text{P not in } E_1 \text{ or } \text{E}_2\right]}. \tag{S28}$$

The next approximation we make is similar: we assume that if the mutant is able to survive, the parent would likely be able to survive in its absence. Mathematically, $\mathbb{P}\left[\text{M in } E_2\right] \gg \mathbb{P}\left[\text{M in } E_2,\ \text{P not in } E_1 \text{ or } \text{E}_2\right]$. While it is possible that a small-effect beneficial mutation could push an organism over the survival threshold, we expect that *most* randomly chosen surviving organisms will not be able to be rescued from extinction by a single mutation. Thus, we write

$$\mathbb{P}_{\text{coex}} \approx \frac{\mathbb{P}\left[\text{P and M in } E_2\right]}{\mathbb{P}\left[\text{M in } E_2\right]}. \tag{S29}$$

Now, our coexistence probability depends only on one ecosystem. But there is still an issue: E_2 is not a traditionally assembled community, since it is formed only from the survivors of an earlier community assembly process plus a mutant, rather than independently sampled organisms. Let's instead consider E_s, an ecosystem obtained from simultaneous community assembly of $\mathcal{S}$ independently sampled species, plus one "mutant" species which is closely related to one of the $\mathcal{S}$ other species. The key difference between E_2 and E_s is that E_s can potentially contain species which were extinct in E_1, but are brought "back to life" by the invasion of the mutant. While such returns are possible (37), we conjecture that

the invasion of a closely-related mutant strain should be a relatively small perturbation to the ecosystem, and as such these re-introductions should be rare – particularly when the number of resources and species are large. So, as a first approximation, we do not expect them to meaningfully impact $\mathbb{P}_{\text{coex}}$. After these approximations, we now have

$$\mathbb{P}_{\text{coex}} \approx \frac{\mathbb{P}\left[\text{P and M in } E_s\right]}{\mathbb{P}\left[\text{M in } E_s\right]}. \tag{S30}$$

This approximation allows us to estimate the mutant-parent coexistence probability in an assembled community by only considering one community assembly process, which includes two closely correlated organisms. (Which one we have labeled the "parent" and "mutant" is essentially arbitrary after making these approximations.) As shown in Fig. S4, simulations confirm that these approximations are valid across most of the parameter regime we study, so our one-ecosystem approximation produces results which match the two-ecosystem one. An exception is at very high expected niche saturation, where species "coming back to life" after mutant invasion becomes more relevant. However, there is a narrow range of parameters where this effect is significant (particularly for ecosystems with many resources), and our approximation where these species are allowed to return nonetheless sets a lower bound on the mutant-parent coexistence probability.

2.2 Partition Function and the Replica Trick

We would now like to argue that our model can be treated similarly to Refs. (36) and (37), meaning that their results describe a baseline assembled community before we add mutations. While we will not reproduce every element of their calculation, we restate key steps and approximations. Since ecological equilibrium can be described as the solution to an optimization problem, we can write a partition function for a system with energy $F(\vec{g}) = \sum_i F_i(g_i) = \sum_i \kappa_i \left(g_i + \frac{g_i^2}{2}\right)$ at temperature β^{-1}. In the zero-temperature limit $\beta \to \infty$, the system will be exactly at ecological equilibrium; at finite values of β, fluctuations from equilibrium correspond to shot noise in the population. The partition function is

$$Z = \int_\Omega \prod_i dg_i e^{-\beta F(\vec{g})}, \tag{S31}$$

where Ω corresponds to the constrained region in our optimization problem, where resource surplus is non-positive. We can write out this region explicitly using the Heaviside function $\theta(x) = \max(0, x)$ as follows:

$$Z = \int_{-\infty}^{1} \prod_i dg_i e^{-\beta F_i(g_i)} \prod_{\mu=1}^{S} \theta\left(\sum_i g_i \alpha_{\mu i} - \tilde{X}_\mu\right) \tag{S32}$$

$$= \int_{-\infty}^{1} \prod_i dg_i e^{-\beta F_i(g_i)} \prod_{\mu=1}^{S} \int d\Delta_\mu \theta(-\Delta_\mu) \delta\left(\Delta_\mu + \sum_i g_i \alpha_{\mu i} - \tilde{X}_\mu\right) \tag{S33}$$

$$= \int_{-\infty}^{1} \prod_i dg_i e^{-\beta F_i(g_i)} \prod_\mu \int \frac{d\Delta_\mu d\hat{\Delta}_\mu}{2\pi} \theta(-\Delta_\mu) e^{i \sum_\mu \hat{\Delta}_\mu \left(\Delta_\mu + \sum_i g_i \alpha_{\mu i} - \tilde{X}_\mu\right)}. \tag{S34}$$

where $\tilde{X}_\mu$ is the pure fitness of species μ, in the gauge described in 6.1.2. We have traded our complicated integration region for two additional integration variables for each species: the resource surplus Δ_μ and its auxiliary Fourier variable $\hat{\Delta}_\mu$.

This partition function is written for a particular choice of S sampled species. However, we are interested not in any particular choices of species, but the behavior of the ecosystem when *typical* species are drawn from some random distribution. These typical ecosystems can be analyzed by calculating $\langle \log Z \rangle$, where the average is over our random draws of species. This average can be calculated through use of the replica trick,

$$\langle \log Z \rangle = \lim_{n \to 0} \frac{\langle Z^n \rangle - 1}{n}. \tag{S35}$$

In accordance with the replica trick, we consider n copies of our system, which share identical resident species but may differ in shot noise (encoded by the temperature β). Then, we treat n as a real number and consider the limit $n \to 0$, comparing our results to simulation to ensure this approximation holds. With n copies of our system, the average partition function is

$$\begin{aligned}
\langle Z^n \rangle = &\int \prod_{i,a} \left[dg_i^a e^{-\beta \sum_{i,a} F_i(g_i^a)} \right] \times \\
&\prod_{\mu=1}^{S} \left\{ \prod_a \left[\int \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a) \right] e^{i \sum_a \hat{\Delta}_\mu^a \Delta_\mu^a} \left\langle e^{-i\tilde{X}_\mu \sum_a \hat{\Delta}_\mu^a} \right\rangle_{\tilde{X}_\mu} \times \right. \\
&\left. \left\langle e^{i \sum_{i,a} \hat{\Delta}_\mu^a g_i^a \alpha_{\mu,i}} \right\rangle_{\alpha_{\mu,i}} \right\},
\end{aligned} \tag{S36}$$

where a runs from 1 to n and indexes the replicas of our system. Next, let's average over our random draws of species. Suppose that the pure fitnesses $\tilde{X}_\mu$ are drawn from a normal distribution with mean $\chi/\mathcal{R}_0$ and variance $(\epsilon/\mathcal{R}_0)^2$. Recall that χ represents our fitness gauge, and will later be set to ensure that $\langle g_i \rangle = 0$. The factors of $1/\mathcal{R}_0$ account for the fact that the $\alpha_{\mu,i}$ in Eq. (S25) are $\sim 1/\mathcal{R}_0$ smaller than the corresponding terms in Refs. (36) and (37) due to the overall normalization of $\vec{\alpha}$. The $\tilde{X}_\mu$-dependent term then becomes

$$\left\langle e^{-i\tilde{X}_\mu \sum_a \hat{\Delta}_\mu^a} \right\rangle_{\tilde{X}_\mu} = \exp\left[-i\frac{\chi}{\mathcal{R}_0} \sum_a \hat{\Delta}_\mu^a - \frac{\epsilon^2}{2\mathcal{R}_0^2} \left(\sum_a \hat{\Delta}_\mu^a \right)^2 \right]. \tag{S37}$$

Now, let's consider the average over the strategy vectors. For binary resource strategies, the $\alpha_{\mu i}$ are defined as

$$\alpha_{\mu i} = \frac{\sigma_{\mu i}}{\sum_j \sigma_{\mu j}}, \tag{S38}$$

where $\sigma_{\mu i}$ are i.i.d. random variables equal to 1 with probability $\mathcal{R}_0/\mathcal{R}$ and 0 otherwise. Unlike the $\sigma_{\mu i}$, the $\alpha_{\mu i}$ are not independent for the same μ with different i's. However, we can approximate them as independent by separating out the dependent portion, provided

that $\mathcal{R}_0 \gg 1$:

$$\alpha_{\mu i} = \frac{\sigma_{\mu i}}{\mathcal{R}_0 + b_\mu \sqrt{\mathcal{R}_0}} \tag{S39}$$

$$\approx \frac{\sigma_{\mu i}}{\mathcal{R}_0}\left(1 - \frac{b_\mu}{\sqrt{\mathcal{R}_0}}\right). \tag{S40}$$

The b_μ are mean-zero $\mathcal{O}(1)$ random variables representing the fluctuations in the total number of resources used by each species. While not technically independent of the $\sigma_{\mu i}$, this dependence would contribute to terms with smaller powers of $\mathcal{R}_0$ and can thus be ignored. So, we consider their averages separately:

$$\left\langle e^{i\sum_{i,a}\hat{\Delta}^a_\mu g^a_i \alpha_{\mu,i}}\right\rangle_{\alpha_{\mu,i}} \approx \left\langle \prod_i \left[1 - \frac{\mathcal{R}_0}{\mathcal{R}} + \frac{\mathcal{R}_0}{\mathcal{R}} e^{i\sum_a \hat{\Delta}^a_\mu g^a_i/\mathcal{R}_0} e^{-i\sum_a \hat{\Delta}^a_\mu g^a_i b_\mu/\mathcal{R}_0^{3/2}}\right]\right\rangle_{b_\mu}$$

$$\approx \left\langle \prod_i \exp\left[i\frac{1}{\mathcal{R}}\sum_a \hat{\Delta}^a_\mu g^a_i\left(1 - \frac{b_\mu}{\mathcal{R}_0}\right) - \frac{(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}_0\mathcal{R}}\left(\sum_a \hat{\Delta}^a_\mu g^a_i\right)^2\right]\right\rangle_{b_\mu} \tag{S41}$$

$$= \left\langle e^{-i\sum_{i,a}\hat{\Delta}^a_\mu g^a_i b_\mu/\mathcal{R}\sqrt{\mathcal{R}_0}}\right\rangle_{b_\mu} \prod_i \exp\left[i\frac{1}{\mathcal{R}}\sum_a \hat{\Delta}^a_\mu g^a_i - \frac{(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}\mathcal{R}_0}\left(\sum_a \hat{\Delta}^a_\mu g^a_i\right)^2\right] \tag{S42}$$

$$\approx \exp\left[-\frac{(\sum_{i,a}\hat{\Delta}^a_\mu g^a_i)^2}{2R^2\mathcal{R}_0}\right] \prod_i \exp\left[i\frac{1}{\mathcal{R}}\sum_a \hat{\Delta}^a_\mu g^a_i - \frac{(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}\mathcal{R}_0}\left(\sum_a \hat{\Delta}^a_\mu g^a_i\right)^2\right] \tag{S43}$$

$$\approx \prod_i \exp\left[i\frac{1}{\mathcal{R}}\sum_a \hat{\Delta}^a_\mu g^a_i - \frac{(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}\mathcal{R}_0}\left(\sum_a \hat{\Delta}^a_\mu g^a_i\right)^2\right]. \tag{S44}$$

Effectively, this approximation treats $\alpha_{\mu i} \approx \sigma_{\mu_i}/\mathcal{R}_0$. Finally, to bring our answer in line with Refs. (36) and (37), we rescale variables $\Delta^a_\mu \to \Delta^a_\mu/\mathcal{R}_0$ and $\hat{\Delta}^a_\mu \to \mathcal{R}_0\hat{\Delta}^a_\mu$. The final partition function is

$$\langle Z^n\rangle = \int \prod_{i,a}\left[dg^a_i e^{-\beta\sum_{i,a}F_i(g^a_i)}\right] \prod_{\mu=1}^{S}\left\{\prod_a\left[\int \frac{d\Delta^a_\mu d\hat{\Delta}^a_\mu}{2\pi}\theta(-\Delta^a_\mu)\right]\times\right.$$
$$\exp\left[i\sum_a \hat{\Delta}^a_\mu\left(\Delta^a_\mu + \frac{\mathcal{R}_0}{\mathcal{R}}\sum_i g^a_i - \chi\right) - \frac{\epsilon^2}{2}\left(\sum_a \hat{\Delta}^a_\mu\right)^2 - \right.$$
$$\left.\left.\frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}}\sum_i\left(\sum_a \hat{\Delta}^a_\mu g^a_i\right)^2\right]\right\}. \tag{S45}$$

This is the same as the partition function calculated in Refs. (36) and (37), except for the presence of χ. We can ignore χ until later in the calculation (for example, by absorbing it into the definition of Δ^a_μ, which causes it to appear only within the Heaviside function), and re-introduce it only when we require that $\langle g_i\rangle = 0$.

2.3 Introducing First-Step Mutations

The above partition function is valid for an ecosystem where all species are sampled randomly. However, in our "simultaneous assembly" approximation of first-step mutation, we would like to include two sampled species which are very related to each other, differing only by a single knock-out mutation. We will index the parent and mutant species P and M; they have the same resource consumption strategy, except that the parent uses resource 1, while the mutant's capacity to use resource 1 is lowered by a factor of γ. For a full knock-out mutation which preserves our assumption about binary resource use, $\gamma = 1$, but this more general type of mutation will aid us when interpreting our conclusions later. In particular, we will find that the magnitude of the phenotypic change, $\|\Delta\vec{\alpha}\| \approx \gamma/\mathcal{R}_0$, is most important when predicting its impact on the community. We assume the mutant's pure fitness is increased by a total amount ΔX relative to its parent (which can be negative to indicate an overall cost to the mutation). With the parent-mutant pair accounted for, the partition function becomes

$$
\begin{aligned}
\langle Z^n \rangle = &\int \prod_{i,a} \left[dg_i^a e^{-\beta \sum_{i,a} F_i(g_i^a)} \right] \prod_{\mu=1}^{S-1} \left\{ \prod_a \left[\int \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a) \right] \times \right. \\
&\exp\left[i \sum_a \hat{\Delta}_\mu^a \left(\Delta_\mu^a + \frac{\mathcal{R}_0}{\mathcal{R}} \sum_i g_i^a - \chi \right) - \frac{\epsilon^2}{2} \left(\sum_a \hat{\Delta}_\mu^a \right)^2 - \right. \\
&\left. \left. \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}} \sum_i \left(\sum_a \hat{\Delta}_\mu^a g_i^a \right)^2 \right] \right\} \times \\
&\prod_a \left[\int \frac{d\Delta_P^a d\hat{\Delta}_P^a}{2\pi} \frac{d\Delta_P^a d\hat{\Delta}_P^a}{2\pi} \theta(-\Delta_P^a)\theta(-\Delta_M^a) \right] \\
&\exp\left\{ i \sum_a \hat{\Delta}_P^a \left(\Delta_P^a + \frac{\mathcal{R}_0}{\mathcal{R}} \sum_i g_i^a + \left(1 - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g_1^a - \chi \right) + \right. \\
&i \sum_a \hat{\Delta}_M^a \left(\Delta_M^a + \frac{\mathcal{R}_0}{\mathcal{R}} \sum_i g_i^a + \left(1 - \gamma - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g_1^a - \chi - \mathcal{R}_0 \Delta X \right) - \\
&\frac{\epsilon^2}{2} \left[\sum_a \left(\hat{\Delta}_P^a + \hat{\Delta}_M^a \right) \right]^2 - \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}} \sum_i \left[\sum_a \left(\hat{\Delta}_P^a + \hat{\Delta}_M^a \right) g_i^a \right]^2 + \\
&\left. \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\mathcal{R}} \left[\sum_a \left(\hat{\Delta}_P^a + \hat{\Delta}_M^a \right) g_1^a \right]^2 \right\}. \qquad \text{(S46)}
\end{aligned}
$$

(Note that the ΔX term has a factor of $\mathcal{R}_0$ to account for the rescaling we did earlier.) The next few steps of the calculation are a series of approximations which allow us to evaluate this complicated integral, paralleling those in Refs. (36) and (37).

2.4 Saddle Point Approximations

For $\mathcal{R} \gg 1$, the central limit theorem dictates that the mean and variance of the equilibrium g_i (which encode the availabilities of each resource) should approach a deterministic value. Introducing the mutant-parent pair is a relatively small change in the ecosystem, so we expect that these deterministic values should be the same as those calculated in previous work, if the mutant-parent pair were not present. This motivates us to define replica-specific "order parameters"

$$m^a = \sum_i g_i^a, \tag{S47}$$

$$q^{ab} = \sum_i g_i^a g_i^b. \tag{S48}$$

As we did with the resource surplus Δ_μ^a, we can introduce these order parameters as new integration variables, defined via delta functions in their Fourier representation. This results in an integral

$$\begin{aligned}
\langle Z^n \rangle = &\int \prod_{a\leq b} \left(\frac{dq^{ab}d\hat{q}^{ab}}{2\pi}\right) \prod_a \left(\frac{dm^a d\hat{m}^a}{2\pi}\right) \exp\left[i\left(\sum_{a\leq b} q^{ab}\hat{q}^{ab} + \sum_a m^a \hat{m}^a\right)\right] \times \\
&\prod_{i=2}^{\mathcal{R}} \left\{ \int_{-\infty}^{1} \prod_a dg_i^a \exp\left[-\sum_a \beta F_i(g_i^a) - i\sum_a \hat{m}^a g_i^a - i\sum_{a\leq b} \hat{q}^{ab} g_i^a g_i^b\right]\right\} \times \\
&\prod_{\mu=1}^{S-1} \left\{ \prod_a \left[\int \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a)\right] \times \right. \\
&\left. \exp\left[i\sum_a \hat{\Delta}_\mu^a \left(\Delta_\mu^a + \frac{\mathcal{R}_0}{\mathcal{R}} m^a - \chi\right) - \frac{1}{2}\sum_{a,b}\left(\frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{\mathcal{R}} q^{ab} + \epsilon^2\right)\hat{\Delta}_\mu^a \hat{\Delta}_\mu^b\right]\right\} \\
&\times \prod_a \left[\int \frac{d\Delta_P^a d\hat{\Delta}_P^a}{2\pi} \frac{d\Delta_M^a d\hat{\Delta}_M^a}{2\pi} \theta(-\Delta_P^a)\theta(-\Delta_M^a) \int_{-\infty}^{1} dg_1^a\right] \exp\left\{\sum_a \beta F_1(g_1^a)\right. \\
&- i\sum_a \hat{m}^a g_1^a - i\sum_{a\leq b} \hat{q}^{ab} g_1^a g_1^b + i\sum_a \hat{\Delta}_P^a \left(\Delta_P^a + \frac{\mathcal{R}_0}{\mathcal{R}} m^a + \left(1 - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g_1^a - \chi\right) \\
&+ i\sum_a \hat{\Delta}_M^a \left(\Delta_M^a + \frac{\mathcal{R}_0}{\mathcal{R}} m^a + \left(1 - \gamma - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g_R^a - \chi - \mathcal{R}_0 \Delta X\right) \\
&\left. - \frac{1}{2}\sum_{a,b}\left[\frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{\mathcal{R}}(q^{ab} - g_1^a g_1^b) + \epsilon^2\right]\left(\hat{\Delta}_P^a \hat{\Delta}_P^b + 2\hat{\Delta}_P^a \hat{\Delta}_M^b + \hat{\Delta}_M^a \hat{\Delta}_M^b\right)\right\}.
\end{aligned} \tag{S49}$$

So far, we have made no approximations from our initial partition function. Now, however, we invoke the saddle-point approximation, stating that the probability distribution for the order parameters and their Fourier conjugates is so heavily peaked that their values are

equivalent across replicas. In short, we assume

$$m^a = m, \tag{S50}$$

$$q^{aa} = q^D \tag{S51}$$

$$q^{ab} = q^O \text{ if } a \neq b, \tag{S52}$$

and likewise for the conjugate (hatted) order parameters. We recall that we must eventually choose our fitness gauge χ such that $m = 0$, a deviation from the analysis in Refs. (36) and (37) due to our choice of consumer resource model. Our saddle point approximation allows us to significantly simplify our integral, by eliminating the integrals over the replica-specific order parameters. This change decouples our large integral expression into a product of three *independent* integrals: one over the resource availabilities g_i for $i > 1$, one over the surpluses Δ^a_μ for all species besides the parent and the mutant, and the final integral over g_1 and the surpluses for the mutant and parent strains.

Because we assume a diverse, many-species ecosystem, the species besides the parent and mutant will be most important in setting the saddle-point values of our order parameters and their Fourier conjugates. Thus, we approximate that these saddle-point values are unchanged from the calculation in Refs. (36) and (37), which includes only the first two sets of integrals (i.e., those independent of the parent and mutant). While we will not restate their calculation here, we will summarize their results. If we assume that the values of q^D and q^O are similar –

$$q^D \approx q^O = q, \tag{S53}$$

$$q^D - q^O = \frac{\mathcal{R}}{\beta} x, \tag{S54}$$

then the saddle-point values of the conjugate variables are

$$\hat{m} \approx i\beta, \tag{S55}$$

$$\hat{q}^O \approx i \left(\frac{\beta}{x} \sqrt{\frac{q}{\mathcal{R}}} \right)^2, \tag{S56}$$

$$\hat{q}^D - \frac{\hat{q}^O}{2} \approx i\beta \frac{x-1}{2x}. \tag{S57}$$

These equations leave q and x undetermined, which can be solved for implicitly and approximated through numerical equations which we will discuss later.

We also make the approximation $q^O - g_1^a g_1^b \approx q^O$. The order parameter q is formed by adding together $g_i^a g_i^b$ for each of $\mathcal{R}$ resources, and we expect each of these quantities to be positive. So, the scale of $q - g_R^a g_R^b \sim q(\mathcal{R}-1)/\mathcal{R} \sim q$. After making this approximation, let us also recall that $F_i(g_i^a) = \kappa_i \left[g_i^a + (g_i^a)^2/2 \right]$. Knowing that $m = 0$ and $\hat{m} \approx i\beta$ causes all the terms linearly dependent on the g_i^a to cancel out, besides the mutation-dependent g_1^a terms: thus, before considering the effects of the mutation, our system is invariant under rotation of the g_i^a's. Practically, this means that our initial choice of a knock-out mutation was arbitrary: a knock-in mutation, or any strategy change $\Delta\vec{\alpha}$ with the same overall

magnitude, can be redefined as a change in a single element through an appropriate rotation. Our results should therefore be valid for mutations which affect multiple resources simultaneously.

2.5 Joint Distribution of Δ_P and Δ_M

Now that we have decoupled our integrals through the saddle-point approximation, we are nearly ready to analyze the phenomenon of parent-mutant coexistence. To do this, we only care about the part of the partition function which depends on those strains. For simplicity, we will also at this point assume uniform resource supply ($\kappa_i = 1$); we will discuss at the end of our calculation what happens when this assumption is relaxed. With the saddle-point approximation for $\kappa_i = 1$, the portion of the integral in Eq. (S49) which depends on the parent and mutant strains is

$$
\begin{aligned}
Z^n_{PM} = \prod_a \left[\int \frac{d\Delta^a_P d\hat{\Delta}^a_P}{2\pi} \frac{d\Delta^a_M d\hat{\Delta}^a_M}{2\pi} \theta(-\Delta^a_P)\theta(-\Delta^a_M) \int_{-\infty}^{1} dg^a_1 \right] \times \\
\exp\Bigg\{ \sum_a \left[-\frac{\beta}{2x}(g^a_1)^2 \right] + \frac{1}{2}\left(\frac{\beta}{x}\sqrt{\frac{q}{\mathcal{R}}} \sum_a g^a_1 \right)^2 \\
+ i \sum_a \hat{\Delta}^a_P \left(\Delta^a_P + \frac{\mathcal{R}_0}{\mathcal{R}} m + \left(1 - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g^a_1 - \chi \right) \\
+ i \sum_a \hat{\Delta}^a_M \left(\Delta^a_M + \frac{\mathcal{R}_0}{\mathcal{R}} m + \left(1 - \gamma - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g^a_1 - \chi - \mathcal{R}_0 \Delta X \right) \\
- \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\beta} x \sum_a \left(\hat{\Delta}^a_P + \hat{\Delta}^a_M \right)^2 \\
- \frac{1}{2} \sum_{a,b} \left[\frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{\mathcal{R}} q + \epsilon^2 \right] \left(\hat{\Delta}^a_P \hat{\Delta}^b_P + 2\hat{\Delta}^a_P \hat{\Delta}^b_M + \hat{\Delta}^a_M \hat{\Delta}^b_M \right) \Bigg\}.
\end{aligned} \tag{S58}
$$

Next, we perform a change of variables to $\hat{\Delta}^a_\pm = \hat{\Delta}^a_P \pm \hat{\Delta}^a_M$:

$$
\begin{aligned}
Z^n_{PM} = \prod_a \left[\int \frac{d\Delta^a_P d\Delta^a_M}{4\pi} \frac{d\hat{\Delta}^a_+ d\hat{\Delta}^a_-}{4\pi} \theta(-\Delta^a_P)\theta(-\Delta^a_M) \int_{-\infty}^{1} dg^a_1 \right] \times \\
\exp\Bigg\{ \sum_a \left[-\frac{\beta}{2x}(g^a_1)^2 \right] + \frac{1}{2}\left(\frac{\beta}{x}\sqrt{\frac{q}{\mathcal{R}}} \sum_a g^a_1 \right)^2 \\
+ i \sum_a \hat{\Delta}^a_+ \left(\frac{\Delta^a_P + \Delta^a_M}{2} + \frac{\mathcal{R}_0}{\mathcal{R}} m + \left(1 - \frac{\gamma}{2} - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g^a_1 - \chi - \mathcal{R}_0 \frac{\Delta X}{2} \right) \\
+ i \sum_a \hat{\Delta}^a_- \left(\frac{\Delta^a_P - \Delta^a_M}{2} + \frac{\gamma}{2} g^a_1 + \mathcal{R}_0 \frac{\Delta X}{2} \right) - \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\beta} x \sum_a \left(\hat{\Delta}^a_+ \right)^2 \\
- \frac{1}{2} \sum_{a,b} \left[\frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{\mathcal{R}} q + \epsilon^2 \right] \left(\sum_a \hat{\Delta}^a_+ \right)^2 \Bigg\}.
\end{aligned} \tag{S59}
$$

We now aim to decouple the replicas, by eliminating terms which involve more than one replica. All these terms are of the form $(\sum_a f^a)^2$, where f^a is some quantity which depends only on one replica. We can decouple these replicas using the identity

$$\exp\left(\frac{1}{2}(Cx)^2\right) = \int d\omega \frac{e^{-\omega^2/2}}{\sqrt{2\pi}} e^{Cx\omega}. \tag{S60}$$

Thus, we can decouple the replicas at the cost of introducing an additional Gaussian integration variable per square term. We can now drop the a indices and replace them with a simple power to n:

$$\begin{aligned} Z_{PM}^n = \int d\omega_1 d\omega_2 \Bigg\{ \int \frac{d\Delta_P d\Delta_M}{4\pi} \frac{d\hat{\Delta}_+ d\hat{\Delta}_-}{4\pi} \theta(-\Delta_P)\theta(-\Delta_M) \int_{-\infty}^{1} dg_1 \\ \times \exp\Bigg[-\frac{\beta}{2x} g_1^2 + \omega_1 \frac{\beta}{x}\sqrt{\frac{q}{\mathcal{R}}} g_1 - \frac{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})}{2\beta} x \hat{\Delta}_+^2 \\ + i\hat{\Delta}_+ \left(\frac{\Delta_P + \Delta_M}{2} + \frac{\mathcal{R}_0}{\mathcal{R}} m + \left(1 - \frac{\gamma}{2} - \frac{\mathcal{R}_0}{\mathcal{R}}\right) g_1 + \omega_2 V_{\text{tot}}^{1/2} - \chi - \mathcal{R}_0 \frac{\Delta X}{2} \right) \\ + i\hat{\Delta}_- \left(\frac{\Delta_P - \Delta_M}{2} + \frac{\gamma}{2} g_1 + \mathcal{R}_0 \frac{\Delta X}{2} \right) \Bigg] \Bigg\}^n \frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi}. \end{aligned} \tag{S61}$$

Here, we have defined $V_{\text{tot}} = q\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})/\mathcal{R} + \epsilon^2$. This parameter, defined as ψ^2 in Refs. (36) and (37), represents the total variance in fitness among sampled species from both strategy and pure fitness. As we will demonstrate later, the integration variables ω_1 and ω_2 do in fact have physical meaning: ω_1 is related to the invasion fitness of the mutant, while ω_2 is related to the abundance of the parent. Next, we perform the integral over g_1. To do so, we assume that most of the weight of the integral is concentrated near small g, meaning that we can approximate the integration bounds as $-\infty$ to ∞ rather than $-\infty$ to 1. We also condense all of the resulting terms which do not depend on $\hat{\Delta}_+$ or $\hat{\Delta}_-$ into a constant term c.

We can write the result as

$$\begin{aligned} Z_{PM}^n = \int d\omega_1 d\omega_2 \Bigg\{ \int \frac{d\Delta_P d\Delta_M}{4\pi} \frac{d\hat{\Delta}_+ d\hat{\Delta}_-}{4\pi} \theta(-\Delta_P)\theta(-\Delta_M) \\ \times \exp\left[-\frac{1}{2}\vec{\Delta}\mathbf{M}\vec{\Delta}^t + i\vec{\Delta}\cdot\vec{v} + c \right] \Bigg\}^n \frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi}, \end{aligned} \tag{S62}$$

where

$$\vec{\Delta} = \begin{bmatrix} \hat{\Delta}_+ & \hat{\Delta}_- \end{bmatrix}, \tag{S63}$$

$$\mathbf{M} = \frac{x}{\beta} \begin{bmatrix} A & B \\ B & C \end{bmatrix}, \tag{S64}$$

$$\vec{v} = \begin{bmatrix} \frac{\Delta_P + \Delta_M}{2} + D & \frac{\Delta_P - \Delta_M}{2} + E \end{bmatrix}, \tag{S65}$$

and we define

$$A = \mathcal{R}_0\left(1-\frac{\mathcal{R}_0}{\mathcal{R}}\right)+\left(1-\frac{\gamma}{2}-\frac{\mathcal{R}_0}{\mathcal{R}}\right)^2, \tag{S66}$$

$$B = \frac{\gamma}{2}\left(1-\frac{\gamma}{2}-\frac{\mathcal{R}_0}{\mathcal{R}}\right), \tag{S67}$$

$$C = \frac{\gamma^2}{4}, \tag{S68}$$

$$D = \omega_2 V_{\text{tot}}^{1/2}+\frac{\mathcal{R}_0}{\mathcal{R}}m+\omega_1\sqrt{\frac{q}{\mathcal{R}}}\left(1-\frac{\gamma}{2}-\frac{\mathcal{R}_0}{\mathcal{R}}\right)-\chi-\mathcal{R}_0\frac{\Delta X}{2}, \tag{S69}$$

$$E = \mathcal{R}_0\frac{\Delta X}{2}+\omega_1\frac{\gamma}{2}\sqrt{\frac{q}{\mathcal{R}}}. \tag{S70}$$

We can now perform the integrals over $\hat{\Delta}_+$ and $\hat{\Delta}_-$. This yields

$$Z_{PM}^n = \int d\omega_1 d\omega_2\left[\frac{1}{\sqrt{\det\mathbf{M}}}\int\frac{d\Delta_P d\Delta_M}{4\pi}\theta(-\Delta_P)\theta(-\Delta_M)\times \exp\left(-\frac{1}{2}\vec{v}\mathbf{M}^{-1}\vec{v}^t+c\right)\right]^n\frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi}. \tag{S71}$$

Now, let's utilize the replica trick limit $n \to 0$. Again following the methodology in Refs. (36) and (37), this can be used to recover the joint probability distribution $\rho(\Delta_P, \Delta_M)$:

$$\lim_{n\to 0} Z_{PM}^n = \int_{-\infty}^0 d\Delta_P d\Delta_M \int d\omega_1 d\omega_2 \frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi}\times \frac{\exp\left(-\frac{1}{2}\vec{v}\mathbf{M}^{-1}\vec{v}^t\right)}{\int_{-\infty}^0 d\Delta_P' d\Delta_M' \exp\left(-\frac{1}{2}\vec{v}'\mathbf{M}^{-1}\vec{v}'^t\right)}, \tag{S72}$$

$$= \int_{-\infty}^0 d\Delta_P d\Delta_M \rho(\Delta_P, \Delta_M). \tag{S73}$$

Recall that $\mathbf{M}^{-1} \sim \beta$, an inverse temperature which scales with the number of cells in our community, reflecting the shot noise in our population at any given time. If we assume that our population is always exactly at equilibrium (i.e., is large enough that shot noise is negligible), we can take the $\beta \to \infty$ limit, in which case the exponentials involving $\mathbf{M}^{-1}$ will be very strongly peaked. In this limit, we can therefore approximate the ratio in Eq. (S72) as a delta function:

$$\rho(\Delta_P, \Delta_M) = \int d\omega_1 d\omega_2 \frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi}\delta(\Delta_P-\Delta_P^*)\delta(\Delta_M-\Delta_M^*), \tag{S74}$$

where $\Delta_P^*(\omega_1, \omega_2)$ and $\Delta_M^*(\omega_1, \omega_2)$ are the non-positive values of Δ_P and Δ_M that minimize $\vec{v}\mathbf{M}^{-1}\vec{v}^t$ for a given ω_1 and ω_2.

2.6 Parent-Mutant Coexistence Probability

The parent-mutant coexistence probability is encoded in the joint distribution $\rho(\Delta_P, \Delta_M)$:

$$\mathbb{P}[\text{Both survive}] = \int d\omega_1 d\omega_2 \frac{e^{-(\omega_1^2+\omega_2^2)/2}}{2\pi} \delta(\Delta_P^*(\omega_1,\omega_2))\delta(\Delta_M^*(\omega_1,\omega_2)). \tag{S75}$$

So, the next step is to uncover the functional forms of Δ_P^* and Δ_M^*, which represent the values of Δ_P and Δ_M that minimize $\vec{v}\mathbf{M}^{-1}\vec{v}^t$, constrained such that both Δ_P and Δ_M are non-positive. The matrix $\mathbf{M}$ and its inverse are positive definite, so this is a convex optimization problem, and either the unconstrained minimum lies in the interior of the allowed region, or the function is optimized on the boundary. If the global minimum happens to be within the allowed region, corresponding to both the parent and the mutant being extinct, these quantities are simple:

$$\Delta_P^*\big|_{\Delta_P^*<0,\ \Delta_M^*<0} = -D - E, \tag{S76}$$

$$\Delta_M^*\big|_{\Delta_P^*<0,\ \Delta_M^*<0} = -D + E. \tag{S77}$$

This unconstrained minimum will be in the allowed region if both of these quantities are negative, which occurs when $D > |E|$. We see that when both the parent and mutant are extinct, D tells us the resource deficit averaged between the parent and mutant, while $2E$ is the difference in resource deficit between mutant and parent. If the average resource deficit is too large, then neither strain can survive, thus

$$\mathbb{P}[\text{Both die}] = \mathbb{P}[D > |E|]. \tag{S78}$$

If the unconstrained minimum lies outside the allowed region, then the function will be minimized on the boundary of the allowed region, when one or both of the parent and mutant are alive. For example, if the mutant survives and the parent goes extinct, the minimum will be found at $\Delta_M^* = 0$ and

$$\Delta_P^*\big|_{\Delta_P^*<0,\ \Delta_M^*=0} = \frac{2B(D+E) - 2(CD+AE)}{A-2B+C}, \tag{S79}$$

$$\approx -2E + \frac{2}{A}(B-C)(D-E). \tag{S80}$$

This expansion comes from the fact that $A \sim \mathcal{R}_0$ (provided that $\mathcal{R}_0$ does not get too close to $\mathcal{R}$), while $B,\ C \sim 1$. Similarly, if the mutant goes extinct while the parent survives, we find

$$\Delta_M^*\big|_{\Delta_P^*=0,\ \Delta_M^*<0} = \frac{2E(A+B) - 2D(B+C)}{A+2B+C}, \tag{S81}$$

$$\approx 2E - \frac{2}{A}(B+C)(D+E). \tag{S82}$$

The last possibility is that the optimum lies on the corner of the allowed region, where the parent and mutant coexist (so $\Delta_P^* = \Delta_M^* = 0$). We can find this probability by requiring

that each of these three potential optima are outside of the allowed region, which occurs when

$$\mathbb{P}[\text{Both survive}] \approx \mathbb{P}[BE + CD < -|AE - BD - CE|,\ D < |E|]. \tag{S83}$$

Recalling that $A \sim \mathcal{R}_0 \gg 1$, the AE term dominates unless $|E| \ll 1$, meaning that E must be small for the inequality to be possible. (Note that the $\mathcal{R}_0 \Delta X$ term is not actually $\mathcal{O}(\mathcal{R}_0)$, because the natural fitness scale for a mutant turns out to be $\Delta X \sim \mathcal{R}_0^{-3/2}$.) When this is the case, the BE and CE terms will be even smaller, allowing us to safely ignore them:

$$\mathbb{P}[\text{Both survive}] \approx \mathbb{P}[CD < -|AE - BD|,\ D < |E|], \tag{S84}$$

$$= \mathbb{P}[CD < -|AE - BD|]. \tag{S85}$$

We have eliminated the second condition because it is now redundant with the first: since $C > 0$, the first inequality requires $D < 0$, trivially satisfying $D < |E|$. Finally, let's recall that in our simultaneous assembly approximation of mutant invasion (Appendix 1.3), we must condition on the probability that the mutant survives:

$$\mathbb{P}_{\text{coex}} \approx \frac{\mathbb{P}[\text{Both survive}]}{\mathbb{P}[\text{Mutant survives}]}, \tag{S86}$$

$$\approx \frac{\mathbb{P}[CD < -|AE - BD|]}{\mathbb{P}[E > D, E > 0]}. \tag{S87}$$

In principle, this is a quantity that can be calculated numerically. However, this result is difficult to intuitively interpret in this form, and it still depends on the saddle-point order parameter q. Nonetheless, we can immediately make two meaningful conclusions. If we send the number of resources consumed by a typical organism $\mathcal{R}_0 \to \infty$, then A becomes large and $\mathbb{P}_{\text{coex}} \to 0$. Likewise, if we send the effect size of our mutation $\gamma \to 0$, then $\frac{|CD|}{|AE-BD|} \sim \gamma \to 0$, and thus $\mathbb{P}_{\text{coex}} \to 0$. Since we are considering a mutation whose effect size is comparable to changing one resource's uptake rate by a factor of γ, both of these observations suggest the same thing: if the phenotypic effect of a mutation is too small, mutant-parent coexistence is impossible. This phenomenon is a signature of "mesoscopic" mutation, suggesting that our results may not be captured by models of mutations with infinitesimal phenotypic effects, such as adaptive dynamics approaches (44, 65).

3 Scaling Analysis of First-Step Mutations

In this section, we aim to restore physical intuition to our replica-theoretic results, finding approximate expressions which indicate how the frequency of inter-niche evolution scales with community parameters such as the number of surviving species.

3.1 Sampling Depth and Fitness Gauge

We must utilize the saddle-point results for an assembled community without mutation, largely from Refs. (36) and (37), to calculate our order parameters m and q. However, we opt for a different set of independent variables, defining our assembly process in terms

of the number of sampled species $\mathcal{S}$ and the typical number of surviving species $\mathcal{S}^*$. The community assembly parameter ϵ, which describes the variation in pure fitness among sampled species, will then be defined implicitly as a function of these parameters. This change of variables is helpful because $\mathcal{S}^*$ is a measurable property of an experimental ecosystem, while ϵ describes a property of *sampled* species, which may be very different from typical surviving species and is thus more difficult to access experimentally.

With this set of variables, a natural ratio to consider is the "sampling permissivity" $\mathcal{S}^*/\mathcal{S}$, roughly the probability that a new independently sampled species would be able to invade the ecosystem. It turns out that it will be useful to define an effective "sampling depth" λ, which is a function of $\mathcal{S}^*/\mathcal{S}$:

$$\lambda \equiv \mathrm{erf}^{-1}\left(1 - 2\frac{\mathcal{S}^*}{\mathcal{S}}\right)\sqrt{2}. \tag{S88}$$

For an undersampled system ($\mathcal{S}^*/\mathcal{S} \to 1$), the sampling depth $\lambda \to -\infty$; for an oversampled system ($\mathcal{S}^*/\mathcal{S} \to 0$), we have $\lambda \to \infty$. From the saddle-point results of Refs. (36) and (37) (which we do not re-derive here), we know that

$$m = \frac{\mathcal{R}}{\mathcal{R}_0}\left(V_{\mathrm{tot}}^{1/2}\lambda + \chi\right). \tag{S89}$$

Recall that χ sets our choice of fitness gauge: the uptake budget of sampled species is drawn with mean $\chi/\mathcal{R}_0$ and standard deviation $\epsilon/\mathcal{R}_0$. But in our community assembly approximations, we have chosen our gauge such that $\langle g_i \rangle = m/\mathcal{R} = 0$. This implies that

$$\chi = -V_{\mathrm{tot}}^{1/2}\lambda. \tag{S90}$$

As we sample more species and λ increases, the uptake budgets or pure fitnesses of surviving species will be further and further into the high-fitness tail of the distribution. Since the typical pure fitness of surviving species is nearly zero in our fitness gauge, χ must become more negative as λ increases to compensate for the increased sampling. In the consumer resource model analyzed in Refs. (36) and (37) (where $\chi = 0$), an increase in sampling depth results in an overall increase in m: as oversampled species become higher and higher in pure fitness, the typical resource becomes less valuable to consume. This effect causes resource knock-out mutations to be more beneficial on average than knock-in mutations in the oversampled limit. In our model, we adjust the "zero" of pure fitness to compensate for this effect, such that $m = 0$ regardless of λ.

For future sections, it is helpful to define a function $I(\lambda)$ such that

$$I(\lambda) = \frac{1+\lambda^2}{2}\mathrm{erfc}\left(\frac{\lambda}{\sqrt{2}}\right) - \frac{\lambda}{\sqrt{2\pi}}e^{-\lambda^2/2}. \tag{S91}$$

3.2 Niche Saturation and σ_{inv}

From saddle-point results in Refs. (36) and (37), we can write expressions for V_{tot}, q, and ϵ in terms of quantities we have already defined, as well as variance in resource supply

$\sigma_K^2 \equiv \mathrm{Var}(\kappa) = \mathrm{Var}(K)/\overline{K}^2$:

$$V_{\mathrm{tot}}^{1/2} = \left(1 - \frac{\mathcal{R}_0}{\mathcal{R}}\right)\left(\frac{2\left(1 - \frac{S}{\mathcal{R}}I(\lambda)\right)}{-\frac{S}{\mathcal{R}}I'(\lambda)} - \lambda\right), \tag{S92}$$

$$q = V_{\mathrm{tot}}\frac{S}{\mathcal{R}}I(\lambda)\frac{\mathcal{R}}{\mathcal{R}_0(1-\mathcal{R}_0/\mathcal{R})} + \mathcal{R}\sigma_K^2\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)^2, \tag{S93}$$

$$\epsilon^2 = V_{\mathrm{tot}}\left(1 - \frac{S}{\mathcal{R}}I(\lambda)\right) - \mathcal{R}_0\sigma_K^2\left(1 - \frac{\mathcal{R}_0}{\mathcal{R}}\right)\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)^2. \tag{S94}$$

The order parameter q describes the total variance in resource availability, allowing us to find the distribution of the resource availabilities h_i: a Gaussian with mean 1 and variance $q/\mathcal{R}$. A mutation changing use of a single resource by a factor γ will have an invasion fitness with magnitude $\gamma|1-h_i|/\mathcal{R}_0$. This leads us to define the quantity

$$\sigma_{\mathrm{inv}} \equiv \frac{\gamma}{\mathcal{R}_0}\sqrt{\frac{q}{\mathcal{R}}}, \tag{S95}$$

which describes the typical magnitude of the invasion fitness of a strategy mutant $\vec{\alpha} \to \vec{\alpha} + \Delta\vec{\alpha}$, where $|\Delta\vec{\alpha}| \approx \gamma/\mathcal{R}_0$. (For example, for a complete knock-out mutant for a system with binary resource use, we have $\gamma = 1$.)

While these equations completely determine all the variables we have defined, they are difficult to interpret physically, having complicated dependence on the ratios $\mathcal{S}/\mathcal{R}$ and $\mathcal{R}_0/\mathcal{R}$ as well as the sampling depth λ. We can make progress by considering how these equations depend on the *niche saturation* $\mathcal{S}^*/\mathcal{R} \leq 1$, describing how close to maximum capacity we expect our final ecosystem to be. By rewriting $\mathcal{S}/\mathcal{R} = (\mathcal{S}/\mathcal{S}^*)(\mathcal{S}^*/\mathcal{R})$, we separate out the parts of these equations that depend on the sampling depth λ (or equivalently, $\mathcal{S}^*/\mathcal{S}$) from the niche saturation $\mathcal{S}^*/\mathcal{R}$. (Technically, rather than using the value of $\mathcal{S}^*$ for any particular ecosystem as an input parameter, we use its expectation $\langle\mathcal{S}^*\rangle$ over random ecosystems with the same underlying parameters. Since these values tend to be close in practice, as shown in Figure 2B, we will simply write $\mathcal{S}^*$ in our scaling relationships.)

To simplify our results, it is helpful to expand our expressions in the undersampled and oversampled limits ($\lambda \to \pm\infty$). It is worth noting that the $\lambda \to -\infty$ limit for fixed $\mathcal{S}^*/\mathcal{R}$ does not correspond to a physical community assembly process: the expression for ϵ in Eqn. (S94) must be positive, but $(\mathcal{S}/\mathcal{R}) \cdot I(\lambda) \to \infty$ as $\lambda \to -\infty$, making the RHS of that equation negative. Effectively, this means that in order to sample species such that a desired fraction of niches are filled, a minimum fraction of sampled species must go extinct in the final ecosystem (on average). However, many quantities we are interested in are still well-defined in the $\lambda \to -\infty$ limit, and indeed scale similarly to the $\lambda \to \infty$ limit: we use this to demonstrate an overall lack of dependence on sampling depth.

We first quote results for $\sigma_K^2 = 0$. From standard expansions of the error function, we obtain

$$\sigma_{\mathrm{inv}} \approx \begin{cases} \gamma\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)\sqrt{\frac{\mathcal{R}-\mathcal{R}_0}{\mathcal{S}^*\mathcal{R}_0^3}} & \text{when } \lambda \to -\infty, \\ \gamma\sqrt{2}\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)\sqrt{\frac{\mathcal{R}-\mathcal{R}_0}{\mathcal{S}^*\mathcal{R}_0^3}} & \text{when } \lambda \to \infty. \end{cases} \tag{S96}$$

Remarkably, the typical invasion fitness scale of a strategy mutation has very weak dependence on the sampling depth λ, changing only by a factor of $\sqrt{2}$ from the extremely undersampled to the extremely oversampled regimes. This result allows us to effectively drop out λ (and thus $\mathcal{S}^*/\mathcal{S}$) as a variable, which was very difficult to see from our saddle-point results above.

3.3 Scaling Analysis of $\mathbb{P}_{\text{coex}}$

We got our expression for σ_{inv} "for free" from the saddle-point value of q, without need to explicitly consider the impact of the parent and mutant on the ecosystem. The same is not true for the mutant-parent coexistence probability: to obtain an interpretable expression for $\mathbb{P}_{\text{coex}}$, we must start from our results in Appendix 2.5, and then undergo a similar procedure of plugging in saddle-point values in the $\lambda \to \pm\infty$ limits. Recall that in our two-ecosystem approximation, we have

$$\mathbb{P}_{\text{coex}} \approx \frac{\mathbb{P}[\text{Both survive}]}{\mathbb{P}[\text{Mutant survives}]}. \tag{S97}$$

To get some intuition, let's start with the simpler calculation of the denominator of this expression: the probability that the mutant species survives in our assembled ecosystem. From our earlier results, this is

$$\begin{aligned}\mathbb{P}_{\text{M lives}} = \mathbb{P}\Bigg[&\Delta X + \mathcal{Z}_1\sigma_{\text{inv}} > 0,\\ &-\mathcal{Z}_2 V_{\text{tot}}^{1/2} + V_{\text{tot}}^{1/2}\lambda + \mathcal{Z}_1\frac{\sigma_{\text{inv}}\mathcal{R}_0}{\gamma}\left(1-\gamma-\frac{\mathcal{R}_0}{\mathcal{R}}\right) - \mathcal{R}_0\Delta X < 0\Bigg],\end{aligned} \tag{S98}$$

where $\mathcal{Z}_i$ are independent standard normal random variables. Note that we have added a negative sign to $\mathcal{Z}_2$; while this leaves its distribution unchanged, it will aid in interpretation. The first two terms in the second inequality scale with the fitness variation in all sampled species $V_{\text{tot}}^{1/2}$, while the remaining terms scale with the fitness change induced by the mutation. If the mutation results in a small change in the organism's overall fitness, the first two terms should dominate, so we can approximate

$$\mathbb{P}_{\text{M lives}} \approx \mathbb{P}\left[\Delta X + \mathcal{Z}_1\sigma_{\text{inv}} > 0,\ \mathcal{Z}_2 - \lambda > 0\right], \tag{S99}$$

$$\approx \frac{\mathcal{S}^*}{2S}\text{erfc}\left(-\frac{\Delta X}{\sigma_{\text{inv}}\sqrt{2}}\right). \tag{S100}$$

Each of these inequalities has a meaning which allows us to interpret $\mathcal{Z}_{1,2}$ as physical quantities. The first inequality represents the probability the mutant is beneficial relative to the parent: the chances that the deterministic fitness change ΔX combined with the stochastic fitness change $\mathcal{Z}_1\sigma_{\text{inv}}$ due to the strategy mutation is positive. The second inequality represents the requirement that the parent background the mutant arises from is well-adapted enough to survive in the population in the first place: only draws of $\mathcal{Z}_2$ which are greater than λ will survive in the population. Indeed, the probability that $\mathcal{Z}_2 > \lambda$ is simply $\mathcal{S}^*/\mathcal{S}$,

the probability a randomly chosen species would survive in the community. Neglecting the other terms in the second inequality neglects the possibility that the parent cannot survive in the initial population, but the mutant can; as discussed in SI Section 2.1, we expect this outcome to be rare.

We posit that $\mathcal{Z}_2$ determines not only the survival of the parent species, but also its approximate relative abundance f_P. Past research (29) suggests that consumer resource models of this form yield a truncated Gaussian distribution of abundances, which (up to an overall normalization factor) coincides with the distribution of $\mathcal{Z}_2$ conditioned on being greater than λ. We can find this normalization factor by enforcing $\langle f_P | f_P > 0 \rangle = 1/\mathcal{S}^*$, which yields

$$f_P \approx \frac{2}{-\mathcal{S} I'(\lambda)} (\mathcal{Z}_2 - \lambda)\theta(\mathcal{Z}_2 - \lambda). \tag{S101}$$

In the oversampled limit $\lambda \gg 1$, the prefactor is approximately $\lambda/\mathcal{S}^*$. Our results in the next section confirm that using this relationship between $\mathcal{Z}_2$ and f_P yields reasonably accurate predictions.

Now, let's calculate the numerator of $\mathbb{P}_{\text{coex}}$, the probability that the mutant and parent will both be alive in our simultaneously assembled ecosystem. Again approximating that terms which scale with overall fitness variation $V_{\text{tot}}^{1/2}$ dominate terms which scale with mutant fitness change ΔX and σ_{inv}, we find that

$$\mathbb{P}_{\text{P,M live}} \approx \mathbb{P}\left[\frac{\gamma V_{\text{tot}}^{1/2}}{\mathcal{R}_0^2}(\mathcal{Z}_2 - \lambda) > \Delta X + \mathcal{Z}_1 \sigma_{\text{inv}} > \frac{\gamma(1 - \gamma - \frac{\mathcal{R}_0}{\mathcal{R}}) V_{\text{tot}}^{1/2}}{\mathcal{R}_0^2 (1 - \frac{\mathcal{R}_0}{\mathcal{R}})}(\mathcal{Z}_2 - \lambda)\right]. \tag{S102}$$

Let us analyze this probability using our physical interpretations of $\mathcal{Z}_1$ and $\mathcal{Z}_2$. The outer inequality only holds when $\mathcal{Z}_2 > \lambda$, again requiring that the parent species survives. The middle term is the leading contribution to the invasion fitness of the mutant, which is $\mathcal{O}(\mathcal{R}_0^{-3/2})$. However, the inequality bounds it between two quantities of size $\mathcal{O}(\mathcal{R}_0^{-2})$. So, we conclude that parent-mutant coexistence is achieved when the invasion fitness of the mutant is sufficiently small.

Since the allowed region for $\mathcal{Z}_1$ is smaller than its typical width by a factor of $1/\sqrt{\mathcal{R}_0}$, we can approximate its probability density as constant over the allowed region. This gives us

$$\mathbb{P}_{\text{P,M live}} \approx \frac{e^{-(\Delta X/\sigma_{\text{inv}})^2/2}}{\sqrt{2\pi}} \frac{\gamma^2 V_{\text{tot}}^{1/2}}{\sigma_{\text{inv}} \mathcal{R}_0^2 (1 - \frac{\mathcal{R}_0}{\mathcal{R}})} \int_\lambda^\infty (\mathcal{Z}_2 - \lambda) \frac{e^{-\mathcal{Z}_2^2/2}}{\sqrt{2\pi}} d\mathcal{Z}_2 \tag{S103}$$

$$= \frac{e^{-(\Delta X/\sigma_{\text{inv}})^2/2}}{\sqrt{2\pi}} \frac{\gamma^2 V_{\text{tot}}^{1/2}}{\sigma_{\text{inv}} \mathcal{R}_0^2 (1 - \frac{\mathcal{R}_0}{\mathcal{R}})} \frac{-I'(\lambda)}{2}. \tag{S104}$$

Finally, we divide by the probability the mutant lives, and consider the undersampled and oversampled limits when $\Delta X = 0$:

$$\mathbb{P}_{\text{coex}} \approx \begin{cases} \gamma \sqrt{\frac{2\mathcal{R}}{\pi \mathcal{S}^*}} \sqrt{\frac{1}{\mathcal{R}_0(1 - \mathcal{R}_0/\mathcal{R})}} & \text{when } \lambda \to -\infty, \\ \gamma \sqrt{\frac{\mathcal{R}}{\pi \mathcal{S}^*}} \sqrt{\frac{1}{\mathcal{R}_0(1 - \mathcal{R}_0/\mathcal{R})}} & \text{when } \lambda \to \infty. \end{cases} \tag{S105}$$

Once again, the dependence on the sampling depth is nearly negligible, contributing only an $\mathcal{O}(1)$ numerical factor. Note, however, that this result is different from that in the main text in Equation (9). This is because in our theoretical calculation, we arbitrarily chose the identity of the parent species and the mutation, effectively giving each surviving species and each mutation an equal chance of being selected. In reality, each surviving species produces mutations at a rate proportional to its abundance, and mutations have a probability proportional to their invasion fitness of surviving genetic drift (52). In the next section, we will discuss how to correct for this difference.

3.4 Dependence of $\mathbb{P}_{\text{coex}}$ on Invasion Fitness and Abundance

Let us reframe our expression for $\mathbb{P}_{\text{coex}}$ in terms of the abundance f_P of the parent and the invasion fitness s_{inv} of the mutant, relating them to $\mathcal{Z}_1$ and $\mathcal{Z}_2$ as we did earlier. We can read off the invasion fitness of a beneficial mutant as the saddle-point value of Δ_M^* when Δ_P is fixed at zero, representing the initial growth rate of the mutant *before* it has a chance to invade:

$$s_{\text{inv}} \approx \Delta X + \mathcal{Z}_1 \sigma_{\text{inv}} + \frac{\gamma}{\mathcal{R}_0^2} V_{\text{tot}}^{1/2} (\mathcal{Z}_2 - \lambda). \tag{S106}$$

We can plug this expression into our earlier probability that both the mutant and parent live to find that

$$\mathbb{P}_{\text{P,M live}} \approx \mathbb{P}\left[0 < s_{\text{inv}} < \frac{\gamma^2 V_{\text{tot}}^{1/2}}{\mathcal{R}_0^2(1-\mathcal{R}_0/\mathcal{R})}(\mathcal{Z}_2 - \lambda)\right]. \tag{S107}$$

Recall that $\mathcal{Z}_2 - \lambda$ is proportional to the abundance of the parent, as we calculated in Eq. (S101) . While $\mathcal{Z}_2$ is not independent of s_{inv}, its contribution to s_{inv} is small, so approximating them as independent is reasonable. Plugging in this relationship, we find that

$$\mathbb{P}_{\text{coex}} \approx \mathbb{P}\left[s_{\text{inv}} < -I'(\lambda)\frac{\gamma^2 V_{\text{tot}}^{1/2}\mathcal{S}}{2\mathcal{R}_0^2(1-\mathcal{R}_0/\mathcal{R})} f_P\right], \tag{S108}$$

$$= \mathbb{P}\left[s_{\text{inv}} < \frac{\gamma^2(1-\mathcal{S}^*/\mathcal{R})}{\mathcal{R}_0^2(\mathcal{S}^*/\mathcal{R})}\mathcal{S}^* f_P\right]. \tag{S109}$$

So, our analysis of mutant-parent coexistence has an easily-interpreted meaning: a mutant strain will coexist with its parent when its (positive) invasion fitness is below

$$s_{\text{coex}}(f_P) = \frac{\gamma^2(1-\mathcal{S}^*/\mathcal{R})}{\mathcal{R}_0^2(\mathcal{S}^*/\mathcal{R})}\mathcal{S}^* f_P. \tag{S110}$$

We can calculate this probability by treating s_{inv} and f_P as independent random variables, accounting for the probability the mutant survives genetic drift. In this case, their probability densities scale as

$$\rho(s_{\text{inv}}) \sim s_{\text{inv}}\theta(s_{\text{inv}})e^{-(s_{\text{inv}}-\Delta X)^2/2\sigma_{\text{inv}}^2}, \tag{S111}$$

$$\rho(f_P) \sim f_P\theta(f_P)e^{-(f_P+D\lambda)^2/2D^2}, \tag{S112}$$

where $D \equiv -2/\mathcal{S}I'(\lambda)$. Finally, $\mathbb{P}_{\text{coex}}$ can be calculated using standard analytic integration methods, giving the results quoted in the main text after applying similar scaling analysis to the previous section. We can also fix either s_{inv} or f_P to calculate a conditional coexistence probability. For example, plugging in a "typical" $f_P = 1/\mathcal{S}^*$ indicates that $\bar{s}_{\text{coex}} = \frac{\gamma^2(1-\mathcal{S}^*/\mathcal{R})}{\mathcal{R}_0^2(\mathcal{S}^*/\mathcal{R})}$ is the typical value of s_{inv} above which the coexistence probability begins to fall.

3.5 Self-Consistency of Replica Approximations

Now that we have expressed the results of our theory in terms of physically interpretable parameters, we can use them to check when the assumptions underlying our analytic approximations are valid. The central assumption we made in SI Section 1.5 was that the resource availabilities h_i are not too far from 1, which should fail at some point in the vicinity of $|g_i| = |1 - h_i| \lesssim 1$. We calculated the typical spread of g_i in terms of σ_{inv} in SI Section 3.2, and found that

$$\text{Std}\,(g_i) \sim (1 - \mathcal{S}^*/\mathcal{R}) \left[\frac{1 - \mathcal{R}_0/\mathcal{R}}{(\mathcal{S}^*/\mathcal{R}) \cdot \mathcal{R}_0} + \sigma_K^2 \right]^{1/2}, \tag{S113}$$

up to an $\mathcal{O}(1)$ factor dependent on $\mathcal{S}^*/\mathcal{S}$. The first term in the brackets should be small over a wide range of parameter values, since $\mathcal{R}_0 \gg 1$. Notably, however, this term causes the $|g_i| \lesssim 1$ approximation to break down when the niche saturation is too small, around

$$\frac{\mathcal{S}^*}{\mathcal{R}} \lesssim \frac{1}{\mathcal{R}_0}. \tag{S114}$$

When this inequality is approached or satisfied, there will not be enough species in the community to cover the space of available resources ($\mathcal{R}_0\mathcal{S}^* \lesssim \mathcal{R}$), so the availability of each resource will be determined by $\mathcal{O}(1)$ rather than many species. Our simulations provide a test of how approaching this bound affects the numerical accuracy of our results: for example, the leftmost points in Fig. 3B (corresponding to $\mathcal{R}_0\mathcal{S}^* = 4\mathcal{R}$) show that the mutant-parent coexistence probability increases more slowly than expected as $\mathcal{S}^*/\mathcal{R}$ declines. Similar quantitative deviations from theory occur in more saturated communities when $\mathcal{R}_0$ is relatively small (e.g. the open circles in Fig. 2D), consistent with our scaling prediction. However, these deviations do not affect our qualitative conclusion that mutant-parent coexistence declines weakly with niche saturation over a wide range of community sizes. Rather, they indicate that in communities where each resource is exploited by only a small number of species, alternative theoretical approaches should be employed.

Self-consistency also requires $\sigma_K \lesssim 1$, so that the spread in resource supply is not too large. We find that our results hold reasonably well even for $\sigma_K = 1$ (Fig. S7B-C). However, different approximations must be employed when the abiotic environment has extremely large asymmetries in resource supply. One such example could be some models of cross-feeding, where an $\mathcal{O}(1)$ number of externally supplied metabolites result in the production of many other resources at much lower rates.

In our model, a consequence of small $|g_i|$ is that the total uptake budgets of surviving species have a maximum spread, $\mathrm{Std}\,(X_\mu|f_\mu > 0) \lesssim 1$, which justified the Taylor expansion in Eq. (S24). In other words, species whose uptake budget differences would drive large changes in their relative abundance over a single generation are unlikely to coexist in the community steady-state. Although this condition places bounds on the spread of X_μ among surviving species, it does not imply that the dynamics are effectively neutral (as in Ref. 32). True neutrality requires that fitness differences are small compared to inverse population sizes and inverse equilibration times; because we focus on large populations and long times, there is a large gap in scales between $X_\mu \lesssim 1$ and the emergence of effective neutrality.

The existence of this gap implies that even modest differences in the total uptake budget can still have a significant impact on community structure. Recall that in the simplest version of our model, the uptake budgets of sampled species are drawn from a normal distribution with standard deviation $\mathrm{Std}(X_\mu) = \epsilon/\mathcal{R}_0$. The total scaled fitness variation V_{tot} among sampled species comes from this variance in uptake budget along with resource consumption strategies. If the ratio $\epsilon^2/V_{\mathrm{tot}} \approx 1$, the fates of species will be strongly dependent on their uptake budgets; if $\epsilon^2/V_{\mathrm{tot}} \ll 1$, species fates will mostly depend on their specific resource consumption strategies. From the replica-theoretic results in SI Section 3.2, for uniform resource supply, we have

$$\frac{\epsilon^2}{V_{\mathrm{tot}}} \approx 1 - \frac{\mathcal{S}}{\mathcal{R}} I(\lambda), \tag{S115}$$

$$\approx 1 - \frac{2}{\lambda^2}\frac{\mathcal{S}^*}{\mathcal{R}} \text{ when } \lambda \gg 1. \tag{S116}$$

So, when enough species are sampled (large λ), total uptake budget will be an important determinant of species fate for any fixed value of $\mathcal{S}^*/\mathcal{R}$. Note that if the spread in pure fitness among sampled species is taken to be very small ($\epsilon \ll 1$), then differences in total uptake budget are negligible until enough species have been added to the community to saturate it ($\mathcal{S}^* \approx \mathcal{R}$), as in the top curve in Fig. 2B (36, 37). However, our replica-theoretic assumptions continue to hold for $\epsilon > 1$, because *surviving* species have a narrower spread of X_μ than *sampled* species:

$$\mathrm{Std}\,(X_\mu|f_\mu > 0) \approx \frac{\epsilon}{\mathcal{R}_0}\sqrt{\frac{\int_\lambda^\infty z^2\mathcal{D}z}{\int_\lambda^\infty \mathcal{D}z} - \left(\frac{\int_\lambda^\infty z\mathcal{D}z}{\int_\lambda^\infty \mathcal{D}z}\right)^2}, \tag{S117}$$

$$\approx \frac{1-\mathcal{S}^*/\mathcal{R}}{\mathcal{S}^*/\mathcal{R}} \cdot \frac{1-\mathcal{R}_0/\mathcal{R}}{\mathcal{R}_0} \text{ when } \lambda \gg 1, \tag{S118}$$

for uniform resource supply, where $\mathcal{D}z \equiv (2\pi)^{-1/2}e^{-z^2/2}dz$ is the standard Gaussian integration element. As expected, the spread in fitnesses becomes large when the spread in the g_i's does, around $\mathcal{S}^*/\mathcal{R} \lesssim 1/\mathcal{R}_0$, at which point the replica theory fails due to the number of species alive being too small for a "large-ecosystem" approximation.

4 Extensions of Model Assumptions

In this section, we test whether our qualitative conclusions strongly depend on the underlying assumptions of our model, such as the distribution of resource supplies and the method of species sampling. We find that many such changes can be captured by straightforward extensions of our theory, or simply by using the result from the binary consumer resource model directly.

4.1 Non-Uniform Resource Supply

First, let's consider non-uniform resource supply (nonzero σ_K^2), as in (37). If each resource has a supply $\kappa_i = 1 + \delta\kappa_i$, then the availability h_i of that resource will have a component which scales with $\delta\kappa_i$ in addition to a community-dependent component. Thus, mutations that knock out a highly-supplied resource will be less beneficial on average. We can determine the strength of this effect by starting with the integrals over the g_i^a in the second line of Eq. (S49), where the resource supply enters through $F_i(g_i) \equiv \sum_i \kappa_i(g_i + g_i^2/2)$. This portion of the partition function was labeled A_i by Ref. (37) (see their Appendix 2, section 1.3). We can follow their calculation to obtain

$$A_i = \int \frac{d\omega}{2\pi} e^{-\omega^2/2} \left\{ \int_{-\infty}^{1} dg_i \exp\beta \left[-\left(\frac{1}{2x} + \frac{\delta\kappa_i}{2}\right) g_i^2 - \left(\frac{\omega b}{\sqrt{R}} + \delta\kappa_i\right) g_i \right] \right\}^n, \tag{S119}$$

where b is a term independent of $\delta\kappa_i$ whose value does not concern us here. The contents of the brackets are a Gaussian integral over the g_i, centered where the integrand attains its maximum value. Taking non-uniform resource supply, $\delta\kappa_i \neq 0$, will shift this saddle-point value of g_i by approximately $x\delta\kappa_i \ll 1$. The saddle-point value of x was calculated in Ref. (36) as $x = 1 - \mathcal{S}^*/\mathcal{R}$, following our typical substitution of $(\mathcal{S}^*/\mathcal{R})(\mathcal{S}/\mathcal{S}^*)$ for $\mathcal{S}/\mathcal{R}$. This shift in the central value of g_i will shift the invasion fitness of a knock-out mutation targeting that resource by a corresponding amount, since the invasion fitness of a knock-out mutation is simply $g_i/\mathcal{R}_0$. Thus, we can write

$$s_{\text{inv}} = \Delta X - \frac{\delta\kappa_i}{\mathcal{R}_0}\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right) + \mathcal{Z}\sigma_{\text{inv}}(\delta\kappa_i = 0). \tag{S120}$$

for a mutation knocking out resource i. The first term represents the mutation's change in pure fitness, the second term reflects that environmental supply makes some mutations more or less beneficial on average, while the stochastic third term depends on the specific draws of species in the community, as we calculated earlier. We can quantify the relative importance of environment and community by considering the invasion fitness of a mutation in a monoculture community, where invasion fitness *only* depends on the environment:

$$s_{\text{inv}}^{\text{mono}} \approx \Delta X - \frac{\delta\kappa_i}{\mathcal{R}_0}. \tag{S121}$$

The correlation between the invasion fitness of a mutant in monoculture and in the community context illustrates how much the impact of environmental supply is felt through

the "shielding" effects of the community: in the oversampled limit ($\lambda \to \infty$), we find that

$$\rho(s_{\text{inv}}^{\text{comm}}, s_{\text{inv}}^{\text{mono}}) \approx \left(1 + \frac{\mathcal{R} - \mathcal{R}_0}{\mathcal{S}^* \mathcal{R}_0 \sigma_K^2}\right)^{-1/2}, \tag{S122}$$

assuming that $\delta\kappa_i$ are drawn from a distribution with mean zero and variance σ_K^2. In general, non-uniform resource supply will increase the variance in invasion fitness σ_{inv}:

$$\sigma_{\text{inv}} \approx \begin{cases} \gamma\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)\sqrt{\frac{\mathcal{R}-\mathcal{R}_0}{\mathcal{S}^*\mathcal{R}_0^3} + \frac{\sigma_K^2}{\mathcal{R}_0^2}} & \text{when } \lambda \to -\infty, \\ \gamma\sqrt{2}\left(1 - \frac{\mathcal{S}^*}{\mathcal{R}}\right)\sqrt{\frac{\mathcal{R}-\mathcal{R}_0}{\mathcal{S}^*\mathcal{R}_0^3} + \frac{\sigma_K^2}{2\mathcal{R}_0^2}} & \text{when } \lambda \to \infty. \end{cases} \tag{S123}$$

Note that σ_{inv} still scales with an overall factor of $1 - \mathcal{S}^*/\mathcal{R}$, leaving our previous qualitative results unaffected. Increasing resource supply variation increases σ_{inv} and thus lowers the coexistence probability, by making it less likely that the invasion fitness of an arbitrary beneficial mutation will be small enough to make coexistence possible. Although the theoretical approach assumes $\sigma_K^2 \ll 1$, its quantitative predictions hold reasonably well even when $\sigma_K^2 \sim 1$, as in the case of exponentially distributed K_i (Fig. S7B). However, when σ_K^2 increases enough that $\sigma_{\text{inv}} \sim 1/\mathcal{R}_0$, the distribution of fitness effects will be somewhat distorted from the Gaussian prediction (Fig. S7C). This is because resource availabilities h_i are non-negative and can only be approximated by a Gaussian when $\text{Std}(h_i) \ll \langle h_i \rangle = 1$.

4.2 Variation in Number of Metabolized Resources

In the binary resource model so far, we have assumed that each resource has an independent probability $\mathcal{R}_0/\mathcal{R}$ of being used by each species. This assumption results in each species using a binomially distributed number of resources, which is tightly peaked when $\mathcal{R}_0 \gg 1$. However, there is no reason to assume that coexisting species cannot differ in their levels of specialism, such that some use many resources and some few. First, it is useful to consider what happens if the parent species uses $\mathcal{R}_P$ resources on average, rather than the $\mathcal{R}_0$ resources that most species in the community use. This extension helps us disentangle what terms in our expressions depend on the typical metabolic overlap between members of the community (encoded by $\mathcal{R}_0$), and what depends on the specific number of resources used by the parent and mutant species. It can also help determine whether generalist or specialist species are more likely to diversify in a given community.

The more resources a given organism uses, the more its growth rate and overall fitness are averaged out between different independent variables, effectively reducing the variance in its abundance as well as the invasion fitness of any particular mutation. So, if $\mathcal{R}_P$ is larger, the parent species is less likely to have particularly high abundance, but the invasion fitness of a strategy mutation it obtains is likely to be smaller in magnitude. The first effect should lower the coexistence probability, while the second should raise it. Performing the

calculation up to our expression for the joint probability of Δ_P and Δ_M gives us

$$A \approx \mathcal{R}_0 \frac{\mathcal{R}_0}{\mathcal{R}_P}\left(1 - \frac{\mathcal{R}_P}{\mathcal{R}}\right), \tag{S124}$$

$$B \approx \frac{\mathcal{R}_0}{\mathcal{R}_P}\frac{\gamma}{2}\left[\frac{\mathcal{R}_0}{\mathcal{R}_P}\left(1 - \frac{\gamma}{2}\right) - \frac{\mathcal{R}_0}{\mathcal{R}}\right], \tag{S125}$$

$$C \approx \frac{\mathcal{R}_0^2}{\mathcal{R}_P^2}\frac{\gamma^2}{4}, \tag{S126}$$

$$D \approx \omega_2 V_P^{1/2} + \lambda V_{\text{tot}}^{1/2}, \tag{S127}$$

$$E \approx \mathcal{R}_0 \frac{\Delta X}{2} + \omega_1 \frac{\mathcal{R}_0}{\mathcal{R}_P}\frac{\gamma}{2}\sqrt{\frac{q}{\mathcal{R}}}. \tag{S128}$$

Here, we have defined $V_P = \frac{\mathcal{R}_0^2(1-\mathcal{R}_P/\mathcal{R})}{\mathcal{R}\mathcal{R}_P}q + \epsilon^2$, reflecting the different variation in overall fitness between the parent and other species in the community. We have once again defined γ such that $\gamma \approx 1$ for a knock-out mutation.

Since most of the community is unaffected by changing $\mathcal{R}_P$, we use the same saddle-point results we calculated before. Thus, the value of σ_{inv} is unchanged, although the invasion fitness of the mutant has a prefactor of $\mathcal{R}_0/\mathcal{R}_P$ to account for the altered number of resources. The inequality we obtain that predicts mutant-parent coexistence (for $\Delta X = 0$) is now given by

$$\mathbb{P}_{\text{P,M live}} \approx \mathbb{P}\left[\mathcal{Z}_1 < \frac{\gamma V_P^{1/2}}{\mathcal{R}_0(1-\mathcal{R}_P/R)\sqrt{q/R}}\left(\mathcal{Z}_2 - \lambda\frac{V_{\text{tot}}^{1/2}}{V_P^{1/2}}\right)\right], \tag{S129}$$

where $\mathcal{Z}_1$ is once again proportional to the invasion fitness of the mutant, and the $\mathcal{Z}_2$-dependent factor on the RHS is proportional to the abundance of the parent strain. The calculations otherwise proceed as before, yielding the results shown in Figure S9A. For these parameter values, it appears that mutant-parent coexistence is more likely for mutations on a background strain which metabolizes many resources, since these mutations tend to have small invasion fitness.

Next, we consider what happens when the number of resources $\mathcal{R}_0$ utilized by *every* species is more broadly distributed, to allow coexistence of generalists and specialists. Specifically, we consider $\mathcal{R}_0 \sim \text{Uniform}(20, 130)$ with $\mathcal{R} = 150$. The actual number of resources used by surviving species will not be uniformly distributed because specialists are more likely to have particularly "lucky" resource strategies; similarly, the typical number of resources used by species which produce a successful knock-out mutant will also be different (Fig. S9C). Rather than analytically calculating all of these distributions for this somewhat artificial example, we can use the observed distribution of $\mathcal{R}_0$ and $\mathcal{R}_P$ as empirical inputs to our analytic predictions of fitness effects and mutant-parent coexistence (in particular, plugging in the mean value of $\mathcal{R}_0$, and integrating over the probability distribution of $\mathcal{R}_P$). Doing so yields reasonable predictions for the mutant-parent coexistence probability and distribution of fitness effects (Fig. S9B,D), confirming that our results do not critically depend on each member of the community utilizing a similar number of resources.

4.3 Metabolic Trade-offs in Sampling

In our model, we have posited a soft metabolic trade-off where each sampled organism μ has an identically distributed "uptake budget" e^{X_μ} which must be divided between the resources it consumes; thus, organisms which consume more resources are generally less effective at consuming any individual one. Other work has considered different forms of such trade-offs (37, 44) or none at all, which can alter the number of species that coexist in the community (29). For example, Refs. (36) and (37) also consider a soft metabolic trade-off, but where the cost of consuming additional resources is associated with a larger average death rate rather than with taking a portion of a fixed overall budget. While these models describe very similar final communities, the main difference is that the death rate trade-off model predicts that as sampling depth $\mathcal{S}/\mathcal{S}^*$ increases, the fitness effect of using an additional resource becomes slightly deleterious on average, such that the mean invasion fitness of knock-out mutations is small but positive (Fig. S8A). Accounting for this small fitness effect through an effective change in pure fitness ΔX in knock-out and knock-in mutants brings the mutant-parent coexistence probabilities in agreement between the models (Fig. S8B).

Another alternative is to eliminate metabolic trade-offs in sampling entirely (29). In terms of the model we have discussed so far, a lack of trade-offs effectively grants organisms which use more resources a larger budget to divide among those resources. This change causes organisms which use more resources to be overrepresented in the population due to their higher effective fitness, and lowers the niche saturation $\mathcal{S}^*/\mathcal{R}$ of resultant communities. However, when these empirical changes in effective $\mathcal{R}_0$ and $\mathcal{S}^*/\mathcal{R}$ are accounted for, the distribution of fitness effects and mutant-parent coexistence probability are well-described by the original model with trade-offs (Fig. S8C-D). It is important to distinguish metabolic trade-offs in *sampled* species from trade-offs in mutations, which can be captured by changing the value of ΔX as described in the main text. These trade-offs may be quite different over short evolutionary timescales: for example, first-step mutations might introduce metabolic inefficiencies which are reduced by subsequent mutations.

4.4 Continuous Resource Usage

We have generally considered a binary model of resource use, where a natural parameter to describe metabolic overlap is the number of resources $\mathcal{R}_0$ used by a typical species. A simple continuous analogue to binary resource usage is drawing resource strategies from a Dirichlet distribution, $\vec{\alpha}_\mu \sim \mathrm{Dirichlet}(\mathcal{R}_0/\mathcal{R}, 1)$. Typical resource consumption vectors sampled from this distribution are qualitatively similar to binary ones; in particular, they will have $\sim \mathcal{R}_0$ entries which are $\mathcal{O}(\mathcal{R}_0^{-1})$, with the remaining entries much smaller. As an analogue to knock-out mutations, we chose a random entry of the parent organism's strategy vector with value at least $\gamma/\mathcal{R}_0$, and reduced it by $\gamma/\mathcal{R}_0$, where $\gamma = 1$ represents a full knock-out. We also considered "global-effect" mutations which slightly perturb each element of the resource use vector. Specifically, a global strategy mutation $\Delta\vec{\alpha}$ was constructed by multiplying each element of $\vec{\alpha}$ by an independent Gaussian random variable with mean 0 and variance $1/\mathcal{R}_0$, then by rescaling such that $\sum_i \Delta\alpha_i = 0$ and $|\Delta\vec{\alpha}| = \gamma/\mathcal{R}_0$.

Using resource consumption vectors drawn from a continuous distribution in this manner left the quantitative properties of our communities relatively unchanged (Fig. S3A-B, Fig. S5).

As a more extreme example of continuous resource usage, we can consider a community where all resources are significantly utilized by all species, but at slightly different levels. For example, we can draw $\vec{\alpha}_\mu \sim \mathrm{Dirichlet}(B, 1)$ for $B \gg 1$. In this case, the distribution of resource consumption strategies will approach a Gaussian, $\alpha_{\mu,i} \sim \mathcal{N}(\mathcal{R}^{-1}, \mathcal{R}^{-1}/\sqrt{B})$; similar sampling schemes have been analyzed before, as in (29). With these Gaussian resource consumption strategies, the number of resources $\mathcal{R}_0$ utilized by a typical species no longer has intuitive meaning. However, we can generalize our conclusions from binary resource usage by defining an effective $\mathcal{R}_0$ such that the binary model and the Gaussian model predict the same level of variance $\|\Delta\vec{\alpha}_{\mathrm{comm}}\|$ in resource consumption strategies:

$$\mathcal{R}_0^{\mathrm{eff}} \approx \mathcal{R} - \mathcal{R}/B. \tag{S130}$$

Thus, a model where organisms use all resources at slightly different levels is equivalent to a binary resource usage model where organisms use *almost* all resources, in terms of the metabolic overlap between species. We can define mutations in this framework by perturbing the parent's resource usage $\vec{\alpha}_P$ toward a different, i.i.d. resource usage profile $\vec{\alpha}'$:

$$\alpha_{M,i} = \mathcal{R}^{-1} + \sqrt{1-\Gamma}\left(\alpha_{P,i} - \mathcal{R}^{-1}\right) + \sqrt{\Gamma}\left(\alpha'_i - \mathcal{R}^{-1}\right), \tag{S131}$$

where $\Gamma \leq 1$ is a mutation effect size parameter describing how much of the initial resource use profile $\vec{\alpha}_P$ is retained.

What should the parent-mutant coexistence probability be for such a mutation? Plugging in the $\|\Delta\vec{\alpha}\|$ and $\|\Delta\vec{\alpha}_{\mathrm{comm}}\|$ for this assembly process into Eq. (9), we find

$$\mathbb{P}_{\mathrm{coex}} \approx \frac{\Gamma}{\mathcal{S}^*/\mathcal{R}} \tag{S132}$$

for $\Gamma \ll 1$. Indeed, we find that the mutant-parent coexistence probability for this community is well-predicted by naively using the binary resource predictions with the value of $\mathcal{R}_0$ defined in Eq. (S130), and has the expected qualitative dependence on Γ and B (Fig. S3C).

4.5 Specialist Community Assembly

As discussed in SI Section 3.5, the replica-theoretic approach involves a "large-ecosystem" approximation which is only valid when $\mathcal{S}^*\mathcal{R}_0 > \mathcal{R}$, such that each resource is consumed by many species. Species coexisting in such ecosystems tend to not have large differences in their uptake budgets X_μ, although these differences can still be ecologically relevant. Here, we consider a simple "specialist assembly" procedure, as a representative example that violates this condition. Imagine that an initial community contains $\mathcal{S}^*$ species, each of which use exactly one resource with no overlap. These species will trivially coexist, at abundances which match the supply of their private resource, $n_\mu = K_\mu$. Then, suppose

that one organism P acquires a partial knock-in mutation which enables it to use another resource ν. We label the mutant M, so that

$$\alpha_{P,\mu} = 1 \qquad \alpha_{P,\nu} = 0, \tag{S133}$$

$$\alpha_{M,\mu} = 1 - \gamma_{\text{spec}}, \qquad \alpha_{M,\nu} = \gamma_{\text{spec}}, \tag{S134}$$

where γ_{spec} describes the effect size of the mutation. While this ecosystem is clearly a caricature, it is a useful example far from the regime where our results should be applicable with which to compare our qualitative results.

What is the impact of the mutation? If no other organism utilizes resource ν, then the mutant will trivially always survive and coexist with its parent. However, if another organism O is already utilizing resource ν, then the mutant will survive if its fitness is greater, $X_M > X_O$. If the mutant survives, it will either replace its parent strain, or coexist with it while replacing O. Specifically, mutant-parent coexistence occurs if the following condition is met:

$$\mathbb{P}_{\text{coex}} = \mathbb{P}\left[\frac{K_\mu}{K_\mu + K_\nu} < \gamma_{\text{spec}}\right], \tag{S135}$$

as well as the trivial case where no other organism utilizes ν. Although the condition for coexistence is different from the more complex assembly procedure we discuss elsewhere, we still find that it is nonzero even in fully saturated communities (Fig. S3D). Thus, while our results do not directly extend to these communities, their existence does not invalidate our main conclusions. Furthermore, we expect that the sparser interaction networks of these communities makes them more amenable to previous analysis focusing on small numbers of species and resources (30).

5 Simulations and Numerics

5.1 Community Simulations

Throughout this work, we have performed numerical simulations of community assembly in order to verify and extend our analytical results. These simulations were performed by randomly sampling species and solving the optimization problem corresponding to ecological equilibrium (described in SI Section 1.5), rather than explicitly solving the abundance dynamics: the final outputs are the resource availabilities h_i and the surpluses Δ_μ. The code for these simulations is based on that in Ref. (37).

Much of our analysis relies on identifying whether individual species are present in the population or extinct. Mathematically, these cases correspond to whether Δ_μ is zero or negative. Numerically, however, Δ_μ is never precisely zero, necessitating the definition of an "extinction threshold." Any species with a surplus Δ_μ below the extinction threshold, a small negative number, is classified as extinct. We typically defined the extinction threshold as $-10^{-3} \cdot \text{Std}(h_i)/\mathcal{R}_0$, several orders of magnitude smaller than the typical fitness effect of a knock-in or knock-out mutation.

In Figure S10, we demonstrate that this is a reasonable choice of threshold. In order to

be numerically sensible, there should be a clear separation between typical values of Δ_μ for surviving and extinct species, and the threshold should be well within this gap. We verified this by plotting the number of putative surviving species as a function of the extinction threshold. In the neighborhood of the threshold we have chosen, the number of putative survivors is independent of the threshold, suggesting that there is a clear separation between alive and extinct species and that our choice of threshold should not heavily impact our results. We also analyzed the values Δ_μ for each strain before and after mutant invasion in order to ensure that we were reliably identifying species that go extinct. A scatterplot of initial vs. final Δ_μ values shows that almost all species are clearly separated into a "surviving" cluster and an "extinct" cluster, again suggesting that our results are not due to numerical artifacts.

5.2 Theory Numerics

While the main text and remainder of the SI describe the novel theory key to our work, some theory curves shown in figures involve minor post-processing steps or slightly more precise calculations than those quoted in the text. For convenience, we summarize these steps (which are also performed in our code supplement) below.

Distribution of fitness effects (e.g. Fig. 2C-D): Our theoretical results predict that the distribution of fitness effects (DFE) follows a Gaussian distribution with mean zero and standard deviation σ_{inv}, which is calculated by plugging our parameters into the equations in SI Section 3.2. The mean of this distribution is taken as a variable ΔX (sometimes assumed to be zero).

Coexistence probability (e.g. Fig. 3): We calculate $\mathbb{P}_{\text{coex}}$ by analytically integrating Eq. (S107) over the probability distributions of s_{inv} and $\mathcal{Z}_2$. Both distributions are truncated Gaussians multiplied by a linear term: $p(s) \sim s\rho(s)\theta(s)$, where $\rho(s)$ is the distribution of fitness effects above and $\theta(s)$ is the Heaviside function. $\mathcal{Z}_2$ is related to the parent relative abundance via Eq. (S101), and thus takes on a similar distribution: $p(z_2) \sim (z_2 - \lambda)e^{-z_2^2/2}\theta(z_2 - \lambda)$. In the inset of Fig. 3D, $\mathbb{P}_{\text{coex}}$ is conditioned on s_{inv}, so it is not integrated over. In Fig. S6, we condition on both s_{inv} and $\mathcal{Z}_2$ (via f_P), so Eq. (S107) determines a line in the s_{inv}-f_P plane above which mutant-parent coexistence is permitted.

Relative abundance (e.g. Fig. 4D): Relative abundance f_μ is calculated from simulation output by identifying surviving species using the resource surplus Δ_μ (as described in SI Section 5.1) and then using a non-negative least-squares solver to invert the relationship $h_i^{-1} = (\mathcal{R}/\kappa_i)\sum_\mu f_{\mu\in\text{Survivors}}\alpha_{\mu,i}e^{X_\mu}$.

Simluating model extensions (e.g. Fig. S3C): For some extensions of the model, the theoretical expectation of $\mathcal{S}^*/\mathcal{R}$ was not explicitly calculated. Rather than choosing parameters to target a specific value of $\mathcal{S}^*/\mathcal{R}$, simulations were run with a fixed value of $\text{Std}(X_\mu)$ and varying log-spaced $\mathcal{S}/\mathcal{R}$ to cover a range of $\mathcal{S}^*/\mathcal{R}$. To make smooth theoretical predictions from this data, the empirical values of $\lambda(\mathcal{S}^*/\mathcal{R})$ were linearly interpolated over the sampled data. In Fig. S8C-D, the sampled $\mathcal{R}_0$ differed from the $\mathcal{R}_0$ among survivors due to the

systematic advantage in using more resources; as such, the empirical average $\mathcal{R}_0$ among surviving species was used as input to the theory.

Figure 2B: We numerically solve the system of equations in SI Section 3.2 for λ and find $\mathcal{S}^*/\mathcal{R}$ by inverting Eq. (S88).
Figure 2E: We calculate the correlation between the expressions in Equations (S120) and (S121) in terms of σ_{inv}.
Figure S1: The set of equations in Eq. (S1a) were numerically solved using a standard MATLAB differential equation solver, with random initial conditions corresponding to low species abundances and resource concentrations at their steady-state in the absence of species.
Figure S8A-B: An effective $\Delta X = m/\mathcal{R}$ is determined by Eq. (S89) with $\chi = 0$.
Figure S9A: In the left panel, $\mathbb{P}_{\text{coex}}$ is determined by integrating Eq. (S129) over its random variables rather than Eq. (S107). $\mathcal{Z}_1$ is a standard Gaussian truncated to take positive values, while $\mathcal{Z}_2$ is a standard Gaussian truncated to take values above $\lambda\sqrt{V_{\text{tot}}/V_P}$.
Figure S9B-D: In addition to the changes discussed in "Simulating model extensions" above, the mean values of $\mathcal{R}_0$ and $\mathcal{R}_P$ among surviving species were linearly interpolated over the sampling regions of each condition to produce the theory curves in panel B. While the mean value of $\mathcal{R}_0$ was used for the DFE in panel D, the entire empirical distribution of $\mathcal{R}_P$ (i.e., the number of resources used by species which were randomly drawn to produce mutants across all simulation trials) was numerically integrated over to account for the tail effects.