

**A collective phase in resource competition in a highly
diverse ecosystem**

Supplemental material

Mikhail Tikhonov and Remi Monasson

S1. THE LYAPUNOV FUNCTION F

Recall that the dynamics of our model are given by

$$\frac{dn_\mu}{dt} = b_\mu n_\mu \Delta_\mu,$$

where Δ_μ is the resource surplus $\Delta_\mu = \sum_i \sigma_{\mu i} H_i(T_i) - \chi_\mu$. This section will show that this dynamics possess a Lyapunov function:

$$F(\{n_\mu\}) = \sum_i \hat{H}_i(T_i) - \sum_\mu n_\mu \chi_\mu. \quad (\text{S1})$$

In other words, F increases on any trajectory of the dynamics above. In addition, we will show that F is convex and bounded from above.

Proposition 1: F increases on any trajectory

We first note that the derivative of F with respect to a species' abundance n_μ is precisely the resource surplus Δ_μ :

$$\frac{\partial F}{\partial n_\mu} = \sum_i H_i(T_i) \frac{\partial T_i}{\partial n_\mu} - \chi_\mu = \Delta_\mu,$$

Therefore, F is indeed a Lyapunov function:

$$\frac{dF}{dt} = \sum_\mu \frac{\partial F}{\partial n_\mu} \frac{dn_\mu}{dt} = \sum_\mu b_\mu n_\mu \Delta_\mu^2 > 0.$$

Proposition 2: F is bounded from above

To see this, recall that $H_i(\cdot)$ was required to be a decreasing function of its argument; moreover, to forbid unbounded growth of any species, we required that for large enough demand T , the resource availability $H_i(T)$ should go to zero. It follows that its integral $\hat{H}_i(x) \equiv \int^x H(T) dT$ grows sub-linearly; in other words, for any $\lambda > 0$ we have $H_i(x) < \lambda x$ if x is large enough. We conclude that $F(\vec{n})$ goes to $-\infty$ as the norm of the abundance vector increases (this precisely corresponds to forbidding infinite population growth). A continuous function defined on the positive quadrant $\{n_\mu \geq 0\}$ and going to $-\infty$ at the boundary of this region is bounded from above, as claimed.

Proposition 3: F is convex

To see this, note that for any function $f(\vec{n})$, the following two operations leave its convexity invariant (M is an arbitrary matrix):

1. adding a linear function of its arguments: $f(\vec{n}) \mapsto g(\vec{n}) = f(\vec{n}) + M\vec{n}$;
2. performing a linear transformation of its arguments: $f(\vec{n}) \mapsto h(\vec{n}) = f(M\vec{n})$.

Given these observations, convexity of F directly follows from the convexity of $\hat{H}_i(x)$ (which is an integral of a decreasing function).

S2. LOCATING THE COMMUNITY EQUILIBRIUM: THE GEOMETRIC INTUITION

The main text shows that the equilibrium of community dynamics is always located at the boundary of the “unsustainable region” Ω defined in the text. Which boundary point is selected? Here we present an intuitive geometric argument, which will be formalized in the following section.

Let h_* be the resource availability at community equilibrium. For concreteness, consider the case $N = 2$, and assume the equilibrium state harbors two species $\{\vec{\sigma}_1, \chi_1\}$ and $\{\vec{\sigma}_2, \chi_2\}$, so that the point \vec{h}_* is the intersection of lines $\vec{h} \cdot \vec{\sigma}_1 = \chi_1$ and $\vec{h} \cdot \vec{\sigma}_2 = \chi_2$. Consider now the vector of total demand \vec{T}_* at this equilibrium. By definition, it is a linear combination of the two strategy vectors: $\vec{T}_* = n_1 \vec{\sigma}_1 + n_2 \vec{\sigma}_2$. Importantly, the coefficients here must be *positive*. We conclude that at equilibrium, the vector \vec{T}_* must point “strictly outward” relative to the region Ω , as in Fig. S1.

This property is sufficient to uniquely determine the equilibrium point. Indeed, consider a vector field $\vec{T}(\vec{h})$, where to each point of the resource availability space \vec{h}_0 we associate the vector of total demand \vec{T}_0 that corresponds to such resource depletion, i.e. such that $H(\vec{T}_0) = \vec{h}_0$. The intuitive argument above suggests that the equilibrium of community dynamics can be found by following this vector field. And indeed, this vector field is a gradient of a certain function, and therefore locating the equilibrium corresponds to maximizing this function. This is formally proven in the following section.

We stress that the vector field $\vec{T}(\vec{h})$ does not describe the dynamics itself; it is merely a tool to find the equilibrium point. The trajectories of the system in the harvest space are *not* integral lines of this vector field.




FIG. S1. The geometric intuition behind the selection of the equilibrium point: at equilibrium, the vector of total demand \vec{T}_* must be pointing “strictly outward” relative to the unsustainable region Ω . Here, at a two-species equilibrium of $\vec{\sigma}_1$ and $\vec{\sigma}_2$, \vec{T}_* must lie within the sector highlighted in blue. The equilibrium point can therefore be found by following the vector field $\vec{T}(\vec{h})$ (in gray).

S3. LOCATING THE COMMUNITY EQUILIBRIUM: THE FORMAL PROOF

Proposition 1: *There exists a function \tilde{F} defined on the harvest plane, such that its gradient at any point \vec{h} is the demand vector \vec{T} that corresponds to this resource availability*

vector:

$$\frac{\partial \tilde{F}}{\partial h_i} = T_i \quad \Leftrightarrow \quad H_i(T_i) = h_i.$$

Proof: Consider $F = \sum_i \hat{H}_i(T_i)$. This function has the property that $\frac{\partial F}{\partial T_i} = H_i(T_i)$. The function \tilde{F} can be explicitly constructed as the Legendre transform of F :

$$\tilde{F}(\vec{h}) = \left[\vec{h} \cdot \vec{T} - F(T) \right] \Big|_{\vec{T}=\vec{T}^*}$$

where \vec{T}^* is defined by the condition $H_i(T_i^*) = h_i$. It is easy to check that this function satisfies the desired requirement. Indeed, for each component i (and omitting this index for simplicity):

$$\frac{\partial \tilde{F}}{\partial h} = T^* + h \frac{\partial T^*}{\partial h} - \frac{\partial T^*}{\partial h} \frac{\partial \hat{H}}{\partial T} \Big|_{T^*} = T^*.$$

Consider now an equilibrium community \mathcal{C} with the total demand \vec{T}^* , and the resource availability vector \vec{h}_* . We already know that \vec{h}_* lies at the boundary of the unsustainable region Ω (see main text). To determine exactly which boundary point is selected, we make the following observation:

Proposition 2: For any other vector $\vec{h}_1 \in \Omega$, we have

$$(\vec{h}_1 - \vec{h}_*) \cdot \vec{T} \leq 0. \quad (\text{S2})$$

Proof: Since \vec{h}_* is the equilibrium state, we can write:

$$\forall \mu : \quad n_\mu \left[\vec{\sigma}_\mu \cdot \vec{h}_* - \chi_\mu \right] = 0.$$

As for \vec{h}_1 , it lies in the unsustainable region Ω , and therefore:

$$\forall \mu : \quad n_\mu \left[\vec{\sigma}_\mu \cdot \vec{h}_1 - \chi_\mu \right] \leq 0.$$

Subtracting the former from the latter, and summing over μ , we conclude:

$$\sum_{\mu} n_\mu \vec{\sigma}_\mu (\vec{h}_1 - \vec{h}_*) \leq 0 \quad \Rightarrow \quad \vec{T} \cdot (\vec{h}_1 - \vec{h}_*) \leq 0 \text{ as claimed. } \square$$

This means that the equilibrium point is such that the value of \tilde{F} cannot be further increased: any movement within the unsustainable region goes against the gradient field of \tilde{F} .

Corollary: The equilibrium community state corresponds to the maximum of \tilde{F} over the unsustainable region Ω .

It is worth contrasting our approach to other situations where community-level objective functions may appear, e.g. community-level flux balance analysis, or similar approaches. In certain contexts, an optimization-based framework is simply postulated, and serves as an exploratory tool to investigate the possible regimes of network performance: for instance, the total metabolic output of a consortium might be taken as a “plausible” global objective function for the community to optimize. Here, we stress that no community-level objectives are postulated; the fact that the ecological dynamics in this model take the form of a global optimization problem is a “lucky” consequence of explicitly specified dynamics of (purely “selfish”) individual species. This special feature makes MacArthur’s resource competition model an especially convenient starting point for investigating the consequences of high dimensionality in an ecological context.

S4. THE ALGEBRA OF “PASSING INTO THE HARVEST SPACE”

Above, we have shown that locating the equilibrium of our ecological dynamics is in fact a convex optimization problem in the P -dimensional space of species abundances. We then described how this optimization problem can be formulated directly in the N -dimensional space of “harvests” h_i . To build intuition, it is instructive to consider the following algebraic argument showing how the two optimization problems are mapped into each other.

In order to locate the maximum of the Lyapunov function F , we investigate the large- β limit of the partition function $Z(\beta)$:

$$\max F = \lim_{\beta \rightarrow \infty} \left(\frac{\log Z}{\beta} \right), \text{ where } Z(\beta) = \int_0^\infty e^{\beta F} \prod_{\mu} dn_{\mu}$$

To compute Z , we first introduce T_i as convenient auxiliary variables. This allows integrating over n_{μ} :

$$\begin{aligned} Z &= \int_0^\infty \prod_{\mu} dn_{\mu} \int_0^\infty dT_i \delta\left(T_i - \sum_{\mu} n_{\mu} \sigma_{\mu i}\right) e^{\beta F(n_{\mu})} \\ &= \int_0^\infty \prod_{\mu} dn_{\mu} \int_0^\infty dT_i \left[\int \frac{d\theta_i}{2\pi/\beta} e^{-i\beta\theta_i(T_i - \sum_{\mu} n_{\mu} \sigma_{\mu i})} \right] e^{\beta(\sum_i \hat{H}_i(T_i) - \sum_{\mu} n_{\mu} \chi_{\mu})} \\ &= \int_0^\infty dT_i \int \frac{d\theta_i}{2\pi/\beta} e^{\beta \sum_i [\hat{H}_i(T_i) - i\theta_i T_i]} \prod_{\mu} \int_0^\infty dn_{\mu} e^{-\beta n_{\mu} [\chi_{\mu} - i \sum_i \theta_i \sigma_{\mu i}]} \\ &= \int_0^\infty dT_i \int \frac{d\theta_i}{2\pi/\beta} e^{\beta \sum_i [\hat{H}_i(T_i) - i\theta_i T_i]} \prod_{\mu} \frac{1/\beta}{\chi_{\mu} - i \sum_i \theta_i \sigma_{\mu i}} \end{aligned}$$

We now focus on the integral over T_i . For large β , it can be computed using saddle-point method. Denoting $i\theta_i \equiv h_i$, we find that the saddle-point T_i^* is defined by the condition:

$$H_i(T_i^*) = h_i,$$

which justifies our suggestive notation (we recognize h_i as the substrate availability at equilibrium demand T_i^*). This condition implicitly defines T_i^* as a function of h_i , so all that remains is the N -dimensional integral over \vec{h} :

$$Z = \text{const} \times \int_{-i\infty}^{i\infty} d\vec{h} \exp\left[-\beta \tilde{F}(\vec{h})\right] \prod_{\mu} \frac{1/\beta}{\chi_{\mu} - \vec{h} \cdot \vec{\sigma}_{\mu}}. \quad (\text{S3})$$

The \tilde{F} in the exponent is precisely the Legendre transform of \hat{H}_i :

$$\tilde{F} \equiv \sum_i \tilde{F}_i = \sum_i \left[h_i T_i - \hat{H}_i(T_i) \right]_{\text{at } T_i = T_i^*}.$$

For large β , this is again a saddle-point integral. We are starting to recognize the problem of extremizing \tilde{F} ; however, here it is computed for purely imaginary arguments, and so a few more steps are needed. The integration contours cannot simply be rotated onto the real axes, since the integrand has a complicated pole structure. Instead, we can convert the integration

contours into piecewise-linear shapes, two of which are purely imaginary, and one is purely real: $-i\infty \rightarrow 0 \rightarrow x \rightarrow x + i\infty$, with $x \in \mathbb{R}$. The deformation of the integration contour is allowed only as long as the poles are not crossed, and the integrand has a pole whenever $\Delta_\mu = 0$ (we note that the denominator in (S4) is $(-\Delta_\mu)$, the negative resource surplus of species μ). Thus in our N -dimensional integral, the shifting of each contour will depend on the exact values of all other variables. Thankfully, the integrand can have an extremum only if all h_i are real, and whenever $N - 1$ variables h_i take real values, the remaining one can vary (on its real-valued portion of the contour) from 0 to the highest value that can be reached without crossing any of the hyperplanes $\Delta_\mu = 0$. The region delimited by these hyperplanes is precisely the “unsustainable region” Ω defined in the main text. We conclude that for the purposes of the saddle-point calculation, our integral becomes:

$$Z = \text{const} \times \int_{\Omega} d\vec{h} \exp \left[-\beta \tilde{F}(\vec{h}) \right] \prod_{\mu} \frac{1/\beta}{\chi_{\mu} - \vec{h} \cdot \vec{\sigma}_{\mu}}. \quad (\text{S4})$$

The certain lack of rigour in our description of the transition from Eq. S3 to Eq. S4 will not be a problem. The purpose of this section is to build additional intuition about the algebraic structure of the problem, and analyzing the expression (S4) will prove instructive. However, the following sections will only use the fact that community equilibrium maximizes \tilde{F} , a result that was rigorously obtained in section titled “Locating the community equilibrium: A formal proof”.

In expression (S4), the exponential term $e^{-\beta \tilde{F}}$ dominates the integrand everywhere, except in the immediate vicinity of the region boundary where $\frac{1}{\Delta_{\mu}}$ diverges. If β is large, but finite, the extremum is achieved at a point \vec{h}^* lying strictly inside the region Ω , at a distance of order $1/\beta$ from the nearest bounding hyperplanes. In this “finite temperature” regime, all species n_{μ} have non-zero abundance: since n_{μ} enters into Z as $e^{-\beta n_{\mu} |\Delta_{\mu}|}$, the observables n_{μ} follow an exponential distribution with mean $\langle n_{\mu} \rangle = \frac{1}{\beta |\Delta_{\mu}|}$. In the zero-temperature limit ($\beta = \infty$), this expected abundance vanishes for all species except a select few, for which Δ_{μ} is precisely zero. Thus, as described in the main text, the extremum \vec{h}^* reaches the boundary of Ω . At this value of harvests, a finite set of species have resource surplus of precisely zero, corresponding to finite-abundance survivors. The resource surplus of all other species is negative, and they go extinct at equilibrium.

To make this argument more precise, we note that at large, but finite β , Eq. (S4) gives us

$$\log Z = \max_{\vec{h} \in \Omega} \left\{ -\beta \tilde{F}(\vec{h}) - \sum_{\mu} \log |\Delta_{\mu}| \right\}.$$

For a large β , the sum over μ is dominated by only a few terms, those corresponding to the closest hyperplanes for which Δ_{μ} tends to zero. Denote their set \mathbb{S} (for “survivors”). The extremum condition:

$$\frac{\partial \tilde{F}}{\partial h_i} = - \sum_{\mu \in \mathbb{S}} \frac{1}{\beta |\Delta_{\mu}|} \frac{\partial |\Delta_{\mu}|}{\partial h_i} = \sum_{\mu \in \mathbb{S}} n_{\mu} \sigma_{\mu i} = T_i. \quad (\text{S5})$$

This of course makes perfect sense given the definition of \tilde{F}_i as the Legendre transform of \hat{H}_i .

But if we are only interested in the identity of the species that survive at community equilibrium, it is wholly encoded in the location of the extremum \vec{h}_* at $\beta = \infty$. Since this

\vec{h}_* is located at the boundary of Ω , the shape of the repulsive potential of interaction with the hyperplanes $\Delta_\mu = 0$ is irrelevant in this limit, and can be replaced by the Heaviside theta-function $\theta(-\Delta_\mu)$. In this limit, the problem reduces to computing the extremum of \tilde{F} over the unsustainable region, as stated in the main text.

S5. RESOURCE SUPPLY MODELS

S5.1. The model of MacArthur

Different models of resource supply correspond to different expressions of the function \tilde{F} . The renewable resource of MacArthur, with renewal rate r and maximum resource availability K , is described by the following resource depletion rule (originally derived in Ref. [14]; see also the Supplementary section A in Ref. [16]):

$$H(T) = K \left(1 - \frac{T}{r}\right) \quad \Rightarrow \quad T^* = r \left(1 - \frac{h^*}{K}\right)$$

Integrating $H(T)$, we find $\hat{H}(T) = KT - \frac{KT^2}{2r}$, and therefore

$$hT^* - \hat{H}(T^*) = hr - \frac{r}{2K}h^2 + \text{const}$$

We conclude that for the resource model of MacArthur:

$$\tilde{F}(\{h_i\}) = \sum_i \left[r_i h_i - \frac{r_i}{2K_i} h_i^2 \right] + \text{const}$$

S5.2. A constantly supplied resource

The constant-supply model of Ref. [16] is a simpler model that postulates that a fixed amount of resource R is evenly divided among all competitors: $H(T) = R/T$. In this model, we have $\hat{H}(T) = R \log T$ and $T^* = R/h^*$. Consequently:

$$\tilde{F} = \sum_i h_i T_i^* - \hat{H}_i(T_i^*) = \sum_i R_i \log h_i + \text{const}$$

S5.3. A general model

Consider the close-to-symmetric scenario, where the supply of all resources is similar. If the cost of all strategies is close to χ_0 , then the availability of resources at equilibrium will be close to χ_0 as well. Linearizing around this point, a general resource supply model $H(T)$ can be characterized with two parameters. First, let τ be the value of demand at which resource availability hits χ_0 : by definition, $H(\tau) = \chi_0$. In the vicinity of this point, let γ be the ‘‘elasticity’’ of supply, describing how quickly resource is depleted by a small increase in demand $\delta\tau \ll \tau$:

$$H(\tau + \delta\tau) = \chi_0 - \gamma \delta\tau.$$

Let us compute the Legendre transform of \hat{H} in this model. We have:

$$H(T) = \chi_0 - \gamma(T - \tau) \quad \Rightarrow \quad \hat{H}(T) = T(\chi_0 + \gamma\tau) - \gamma\frac{T^2}{2}$$

The demand that corresponds to a particular value of resource availability close to χ_0 :

$$H(T^*) = h \quad \Rightarrow \quad T^* = \tau + \frac{1}{\gamma}(\chi_0 - h).$$

In the vicinity of $h = \chi_0$ it is convenient to work with shifted variables: $h \equiv \chi_0 - \frac{g}{N}$ (the N in the denominator reminds that the deviation is small). After a little algebra we find:

$$\tilde{F}_i(g_i) = -\frac{1}{2}\gamma\tau_i^2 - \tau_i\frac{g_i}{N} - \frac{1}{2\gamma_i}\left(\frac{g_i}{N}\right)^2,$$

where the index i reminds us that parameters τ_i and γ_i could be different for different resources. Omitting the irrelevant global constant, we find the expression for \tilde{F} in this general cost model:

$$\tilde{F}_{\text{general}} = -\sum_i \left[\tau_i\frac{g_i}{N} + \frac{1}{2\gamma_i}\left(\frac{g_i}{N}\right)^2 \right]. \quad (\text{S6})$$

Of course, the two particular models we considered above reduce to this same form in the vicinity of $h_i \approx \chi_0$. Specifically, for MacArthur's model of renewable resource:

$$\tilde{F}_{\text{MacArthur}}(\{g_i\}) = -\sum_i \left[r_i \left(1 - \frac{\chi_0}{K_i} \right) \frac{g_i}{N} + \frac{r_i}{2K_i} \left(\frac{g_i}{N} \right)^2 \right].$$

Similarly, for the constant supply model:

$$\tilde{F}_{\text{const}}(\{g_i\}) = -\sum_i R_i \left[\frac{g_i}{N} + \frac{1}{2} \left(\frac{g_i}{N} \right)^2 \right].$$

Our calculation below will be for the general close-to-symmetric case where the supply of resources is similar:

$$\tau_i \equiv \bar{\tau} + \frac{\delta\tau_i}{\sqrt{N}} \quad \gamma_i \equiv \bar{\gamma} + \frac{\delta\gamma_i}{\sqrt{N}}. \quad (\text{S7})$$

Here $\sum_i \delta\tau_i = \sum_i \delta\gamma_i = 0$ by definition, and as N becomes large, $\delta\tau_i$ and $\delta\gamma_i$ remain of order 1. Note that the function \tilde{F} can be rescaled by a constant positive factor, leaving the maximization problem unchanged (we seek the location of the maximum, not its magnitude). Without restricting generality, therefore, we can set $\bar{\tau} = 1$.

S6. THE COST MODEL

S6.1. Only low-cost outliers matter

The main text made the argument that since competition is restricted to only a subset of species, the self-selected pool of low-cost outliers with similar costs χ_μ , the details of the cost model matter only inasmuch as they determine the properties of this subset. To illustrate

this point, consider a scenario at $N = 2$ where the cost of strategy $\{x, 1 - x\}$ is drawn out of a normal distribution with mean $\chi_0(1 + \sin \pi x)$ and width $\frac{1}{2}\chi_0 \sin \pi x$. One such realization for 20 equally spaced values of $0 \leq x \leq 1$ is shown in Fig. S2. In this illustration, mixed strategies tend to be expensive; as a result, all three low-cost outliers (in bold) are close to specialists, and only these species are competitive (this is exactly the scenario depicted in Fig. 2b). However, the exact details of the cost model (the precise shape of the solid red curve in Fig. S2) are otherwise irrelevant for the coexistence problem.




FIG. S2. **The cost model: an illustration.** Competition is restricted to a self-selected pool of low-cost outliers, so we only need to model this pool, where all strategies have similar costs. This illustration shows an example of a cost model at $N = 2$ yielding the competition scenario depicted in Fig. 2b of the main text (the cost of strategy $\{x, 1 - x\}$ (at $N = 2$) is plotted as a function of x). Since metabolic diversification is penalized, all three competitive species (highlighted) are close to being specialists, but otherwise the details of the red curve have no effect on the coexistence problem.

S6.2. The assumption of uncorrelated costs

The main text made a strong assumption, namely that the cost of a species is uncorrelated with its metabolic strategy. The aim of this section is to show that this assumption is not overly restrictive.

Let us imagine a more realistic cost model. Intuitively, some functional capabilities are naturally more compatible than others. If two metabolic enzymes require starkly different conditions for optimal function, or if their enzymatic activity is characterized by an undesirable cross-talk, expressing them in the same organism would require maintaining two separate compartments and incur extra cost. For instance, the nitrogen-fixing enzyme nitrogenase is inactivated by oxygen; as a result, performing oxygen respiration and nitrogen fixation in the same cell would be extremely costly and is never observed. Conversely, an organism can efficiently make use of several enzymes requiring similar specific conditions (e.g. specific levels of pH) by investing into the maintenance of appropriate intracellular conditions only once, but harvesting the benefit from all.

These considerations motivate the following model for the cost of an organism μ with metabolic strategy $\sigma_{\mu i}$:

$$\chi_{\mu} = c_0 + c_1 \sum_i \sigma_{\mu i} + \sum_{ij} J_{ij} \sigma_{\mu i} \sigma_{\mu j}. \quad (\text{S8})$$

This cost is composed of a baseline c_0 (e.g. maintaining the cell wall, ribosomes, replication machinery), a fixed cost per pathway c_1 , and a correction that arises from (positive or

negative) interactions between pathways, specified by an $N \times N$ symmetric matrix J_{ij} , which summarizes the complexity of enzymatic biochemistry. The scatter of J_{ij} causes the species' cost to explicitly depend on its pathway content. A non-zero mean of J_{ij} , and the parameter c_0 , both introduce a global trend, causing the cost of a typical species to depend on its degree of specialization (the number of carried pathways). The main text assumed a model where neither generalists nor specialists are given any preference; the more structured model above breaks this degeneracy.

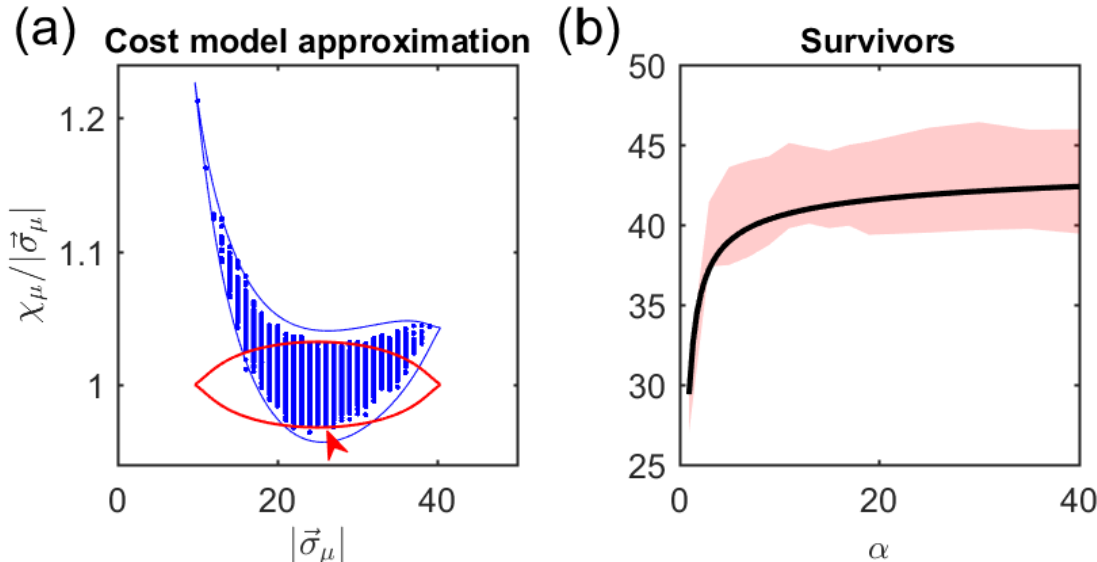


FIG. S3. **A random-cost model provides a good approximation even if the costs are structured.** (a) Blue points: The cost per carried pathway (characterizes individual performance of a species [16]), shown for 5×10^4 randomly generated species at $N = 50$. Cost model parameters: $c_0 = 10$, $c_1 = 1$, and the elements of J_{ij} were drawn out of a Gaussian distribution with mean 0.015 and width 0.02. The specific values are not important, and are used purely for the sake of concreteness. Blue lines: the analytic prediction for the envelope of this point cloud. The surviving species are found among the low-cost outliers (arrowhead). The approximation of uncorrelated costs substitutes the true cost model (blue lines) by a model that matches cost statistics in the vicinity of the low-cost region (red lines). (b) The number of survivors at equilibrium for the correlated cost model of panel (a), shown as a function of the species pool size α . Numerical simulations (mean \pm standard deviation over 50 replicates, shaded) overlaid with the theoretical prediction computed under the random-cost approximation (black line).

The model (S8) is fully deterministic with parameters $\{c_0, c_1, J_{ij}\}$, and these $\simeq N^2$ parameters determine the costs of a very large number of metabolic strategies (up to 2^N). Fig. S3a shows an example of this deterministic model for one particular set of parameters. For each of 5×10^4 randomly generated species, the plot shows its *cost per carried pathway* as a function of the number of carried pathways $|\vec{\sigma}_\mu| \equiv \sum_i \sigma_{\mu i}$. The former quantity is the best predictor of the likelihood of a species to survive competition (high-cost species are out-competed [16]).

We see that high specialization (small $|\vec{\sigma}_\mu|$) is costly; this is due to the c_0 term in Eq. S8. In the opposite extreme, the wide-spectrum generalists (large $|\vec{\sigma}_\mu|$) are negatively affected by pathway cross-talk which is, on average, detrimental. The parameters above were chosen to

be representative of the regime where the lowest-cost species are found at an intermediate value of $|\vec{\sigma}_\mu|$, in this case $|\vec{\sigma}_\mu| \approx N/2$. Under the model (S8), the costs of species are certainly correlated with the metabolic strategy: the latter determines the former. However, the number of surviving species is only of order N , and these are all selected from the low-cost outliers, a region indicated by arrowhead in Fig. S3a. For this restricted set (compared to 2^N), the assumption of effectively uncorrelated costs proves to be a very good approximation, and can be captured by just three parameters used in the main text:

- p : the sparsity of pathway content, chosen to match the average pathway count in the surviving species;
- ϵ : the width of the cost distribution, chosen to match the cost statistics of the low-cost outliers;
- α : the effective number of species in the pool, chosen to match the probability of generating a strategy with pN pathways; for simplicity, the parameters in Fig. S3a were chosen to set $p \approx 1/2$, in which case the number of species in the pool need not be renormalized.

If the costs of the same 5×10^4 species as shown in Fig. S3a were generated using the uncorrelated cost model with the best matching parameters, the datapoints would fill out the region indicated with red solid lines. Although this is a poor approximation elsewhere, in the only region that matters (arrowhead) the quality of the 3-parameter fit proves sufficient. Indeed, Fig. S3b shows the number of survivors at equilibrium, as a function of an increasing species pool size α . Numerical simulations (50 replicates, shaded) are overlaid with the theoretical prediction computed under the random-cost model used in the text (black line). For small magnitude of the effective “noise” ϵ , once the number of survivors hits N , the community enters the S -phase described in the main text.

Summarizing the above, the assumption of uncorrelated costs is not overly restrictive, because even in a more complex cost model, it can be invoked as an effective characterization of a small number of low-lying “energy levels”. This argument, presented here at an intuitive level, has been formally justified in other contexts, such as the integer partitioning problem [23, 24]. Of course, it is possible to make the structure in the species costs sufficiently strong to render the uncorrelated approximation invalid. For example, for certain parameter values the convexity of the “banana” shape in Fig. S3a can be inverted, so that extreme specialists and extreme generalists both acquire an advantage (there is no longer one optimal number of pathways, and instead the low-cost outliers fall into two categories). This regime cannot be adequately described by our uncorrelated model that assumes a single value for the metabolic “sparsity” p .

S7. THE REPLICA-THEORETIC CALCULATION

This section demonstrates how the geometrical problem formulated above can be solved using methods of statistical physics, specifically an approach termed “replica theory”. An attempt is made to present this computation in a detailed and self-contained way, i.e. not assuming familiarity with statistical physics of disordered systems. For a more comprehensive introduction to this powerful technique, we refer the reader to Ref. [25].

S7.1. The basic idea

We seek to compute:

$$Z = \int_0^\infty \prod_i dh_i e^{\beta \tilde{F}} \prod_{\mu=1}^P \theta(\chi_\mu - \vec{h} \cdot \vec{\sigma}_\mu) \quad (\text{S9})$$

The pool of competitors is modeled as follows. For each species μ , we pick $\vec{\sigma}_\mu$ as a random binary vector, where each component $\sigma_{\mu i}$ is 1 with probability p , and 0 otherwise. We then draw a random cost $\chi_\mu = \sum_i \sigma_{\mu i} + \epsilon x_\mu$, where x_μ is a Gaussian random variable of variance 1.

The argument of the Heaviside θ -function in Eq. (S9) is the negative resource surplus $-\Delta_\mu$. Under the cost model described above, we have

$$\Delta_\mu = \vec{h} \cdot \vec{\sigma}_\mu - \chi_\mu = -\epsilon x_\mu - \sum_i \sigma_{\mu i} (1 - h_i).$$

Change variables $h_i \equiv 1 - \frac{g_i}{N}$ (g_i runs from $-\infty$ to N), and introduce Δ_μ as an explicit auxiliary integration variable:

$$\begin{aligned} Z &= \int_{-\infty}^N \prod_i \frac{dg_i}{N} e^{\beta \tilde{F}(\{g_i\})} \prod_{\mu=1}^P \int d\Delta_\mu \theta(-\Delta_\mu) \delta\left(\Delta_\mu + \epsilon x_\mu + \frac{1}{N} \sum_i g_i \sigma_{\mu i}\right) \\ &= \int_{-\infty}^N \prod_i \frac{dg_i}{N} e^{\beta \tilde{F}(\{g_i\})} \prod_{\mu=1}^P \int \frac{d\Delta_\mu d\hat{\Delta}_\mu}{2\pi} \theta(-\Delta_\mu) \exp\left[i \sum_\mu \hat{\Delta}_\mu \left(\Delta_\mu + \epsilon x_\mu + \frac{1}{N} \sum_i g_i \sigma_{\mu i}\right)\right] \end{aligned}$$

In this expression, x_μ and $\sigma_{\mu i}$ are “frozen disorder”: they are drawn randomly, but are then kept fixed, while other variables relax to their equilibrium values. Computing this integral for a particular realization of the disorder, even if it were possible to do so, would not be very informative. Instead, we are interested in the behavior of the “typical” realization of the system. This means that we are interested in the typical free energy $\langle F \rangle = \langle \log Z \rangle$ (angular brackets denote averaging over disorder). This quantity is hard to compute directly, because the average is outside of the logarithm. The opposite case, the logarithm of the average, would be very simple to compute; unfortunately, unlike the free energy F , the partition function Z is dominated not by typical realizations of the disorder, but by extreme ones. The logarithm of the average would capture the behavior of the system in highly improbable extreme cases, which is of no use to us.

The gist of the “replica trick” is summarized in the following formula:

$$\langle \log Z \rangle = \lim_{n \rightarrow 0} \frac{\langle Z^n \rangle - 1}{n} \quad (\text{S10})$$

This trick makes it possible to formally derive an expression for $\langle \log Z \rangle$ by computing only expressions of type Z^n , which is the partition function of n copies (“replicas”) of the system, and then formally sending n to zero. The replicas are identical (have the same disorder), but independent (each has its own set of degrees of freedom). There are, of course, mathematical subtleties related to taking this limit, and we refer the reader to Ref. [25]. This reference also provide some intuition for the physical basis of the argument and the interpretation of the auxiliary variables that appear along the way. Here, the analytical result we derive will be validated by an excellent agreement with numerical simulations.

S7.2. Averaging over disorder

Proceeding with our argument, we write the partition function of n copies of our system. Each degree of freedom is copied n times, labeled by the replica index a running from 1 to n :

$$Z^n = \int_{-\infty}^{\infty} \prod_{i,a} \frac{dg_i^a}{N} e^{\beta \sum_a \tilde{F}(\{g_i^a\})} \prod_{\mu,a} \int \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a) \exp \left[i \sum_{\mu,a} \hat{\Delta}_\mu^a \left(\Delta_\mu^a + \epsilon x_\mu + \frac{1}{N} \sum_i g_i^a \sigma_{\mu i} \right) \right]$$

We stress that all n replicas have the same disorder: quantities x_μ and $\sigma_{\mu i}$ have no index a . Thanks to our strategic choice of cost model, averaging Z^n over these disorder variables is separable:

$$\begin{aligned} \langle Z^n \rangle_{x_\mu, \bar{\sigma}_\mu} &= \int \prod_{i,a} \frac{dg_i^a}{N} e^{\beta \sum_a \tilde{F}(g_i^a)} \prod_{\mu,a} \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a) e^{i \sum_{\mu,a} \hat{\Delta}_\mu^a \Delta_\mu^a} \\ &\quad \times \underbrace{\prod_{\mu} \langle e^{i\epsilon \sum_a \hat{\Delta}_\mu^a x_\mu} \rangle_{x_\mu}}_{(1)} \times \underbrace{\prod_{i,\mu} \langle e^{\frac{i}{N} \sum_a \hat{\Delta}_\mu^a g_i^a \sigma_{\mu i}} \rangle_{\sigma_{\mu i}}}_{(2)} \end{aligned}$$

If x is a Gaussian random variable of unit variance, then $\langle e^{\alpha x} \rangle_x = e^{\frac{1}{2}\alpha^2}$, and therefore

$$(1) \equiv \prod_{\mu} \langle e^{i\epsilon \sum_a \hat{\Delta}_\mu^a x_\mu} \rangle_{x_\mu} = \exp \left[-\frac{1}{2} \epsilon^2 \sum_{\mu} \left(\sum_a \hat{\Delta}_\mu^a \right)^2 \right].$$

To compute the second term, recall that $\sigma_{\mu i}$ is either 1 or 0 with probabilities p and $1-p$:

$$\begin{aligned} (2) &= \prod_{i,\mu} \langle e^{\frac{i}{N} \sum_a \hat{\Delta}_\mu^a g_i^a \sigma_{\mu i}} \rangle_{\sigma_{\mu i}} \\ &= \prod_{i,\mu} \left((1-p) + p e^{\frac{i}{N} \sum_a \hat{\Delta}_\mu^a g_i^a} \right) \\ &= \prod_{i,\mu} \left(1 + p \left[\frac{i}{N} \sum_a \hat{\Delta}_\mu^a g_i^a \right] + \frac{p}{2} \left[\frac{i}{N} \sum_a \hat{\Delta}_\mu^a g_i^a \right]^2 + o(1/N^2) \right) \\ &= \exp \left[\frac{ip}{N} \sum_{i,\mu,a} \hat{\Delta}_\mu^a g_i^a - \frac{p(1-p)}{2N^2} \sum_{i,\mu} \left(\sum_a \hat{\Delta}_\mu^a g_i^a \right)^2 + o(1/N^2) \right] \end{aligned}$$

Here we used $1 + p\epsilon + \frac{p\epsilon^2}{2} + \dots = \exp(p\epsilon + p(1-p)\epsilon^2/2 + \dots)$. Putting everything together:

$$\begin{aligned} \langle Z^n \rangle_{\text{disorder}} &= \int \prod_{i,a} \frac{dg_i^a}{N} e^{\beta \sum_a \tilde{F}(\{g_i^a\})} \prod_{\mu,a} \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \theta(-\Delta_\mu^a) \\ &\quad \times \exp \left\{ i \sum_{\mu,a} \hat{\Delta}_\mu^a \left(\Delta_\mu^a + \frac{p}{N} \sum_i g_i^a \right) - \frac{1}{2} \epsilon^2 \sum_{\mu} \left(\sum_a \hat{\Delta}_\mu^a \right)^2 - \frac{p(1-p)}{2N^2} \sum_{i,\mu} \left(\sum_a \hat{\Delta}_\mu^a g_i^a \right)^2 \right\} \end{aligned}$$

S7.3. Decoupling indices i and μ

In order to make progress, we need to eliminate terms that directly couple indices μ and i . To achieve this, we introduce a yet another set of variables, using the same trick of inserting a delta-function into our integral:

$$\begin{aligned} \text{Introduce } m^a &\equiv \frac{1}{N} \sum_i g_i^a &\Rightarrow \text{insert } 1 &= \int \prod_a \frac{dm^a d\hat{m}^a}{2\pi} e^{i\hat{m}^a(m^a - \frac{1}{N} \sum_i g_i^a)} \\ \text{Introduce } q^{ab} &\equiv \frac{1}{N^2} \sum_i g_i^a g_i^b &\Rightarrow \text{insert } 1 &= \int \prod_{a \leq b} \frac{dq^{ab} d\hat{q}^{ab}}{2\pi} e^{i\hat{q}^{ab}(q^{ab} - \frac{1}{N^2} \sum_i g_i^a g_i^b)} \end{aligned}$$

These auxiliary variables will become our order parameters, as they capture the mean and variance of the deviation of resource availability from 1. The indices i and μ are now decoupled, and we can split the integral accordingly. Recall that we have index i labels N different resources, while μ labels P different strategies. Recall also that \tilde{F} is a sum over N terms: $\tilde{F} \equiv \sum_{i=1}^N \tilde{F}_i$.

$$\begin{aligned} \langle Z^n \rangle &= \int \prod_{a \leq b} \frac{dq^{ab} d\hat{q}^{ab}}{2\pi} \int \prod_a \frac{dm^a d\hat{m}^a}{2\pi} \exp \left[i \sum_{a \leq b} q^{ab} \hat{q}^{ab} + i \sum_a \hat{m}^a m^a \right] \\ &\times \prod_i \left\{ \int_{-\infty}^N \prod_a \frac{dg_i^a}{N} \exp \left[\sum_a \beta \tilde{F}_i(\{g_i^a\}) - \frac{i}{N} \sum_a \hat{m}^a g_i^a - \frac{i}{N^2} \sum_{a \leq b} \hat{q}^{ab} g_i^a g_i^b \right] \right\} \\ &\times \prod_\mu \left\{ \int \prod_a \frac{d\Delta_\mu^a d\hat{\Delta}_\mu^a}{2\pi} \prod_a \theta(-\Delta_\mu^a) \exp \left[i \sum_a \hat{\Delta}_\mu^a (\Delta_\mu^a + pm^a) - \frac{1}{2} \sum_{a,b} (p(1-p)q^{ab} + \epsilon^2) \hat{\Delta}_\mu^a \hat{\Delta}_\mu^b \right] \right\} \\ &= \int \prod_{a \leq b} \frac{dq^{ab} d\hat{q}^{ab}}{2\pi} \int \prod_a \frac{dm^a d\hat{m}^a}{2\pi} \exp \left[i \sum_{a \leq b} q^{ab} \hat{q}^{ab} + i \sum_a \hat{m}^a m^a \right] \times \prod_{i=1}^N A_i \times B^P, \end{aligned}$$

with A_i and B given by:

$$\begin{aligned} A_i &= \int_{-\infty}^N \prod_a \frac{dg^a}{N} \exp \left[\sum_a \beta \tilde{F}_i(g^a) - \frac{i}{N} \sum_a \hat{m}^a g^a - \frac{i}{N^2} \sum_{a \leq b} \hat{q}^{ab} g^a g^b \right] \\ B &= \int \prod_a \frac{d\Delta^a d\hat{\Delta}^a}{2\pi} \prod_a \theta(-\Delta^a) \exp \left[i \sum_a \hat{\Delta}^a (\Delta^a + pm^a) - \frac{1}{2} \sum_{a,b} (p(1-p)q^{ab} + \epsilon^2) \hat{\Delta}^a \hat{\Delta}^b \right] \end{aligned}$$

After averaging over disorder, the problem becomes fully symmetric in indices μ (all strategies are drawn from the same distribution, so there is no inherent difference in how they contribute). The same would be true for resources, except in the interest of generality, we allowed the supply functions H_i to be different for different resources.

S7.4. Decoupling replicas: the replica-symmetric ansatz

The idea now is to treat the integrals over m and q as saddle-point integrals. We will shortly introduce a rescaling of variables that will make β appear in the exponent to serve as

the large parameter, making the saddle-point approximation appropriate. In this approximation, the integral is replaced by the value of the integrand at one location, the saddle point (up to a multiplicative prefactor, which, as explained below, is irrelevant for our purposes). When looking for the saddle point, we will make the assumption that it is symmetric under a permutation of replicas. This is the so-called “replica-symmetric ansatz”. The validity of this assumption will be justified *a posteriori* by the fact that the saddle-point we will find is “well-behaved” and the analytical results match the numerical simulations.

At a fully replica-symmetric saddle point, all components of m^a must coincide. As for the matrix q^{ab} , all its diagonal elements must be equal, and all of the off-diagonal ones must be equal as well. The same holds for the conjugate variables \hat{m}^a , \hat{q}^{ab} . We therefore look for a saddle point of the following form:

$$\hat{q}^{ab} = \begin{cases} \hat{q}_D & \text{if } a = b \\ \hat{q}_O & \text{if } a \neq b \end{cases}, \quad \hat{m}^a = \hat{m}^*,$$

and similarly for q^{ab} and m^a . With these assumptions:

$$\log \langle Z^n \rangle = \text{extr} \left\{ i n q_D \hat{q}_D + i \frac{n(n-1)}{2} q_O \hat{q}_O + i n \hat{m}^* m^* + \sum_i \log A_i + P \log B \right\}.$$

(Why did we take the logarithm? Recall from Eq. S10 that our ultimate goal is to compute $Z^n - 1$ in the $n \rightarrow 0$ limit. Conveniently, in this limit, subtracting 1 is the same as taking the logarithm.) We stress that all expressions need only be computed to the leading exponential order. In particular, multiplicative constants are irrelevant as they only amount to an additive constant under the logarithm. Below, such constants will be omitted, and the “equal” signs will mean “up to a constant additive factor”.

S7.5. The limit $n \rightarrow 0$

Note that in the expression we just found, n enters in a way that allows taking the formal limit $n \rightarrow 0$, which is our final goal. It turns out that this limit also makes the expressions for $\log A_i$ and $\log B$ somewhat easier to compute. Therefore, we write:

$$\lim_{n \rightarrow 0} \frac{\log \langle Z^n \rangle}{n} = \lim_{n \rightarrow 0} \text{extr} \left\{ i q_D \hat{q}_D - \frac{i}{2} q_O \hat{q}_O + i \hat{m}^* m^* + \frac{1}{n} \sum_i \log A_i + \frac{P}{n} \log B \right\}.$$

For reasons that will become clear later, we note that this can also be written as

$$\lim_{n \rightarrow 0} \frac{\log \langle Z^n \rangle}{n} = \lim_{n \rightarrow 0} \text{extr} \left\{ i \left(\hat{q}_D - \frac{1}{2} \hat{q}_O \right) q_D - \frac{q_D - q_O}{2} (-i \hat{q}_O) + i \hat{m}^* m^* + \frac{1}{n} \sum_i \log A_i + \frac{P}{n} \log B \right\} \quad (\text{S11})$$

S7.6. Computing $\log A_i$

Recall the expression we denoted A_i :

$$A_i = \int_{-\infty}^N \prod_a \frac{dg^a}{N} \exp \left[\sum_a \beta \tilde{F}_i(g^a) - \frac{i}{N} \sum_a \hat{m}^a g^a - \frac{i}{N^2} \sum_{a \leq b} \hat{q}^{ab} g^a g^b \right]$$

We first write:

$$\begin{aligned}
\sum_{a \leq b} \hat{q}^{ab} g^a g^b &= \hat{q}_D \sum_a (g_a)^2 + \frac{1}{2} \hat{q}_O \sum_{a \neq b} g^a g^b \\
&= \hat{q}_D \sum_a (g_a)^2 + \frac{1}{2} \hat{q}_O \left[\left(\sum_a g_a \right)^2 - \sum_a (g_a)^2 \right] \\
&= \left(\hat{q}_D - \frac{1}{2} \hat{q}_O \right) \sum_a (g_a)^2 + \frac{1}{2} \hat{q}_O \left(\sum_a g_a \right)^2
\end{aligned}$$

Now use Feynman's trick of removing the square by introducing an extra Gaussian variable:

$$\exp\left(\frac{1}{2} C x^2\right) = \int \mathcal{D}z e^{z x \sqrt{C}}$$

(the curly \mathcal{D} denotes the standard Gaussian measure with variance 1). This lets us write:

$$\exp\left[-\frac{i}{N^2} \sum_{a \leq b} \hat{q}^{ab} g^a g^b\right] = \int \mathcal{D}z \exp\left[z \frac{\sqrt{-i\hat{q}_O}}{N} \sum_a g^a - \frac{i}{N^2} \left(\hat{q}_D - \frac{1}{2} \hat{q}_O\right) \sum_a (g^a)^2\right].$$

At the price of introducing an extra Gaussian variable, all replicas are now fully decoupled. Plugging this into our expression for A_i :

$$A_i = \int \mathcal{D}z \left[\int_{-\infty}^N \frac{dg}{N} \exp\left(\beta \tilde{F}_i(g) - \frac{i}{N^2} \left(\hat{q}_D - \frac{1}{2} \hat{q}_O\right) g^2 - \frac{1}{N} \left(i\hat{m}^* - z\sqrt{-i\hat{q}_O}\right) g\right) \right]^n$$

Conveniently, for small n :

$$\log \int \mathcal{D}z x^n = \log \left[\int \mathcal{D}z (1 + n \log x + \dots) \right] = \log \left[1 + n \int \mathcal{D}z \log x + \dots \right] = n \int \mathcal{D}z \log x + \dots$$

Therefore:

$$\lim_{n \rightarrow 0} \frac{\log A_i}{n} = \int \mathcal{D}z \log \int_{-\infty}^N \frac{dg}{N} \exp\left[\beta \tilde{F}_i(g) - \frac{i}{N^2} \left(\hat{q}_D - \frac{1}{2} \hat{q}_O\right) g^2 - \frac{1}{N} \left(i\hat{m}^* - z\sqrt{-i\hat{q}_O}\right) g\right]$$

Introduce rescaled variables as follows:

$$\begin{aligned}
i \left(\hat{q}_D - \frac{1}{2} \hat{q}_O \right) &\equiv \beta a \\
\sqrt{-i\hat{q}_O} &\equiv \frac{\beta b}{\sqrt{N}} \\
i\hat{m}^* &\equiv \beta \hat{m}
\end{aligned}$$

In the new variables:

$$\lim_{n \rightarrow 0} \frac{\log A_i}{n} = \int \mathcal{D}z \log \int_{-\infty}^N \frac{dg}{N} \exp \beta \left[\tilde{F}_i(g) - a \frac{g^2}{N^2} - \left(\hat{m} - \frac{zb}{\sqrt{N}} \right) \frac{g}{N} \right]$$

Substitute the general form of \tilde{F}_i from (S6), for a close-to-symmetric resource supply (S7):

$$\lim_{n \rightarrow 0} \frac{\log A_i}{n} = \int \mathcal{D}z \log \int_{-\infty}^N \frac{dg}{N} \exp \beta \left[- \left(a + \frac{1}{2\gamma_i} \right) \left(\frac{g}{N} \right)^2 + \left(-\bar{\tau} - \hat{m} + \frac{zb + \delta\tau_i}{\sqrt{N}} \right) \frac{g}{N} \right]$$

In the limit $\beta \rightarrow \infty$:

$$\lim_{n \rightarrow 0} \frac{\log A_i}{n} = \beta \int \mathcal{D}z \max_{y < 1} \left[- \left(a + \frac{1}{2\gamma_i} \right) y^2 + \left(-\bar{\tau} - \hat{m} + \frac{zb + \delta\tau_i}{\sqrt{N}} \right) y \right]$$

Here $y \equiv \frac{g}{N}$. The quadratic form $-Py^2 + Qy$ is maximized at $y = \frac{Q}{2P}$, reaching the maximal value of $\frac{Q^2}{4P}$. In our case, therefore, for a given z the quadratic form reaches its maximum at

$$y_i^*(z) = \frac{-\bar{\tau} - \hat{m} + (zb + \delta\tau_i)/\sqrt{N}}{2a + 1/\gamma_i}.$$

Let us shift the variable \hat{m} by defining $\hat{m} \equiv -\bar{\tau} - \frac{\delta\hat{m}}{\sqrt{N}}$. For now, we can treat this as a simple change of variables. We find:

$$y_i^*(z) = \frac{1}{\sqrt{N}} \frac{zb + \delta\tau_i + \delta\hat{m}}{2a + 1/\gamma_i} = \frac{1}{\sqrt{N}} \frac{zb + \delta\tau_i + \delta\hat{m}}{2a + 1/\bar{\gamma}} + o\left(\frac{1}{\sqrt{N}}\right).$$

Recall that \hat{m} is one of the variables over which the extremum is computed in (S11). The shift of \hat{m} that we just did constitutes an assumption, namely that the extremum is located close to $\hat{m} \approx \bar{\tau}$ (i.e. that the difference is at most of order $1/\sqrt{N}$). We will check the consistency of this assumption below. Note that if we are correct to assume this, then $y_i^*(z)$ is small, justifying the expansion of resource depletion functions to first order in y .

The integral over z is now simple to compute:

$$\lim_{n \rightarrow 0} \frac{\log A_i}{n} = \frac{\beta}{2N} \int \mathcal{D}z \frac{(zb + \delta\tau_i + \delta\hat{m})^2}{2a + 1/\bar{\gamma}} = \frac{\beta}{2N} \frac{b^2 + (\delta\tau_i + \delta\hat{m})^2}{2a + 1/\bar{\gamma}}$$

Finally, performing the sum over i and recalling that $\sum_i \delta\tau_i = 0$, we find a very simple final expression:

$$\sum_i \lim_{n \rightarrow 0} \frac{\log A_i}{n} = \beta \frac{b^2 + \delta\hat{m}^2 + \overline{\delta\tau^2}}{4a + 2/\bar{\gamma}},$$

where $\overline{\delta\tau^2} \equiv \frac{1}{N} \sum_i (\delta\tau_i)^2$ is the variance of resource supply across i .

S7.7. Computing $\log B$

Recall the definition of B :

$$B = \int \prod_a \frac{d\Delta^a d\hat{\Delta}^a}{2\pi} \prod_a \theta(-\Delta^a) \exp \left[i \sum_a \hat{\Delta}^a (\Delta^a + pm^a) - \frac{1}{2} \sum_{a,b} (p(1-p)q^{ab} + \epsilon^2) \hat{\Delta}^a \hat{\Delta}^b \right]$$

Proceeding as above, we decompose

$$- \sum_{a,b} \frac{p(1-p)q^{ab} + \epsilon^2}{2} \hat{\Delta}^a \hat{\Delta}^b = - \frac{p(1-p)}{2} (q_D - q_O) \sum_a (\hat{\Delta}^a)^2 - \frac{p(1-p)q_O + \epsilon^2}{2} \left(\sum_a \hat{\Delta}^a \right)^2.$$

We then remove the square by introducing an extra Gaussian variable, making all replicas fully decoupled:

$$B = \int \mathcal{D}w \left[\int \frac{d\Delta d\hat{\Delta}}{2\pi} \theta(-\Delta) \exp \left(i\hat{\Delta}(\Delta + pm^*) - \frac{p(1-p)}{2}(q_D - q_O)\hat{\Delta}^2 + iw\hat{\Delta}\sqrt{p(1-p)q_O + \epsilon^2} \right) \right]^n.$$

Note that the integral over $\hat{\Delta}$ inside the square brackets is a simple Gaussian integral, and we can write:

$$B = \int \mathcal{D}w \left[\int_{-\infty}^0 \frac{d\Delta}{\sqrt{2\pi p(1-p)(q_D - q_O)}} \exp \left\{ -\frac{1}{2} \frac{\left(\Delta + pm^* + w\sqrt{p(1-p)q_O + \epsilon^2} \right)^2}{p(1-p)(q_D - q_O)} \right\} \right]^n \quad (\text{S12})$$

Introduce a notation $E(x) \equiv \int_x^\infty \frac{dy}{\sqrt{2\pi}} e^{-y^2/2}$. This is essentially the error function, up to a couple constants that would be a nuisance to carry around: $E(x) = \frac{1}{2} \operatorname{erfc} \frac{x}{\sqrt{2}}$. We can then write:

$$\lim_{n \rightarrow 0} \frac{\log B}{n} = \int \mathcal{D}w \log E \left[-\frac{pm^* + w\sqrt{p(1-p)q_O + \epsilon^2}}{\sqrt{p(1-p)(q_D - q_O)}} \right].$$

As above for A , we now introduce a rescaled variable x , and some convenient notations:

$$\begin{aligned} q_D - q_O &\equiv \frac{Nx}{\beta} \\ q_D &\approx q_O \equiv q \\ \sqrt{p(1-p)q + \epsilon^2} &\equiv \psi \\ pm^*/\psi &\equiv \lambda \end{aligned}$$

In the new variables:

$$\lim_{n \rightarrow 0} \frac{\log B}{n} = \int \mathcal{D}w \log E \left[-\sqrt{\frac{\beta}{N}} \frac{(w + \lambda)\psi}{\sqrt{p(1-p)x}} \right].$$

The logarithm of $E(x) \equiv \int_x^\infty \frac{dy}{\sqrt{2\pi}} e^{-y^2/2}$ in the large-argument limit is very simple. Indeed:

$$\lim_{\beta \rightarrow \infty} E(\sqrt{\beta}x) \simeq \begin{cases} 1 - C \exp(-\beta x^2/2) & \text{if } x < 0 \\ C \exp(-\beta x^2/2) & \text{if } x > 0 \end{cases}$$

Therefore (omitting additive constants as always):

$$\lim_{\beta \rightarrow \infty} \log E(\sqrt{\beta}x) \simeq \begin{cases} 0 & \text{if } x < 0 \\ -\beta x^2/2 & \text{if } x > 0 \end{cases}$$

Plugging this into our expression for $\log B$, we find

$$\lim_{n \rightarrow 0} \frac{\log B}{n} = -\frac{\beta\psi^2}{2Np(1-p)x} \int_{-\infty}^{-\lambda} \mathcal{D}w (w + \lambda)^2 = -\frac{\beta\psi^2}{2Np(1-p)x} I(\lambda),$$

where $I(\lambda)$ can be expressed in terms of the error function:

$$I(\lambda) \equiv \int_0^\infty e^{-\frac{(w-\lambda)^2}{2}} w^2 \frac{dw}{\sqrt{2\pi}} = -\frac{\lambda}{\sqrt{2\pi}} e^{-\frac{\lambda^2}{2}} + \frac{1 + \lambda^2}{2} \operatorname{erfc} \left(\frac{\lambda}{\sqrt{2}} \right).$$

S7.8. Putting everything together

Combining the results above, plugging them into (S11), and recalling that in the large- N limit, P also goes to infinity with $\frac{P}{N} \equiv \alpha$, we find:

$$\begin{aligned} \langle \log Z \rangle &= \lim_{n \rightarrow 0} \frac{\langle Z^n - 1 \rangle}{n} = \lim_{n \rightarrow 0} \frac{\log \langle Z^n \rangle}{n} \\ &= \lim_{n \rightarrow 0} \text{extr} \left\{ i \left(\hat{q}_D - \frac{1}{2} \hat{q}_O \right) q_D - \frac{q_D - q_O}{2} (-i \hat{q}_O) + i \hat{m}^* m^* + \frac{1}{n} \sum_i \log A_i + \frac{P}{n} \log B \right\} \\ &= \beta \text{extr} \left\{ aq - \frac{b^2 x}{2} + \left(-\bar{\tau} - \frac{\delta \hat{m}}{N} \right) \frac{\psi \lambda}{p} + \frac{b^2 + \delta \hat{m}^2 + \overline{\delta \tau^2}}{4a + 2/\bar{\gamma}} - \frac{\alpha \psi^2}{2p(1-p)x} I(\lambda) \right\}. \end{aligned}$$

Recall that $\psi \equiv \sqrt{p(1-p)q + \epsilon^2}$, so that the extremum is taken over six variables: $\delta \hat{m}$, a , b , q , λ and x . Consider the extremum condition for $\delta \hat{m}$:

$$\frac{\delta \hat{m}}{2a + 1/\bar{\gamma}} = \frac{1}{N} \frac{\psi \lambda}{p}.$$

As $N \rightarrow \infty$, we therefore have $\delta \hat{m} \rightarrow 0$, demonstrating that the approximation $\hat{m} \approx \bar{\tau}$ was indeed self-consistent. Setting $\delta \hat{m} = 0$, we find:

$$\langle \log Z \rangle = \beta \text{extr} \left\{ aq - \frac{b^2 x}{2} - \bar{\tau} \frac{\psi \lambda}{p} + \frac{b^2 + \overline{\delta \tau^2}}{4a + 2/\bar{\gamma}} - \frac{\alpha \psi^2}{2p(1-p)x} I(\lambda) \right\}.$$

The powers of N and β in our rescaled variables were chosen to ensure that this expression no longer depends on N and is proportional to β . Conveniently, the extremum conditions for variables a and b can also be solved, and these variables eliminated:

$$\begin{aligned} a &= \frac{\bar{\gamma} - x}{2x\bar{\gamma}} \\ b^2 &= \frac{q}{x^2} - \overline{\delta \tau^2} \end{aligned}$$

Our final expression for the partition function:

$$\langle \log Z \rangle = \beta \text{extr} \left\{ \frac{\bar{\gamma} - x}{2x\bar{\gamma}} q + \frac{\overline{\delta \tau^2}}{2} x - \frac{\lambda \tau \psi(q)}{p} - \frac{\alpha \psi^2(q)}{2xp(1-p)} I(\lambda) \right\} \quad (\text{S13})$$

The extremum is to be computed over q , x and λ . Here α , p and ϵ (hidden in $\psi \equiv \sqrt{p(1-p)q + \epsilon^2}$) are parameters characterizing the pool of competitors (number of strategies, typical functional sparsity, and intrinsic scatter cost, respectively). Parameters $\bar{\tau}$, $\overline{\delta \tau^2}$ and $\bar{\gamma}$ characterize resource supply (respectively: average capacity, variability across resources, and average ‘‘elasticity’’). For the simplest single-parameter resource model used in the text, the influx of resource i is fixed at R_i , whose average, without loss of generality, can be set to 1. In this case we have $\bar{\tau} = 1$, $\bar{\gamma} = 1$ and the remaining parameter is the variance of resource supply $\overline{\delta \tau^2}$, denoted $\overline{\delta R^2}$ in the main text.

S8. THE SADDLE-POINT EQUATIONS

S8.1. Simplifying the equations to solve them numerically

Hidden in $I(\lambda)$ is the error function erfc , which means that the extremum of (S13) cannot be found analytically. However, the equations can be simplified to a form where they can either be solved numerically, or investigated analytically in certain limits.

The extremum conditions:

$$\frac{\partial}{\partial q} : \quad \frac{1 - \alpha I(\lambda)}{x} = \frac{1}{\bar{\gamma}} + \frac{(1-p)\lambda\bar{\tau}}{\psi} \quad (\text{S14})$$

$$\frac{\partial}{\partial x} : \quad \bar{\delta\tau^2} - \frac{q}{x^2} + \frac{\alpha\psi^2}{x^2 p(1-p)} I(\lambda) = 0 \quad (\text{S15})$$

$$\frac{\partial}{\partial \lambda} : \quad -2 - \frac{\alpha\psi}{x\bar{\tau}(1-p)} \frac{dI}{d\lambda} = 0$$

Using the first and the third equations, we write:

$$\frac{1 - \alpha I(\lambda)}{\frac{1}{\bar{\gamma}} + \frac{(1-p)\lambda\bar{\tau}}{\psi}} = x = -\frac{\alpha\psi}{2\bar{\tau}(1-p)} \frac{dI}{d\lambda}$$

Rearranging, we find a way to express ψ (and thus q) in terms of λ only:

$$\frac{\psi}{\bar{\gamma}\bar{\tau}} = \frac{2(1-p)(1 - \alpha I(\lambda))}{-\alpha I'(\lambda)} - (1-p)\lambda \quad (\text{S16})$$

Plugging this into the first equation, and recalling the definition of $I(\lambda)$, we find a very simple expression for x :

$$\frac{x}{\bar{\gamma}} = 1 - \alpha I(\lambda) + \frac{\alpha\lambda}{2} \frac{dI}{d\lambda} = 1 - \alpha E(\lambda). \quad (\text{S17})$$

(Recall that $E(\lambda) \equiv \frac{1}{2} \text{erfc}(\lambda/\sqrt{2})$.) This result makes it possible to eliminate x from the equations. Plugging all this into (S15), we find an equation that involves λ only, and can easily be solved numerically. Once λ is known, equations (S16) and (S17) determine q and x .

S8.2. Investigating the limit $\epsilon \rightarrow 0$

To study the equations analytically in the limit $\epsilon \rightarrow 0$, we plug the expression for x into (S14), and reorganize the terms in equation (S15), putting it into the form cited in the main text:

$$\begin{cases} \frac{1 - \alpha I(\lambda)}{1 - \alpha E(\lambda)} = 1 + \frac{\lambda}{\psi} (1-p)\bar{\gamma} \\ \psi^2(1 - \alpha I(\lambda)) = \epsilon^2 + (1 - \alpha E(\lambda))^2 \bar{\gamma}^2 p(1-p) \bar{\delta\tau^2} \end{cases}$$

The easiest way to derive the expression for the critical line given in the main text is to observe the following. At $\epsilon = 0$, we see that simultaneously setting $\psi = 0$ and $1 - \alpha E(\lambda) = 0$

yields a solution. One can check that in the vicinity of the transition, both go to zero linearly, so that their ratio remains well-defined. Omitting the negligible first terms in the right-hand sides of both equations, we rewrite them as follows:

$$\begin{cases} \frac{\psi}{\bar{\gamma}} \frac{1 - \alpha I(\lambda)}{1 - \alpha E(\lambda)} = \lambda (1 - p) \\ \frac{\psi^2}{\bar{\gamma}^2} \frac{(1 - \alpha I(\lambda))^2}{(1 - \alpha E(\lambda))^2} = p(1 - p) \overline{\delta\tau^2} (1 - \alpha I(\lambda)) \end{cases}$$

At the critical line, the system is degenerate, i.e. the two equations are proportional. We immediately read off the condition that must be satisfied:

$$p(1 - p) \overline{\delta\tau^2}_{\text{crit}} (1 - \alpha_{\text{crit}} I(\lambda)) = \lambda^2 (1 - p)^2 \quad \Rightarrow \quad \overline{\delta\tau^2}_{\text{crit}} = \frac{1 - p}{p} \frac{\lambda^2}{1 - \alpha_{\text{crit}} I(\lambda)}.$$

Here λ is fixed by the condition $1 - \alpha_{\text{crit}} E(\lambda) = 0$. This is the expression quoted in the main text.

S9. COMPUTING THE NUMBER OF SURVIVORS AT EQUILIBRIUM

To find the number of species that survive at equilibrium in our model, we set out to compute the distribution of the observable Δ . Recall that Δ is the ‘‘resource surplus’’; for a given set of competitors, there is a discrete set of values of Δ_μ , the resource surplus experienced by each species. However, after we average our partition function over disorder, Δ becomes a random variable, drawn out of a certain distribution, whose shape is encoded in the partition function. This distribution is what we now set out to compute.

The Δ -dependent part of the partition function is fully contained in the expression for B . Recall the intermediate expression (S12) derived earlier, when computing B :

$$\begin{aligned} B &= \int \mathcal{D}w \left[\int_{-\infty}^0 \frac{d\Delta}{\sqrt{2\pi p(1-p)Nx/\beta}} \exp \left\{ -\frac{\beta}{2} \frac{(\Delta + pm + w\psi)^2}{p(1-p)Nx} \right\} \right]^n \\ &\equiv \int \mathcal{D}w \left[\int_{-\infty}^0 \frac{d\Delta}{\sqrt{2\pi/\beta'}} \exp \left\{ -\frac{\beta'}{2} (\Delta + pm + w\psi)^2 \right\} \right]^n, \end{aligned}$$

Here we introduced $\beta' \equiv \frac{\beta}{p(1-p)Nx}$ to make the notations slightly less heavy. Let us put this expression in the following form, retaining one copy of the integral over Δ , while evaluating the remaining $n - 1$ copies as before:

$$B = \int_{-\infty}^0 \frac{d\Delta}{\sqrt{2\pi/\beta'}} \int \mathcal{D}w \exp \left(-\frac{\beta'}{2} (\Delta + pm + w\psi)^2 \right) \left[E \left(-\sqrt{\beta'} (w\psi + pm) \right) \right]^{n-1}$$

Once again, $E(x)$ is a short-hand for $\frac{1}{2} \text{erfc}(x/\sqrt{2})$. Sending $n \rightarrow 0$ and recalling the notation $\lambda \equiv \frac{pm}{\psi}$:

$$B = \int_{-\infty}^0 \frac{d\Delta}{\sqrt{2\pi/\beta'}} \int \mathcal{D}w \exp \left\{ -\frac{\beta'}{2} (\Delta + (w + \lambda)\psi)^2 \right\} \frac{1}{E \left(-\sqrt{\beta'} \psi (w + \lambda) \right)}.$$

From this we infer the distribution of Δ for $\Delta \leq 0$ (positive Δ are forbidden):

$$p(\Delta) = \int \mathcal{D}w \frac{1}{\sqrt{2\pi/\beta'}} \frac{\exp\left\{-\frac{\beta'}{2}(\Delta + (w + \lambda)\psi)^2\right\}}{E(-\sqrt{\beta'}\psi(w + \lambda))}$$

This is a complicated-looking expression, but its $\beta' \rightarrow \infty$ limit can in fact be computed very easily, using the following trick. By definition of the function $E(x)$, this expression can be rewritten as follows:

$$p(\Delta) = \int \mathcal{D}w \left[\frac{\frac{1}{\sqrt{2\pi/\beta'}} \exp\left\{-\frac{\beta'}{2}(\Delta + (w + \lambda)\psi)^2\right\}}{\int_{-\infty}^0 \frac{dy}{\sqrt{2\pi/\beta'}} \exp\left\{-\frac{\beta'}{2}(y + (w + \lambda)\psi)^2\right\}} \right] \equiv \int \mathcal{D}w p(\Delta|w).$$

The key observation is that here, the ‘‘conditional distribution’’ $p(\Delta|w)$ is (by inspection) a properly normalized distribution, for all w :

$$\forall w : \int_0^{\infty} d\Delta p(\Delta|w) = 1 \quad (\text{S18})$$

We now compute the limit $\lim_{\beta' \rightarrow \infty} p(\Delta|w)$. For $w > -\lambda$, this is a Gaussian distribution of width $\frac{1}{\sqrt{\beta'}} \rightarrow 0$ centered at $\Delta = -\psi(\lambda + w) < 0$. For $w < -\lambda$, the probability density is highest at $\Delta = 0$ and goes to zero everywhere else (recall that positive Δ are forbidden). Therefore, the normalization condition (S18) immediately tells us that:

$$\lim_{\beta' \rightarrow \infty} p(\Delta|w) = \begin{cases} \delta(\Delta + \psi(\lambda + w)) & \text{if } w > -\lambda \\ \delta(\Delta) & \text{if } w < -\lambda \end{cases}$$

As a result:

$$\begin{aligned} p(\Delta) &= \int_{-\lambda}^{\infty} \frac{dw}{\sqrt{2\pi}} e^{-\frac{w^2}{2}} \delta(\Delta + \psi(\lambda + w)) + \int_{-\infty}^{-\lambda} \frac{dw}{\sqrt{2\pi}} e^{-\frac{w^2}{2}} \delta(\Delta) \\ &= \frac{1}{\sqrt{2\pi\psi^2}} e^{-\frac{(\Delta + \lambda\psi)^2}{2\psi^2}} \cdot \theta(-\Delta) + E(\lambda)\delta(\Delta). \end{aligned} \quad (\text{S19})$$

This is the expression quoted in the main text. The weight of the delta-shaped peak at $\Delta = 0$ corresponds to the species whose resource balance is met; these are the species who survive competition.

S10. VALIDATION OF ANALYTICAL RESULTS AGAINST SIMULATIONS

Numerical simulations were performed in MatLab; a script generating all figures is available as Supplementary File 1. Random instances were generated by randomly drawing strategy vectors and costs as described in the main text. All simulations used $N = 50$ for the number of resources. The equilibrium of a community was determined by a direct N -dimensional numerical optimization of $\tilde{F} = \sum_i R_i/T_i$ subject to the linear constraints defining the ‘‘unsustainable region’’ Ω . The uneven resource supply was always implemented as a bimodal distribution shown in Fig. 3b (upper panel); the amplitude of the step was adjusted to match the required magnitude of δR^2 .

S10.1. Figure 2b

500 simulations were performed at $\alpha = 10$, $\epsilon = 10^{-3}$ and $\overline{\delta R^2} = 1$ (reusing the dataset computed for Fig. 3b). These parameters are comfortably in the S -phase, so the expected number of survivors is N . After numerical equilibration, all but the top N values of the resource surplus Δ were recorded (the top N are within numerical error of 0 and correspond to survivors). Panel 2b shows a joint histogram of these values recorded over 500 instances, appropriately normalized to be comparable with the theoretical distribution. For visualization purposes, the delta-shaped peak is shown as a rectangle, its height is meaningless. The theoretical prediction for the number of survivors is verified in panel 2c.

S10.2. Figure 2c

The theoretical curves are overlaid with simulations obtained as follows. For 10 values of α equispaced between 2 and 5, and for three values of ϵ indicated on the plot (0.1, 0.03 and 10^{-4}), communities were equilibrated using MatLab solver `fmincon` with numerical precision parameter 10^{-10} . Species whose resource surplus was within 10^{-8} of zeros were declared as survivors (for comparison, the resource surplus of the first extinct species was typically $\approx 10^{-6}$). The panel shows the mean number of survivors over 500 instances for each (α, ϵ) pair.

S10.3. Figure 3a

For each indicated epsilon, 10 simulations were performed for each α in the list: $\{1, 2, 3, 4, 5, 6, 8, 10, 15, 20, 50\}$.

S10.4. Figure 3b

For $\epsilon = 10^{-3}$ and $\overline{\delta R^2} = 1$, we performed 500 simulations at $\alpha = 2$ for V phase, and at $\alpha = 10$ for S -phase. The latter dataset was also used for generating the resource surplus histogram shown in 2b.

S11. INTERPRETATION OF THE V-PHASE AND THE S-PHASE

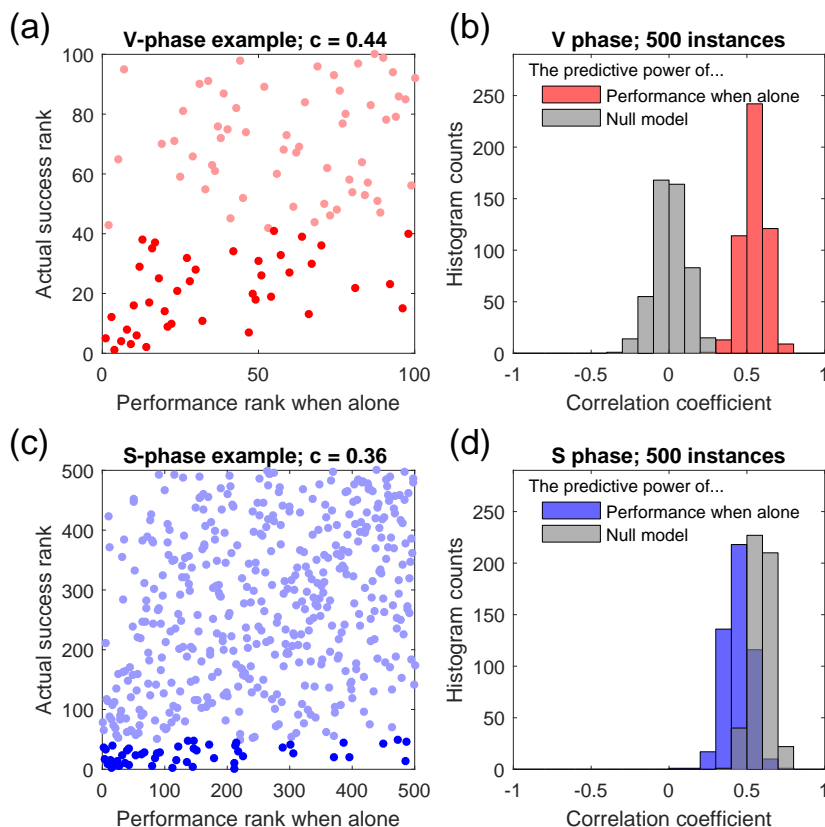


FIG. S4. Define a species “success rank” by ordering all survivors by decreasing abundance at equilibrium, followed by all the species that went extinct, in order of increasing resource insufficiency $|\Delta_\mu|$. Now, consider two quantities that could potentially be used to predict this success. To obtain the first, we measure the growth rate of each species when it is placed into the externally supplied conditions, with no competitors present. Alternatively, the “null model” performance predictor is simply the species’ cost per pathway: clearly, the high-cost species are less likely to survive. We call the latter quantity the “null model”, because it includes no information about the environment to which the community is subjected. We now ask if either of the two quantities are in fact predictive of the true success of a species. We measure predictive power by the Spearman (rank-order) correlation coefficient between the predictor being tested and the true success rank. In the V-phase, the species’ performance measured in external conditions is indeed predictive of the success rank (panel (a)) A histogram of the correlation coefficients over 500 instances (panel (b)) confirms that this predictor significantly outperforms the null model. In contrast, in the S-phase, the environment-specific performance measured in the externally imposed conditions becomes irrelevant (panels (c) and (d)). Although the correlation observed in panel C is non-zero, this is due to the generic fact that low-cost organisms are generally more likely to survive than high-cost ones. This is demonstrated in panel D: unlike the more intuitive V-phase, in the S-phase the “environment-aware” predictor performs *worse* than the null model.