Emergent neutrality in consumer-resource dynamics

Rafael D'Andrea[1,*], Theo Gibbs[1, 2], James P. O'Dwyer[1]

1 Department of Plant Biology, University of Illinois, Urbana IL USA
2 Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton NJ USA

* rdandrea@illinois.edu

Abstract

Neutral theory assumes all species and individuals in a community are ecologically equivalent. This controversial hypothesis has been tested across many taxonomic groups and environmental contexts, and successfully predicts species abundance distributions across multiple high-diversity communities. However, it has been critiqued for its failure to predict a broader range of community properties, particularly regarding community dynamics from generational to geological timescales. Moreover, it is unclear whether neutrality can ever be a true description of a community given the ubiquity of interspecific differences, which presumably lead to ecological inequivalences. Here we derive analytical predictions for when and why non-neutral communities of consumers and resources may present neutral-like outcomes, which we verify using numerical simulations. Our results, which span both static and dynamical community properties, demonstrate the limitations of summarizing distributions to detect non-neutrality, and provide a potential explanation for the successes of neutral theory as a description of macroecological pattern.

Author Summary

The neutral theory of biodiversity assumes that species are ecologically equivalent. Given the natural history observation of ubiquitous phenotypic differences between species, it is surprising that neutral theory has successfully predicted a broad range of biodiversity patterns, and simultaneously unsurprising that these results have not convinced ecologists that the natural world is neutral. However, we have lacked a description of how neutrality can emerge in a natural way from ecological mechanisms and species differences. Our study sheds light on this question, providing a theoretical backdrop for the success of neutral theory as a description of macroecological pattern. We derive a prediction for the degree to which consumers must differ in preferences for different resources before the resulting biodiversity patterns become distinguishable from neutrality. These predictions, which we confirm using simulations, show that neutral-like outcomes are possible even when resource requirements across consumers are very far from neutral. Our results can be tested in experimental microbial communities, where, equipped with an inferred consumption network, our analysis can yield predictions for biodiversity patterns and community turnover at different taxonomic levels.

Introduction

One of the central questions in community ecology is how species interactions affect community structure and ecological dynamics, and conversely whether summarized representations of the latter can be used to make inferences about the former. Macroecological patterns such as species abundance distributions in particular have been extensively studied, and often used in inference approaches [1–3].

However, these inference approaches have been called into question, especially following intense debate over the merits of the neutral theory of biodiversity [1, 4–10]. Neutral theory assumes that all species and individuals in a community are ecologically equivalent, and therefore community assembly is a purely stochastic process, with differences in species abundances resulting from demographic noise [11, 12]. This hypothesis has been heavily tested across many taxonomic groups and environmental contexts [13–16], and despite its radical assumptions, has had some success in describing macroecological patterns. Prominent among these successes is the prediction of a broad range of species abundances, matching the functional form for the species abundance distribution observed in multiple high-diversity communities [12, 16–18]. On the other hand, the theory has been criticized on two fronts: first, it is unclear whether neutral theory could successfully predict a broader range of community properties, particularly regarding community dynamics from generational to geological timescales [19–27]. Second, it has been argued that the patterns that neutral theory successfully fits do not uniquely reflect the underlying ecological properties of the species involved; in particular, it is possible that summarized indices of community structure and dynamics may look neutral even when species are not ecologically equivalent [2, 27–31].

Some previous studies of neutral outcomes in non-neutral dynamics have focused on phenomenological models of species interactions, such as Lotka-Volterra competition [32, 33]. One drawback of this approach is that the influence of demographic noise relative to ecological selection can be tuned as a free parameter, rather than resulting endogenously from the interactions between species and the resources they compete for. A study of consumer-resource dynamics has shown that neutral behavior may result from metabolic tradeoffs guaranteeing ecological equivalence between consumers [34]. However, consumers vastly outnumber resources in those ecosystems, and therefore coexistence—and the accompanying neutrality—requires those tradeoffs to be exact. It remains an open question whether species competing for resources may appear neutral when their resource requirements are not fine tuned, and we still lack a quantitative prediction for how large niche differences between consumers must be in order for those differences to be reflected at the community scale.

Here we present a model of stochastic, non-neutral consumer-resource dynamics, with which we investigate the potential for neutral-like outcomes from two perspectives: snapshots of community structure, and dynamics over multiple timescales. Using a deterministic version of our model, we derive predictions for a threshold between neutral and non-neutral behavior, which we test against numerical simulations of the full model. Specifically, we hypothesize that when the timescales of relaxation to equilibrium under non-neutral dynamics are commensurate with the timescales of drift to extinction under neutral dynamics, the system will appear neutral. We investigate two qualitatively different departures from neutrality to test this hypothesis: a generalist scenario, where consumers have different preferences for resources but without a marked preferences for any one resource over others; and a specialist scenario, where each consumer has its own preferred resource, while consuming all other resources at a lower rate. Our analysis is based on these two particular ways to 'break' neutrality, but is not confined to them. However, we demonstrate that these cases support our hypotheses: consumer resource models with a degree of non-neutrality anywhere below our predicted threshold display neutral-like dynamics and patterns.

Methods

Consumer-Resource Model

We base our analysis on a stochastic model of S consumers, whose abundances we denote by N_i for $i \in \{1, ..., S\}$, that compete for K abiotic, substitutable resources, whose concentrations are R_k for $k \in \{1, ..., K\}$. Resources are externally supplied, and depleted only by consumption. Consumers grow purely through the consumption of these resources, and also undergo density-dependent mortality. While this is a challenging model to solve exactly, we can straightforwardly simulate a stochastic process with these three event types (resource inflow, consumption, and consumer mortality). In this simulation approach, we define $T\{X \to X + 1\}$ as the transition rate at which the species with abundance X gains one individual, while all other abundances stay precisely the same. In this notation, the following transition rates represent the events that occur in our stochastic process:

$$
\begin{aligned}
T\{R_k \to R_k + 1\} &= \rho_k \\
T\{R_k \to R_k - 1\} &= (1 - \epsilon)R_k \sum_j C_{kj} N_j \\
T\{R_k \to R_k - 1, N_i \to N_i + 1\} &= \epsilon N_i C_{ki} R_k \\
T\{N_i \to N_i - 1\} &= \eta_i N_i \ .
\end{aligned}
\tag{1}
$$

The basic event types that lead to the model in Eq 1 are conceptually represented in Fig 1. ρ_k is the inflow rate of the k-th resource. C_{ij} are non-negative coefficients which measure the rate of consumption of resource i by consumer j, forming a $K \times S$ matrix, C. In this study we set $K = S$ for simplicity of the analysis, but our main mathematical results hold in general (see S1 Appendix section 3). ϵ measures the consumers' efficiency in converting resources into consumer biomass, being the probability that a consumption event between resource k and consumer i results in a new individual of consumer i (in other words, every consumption event results in a loss of one resource unit, and a fraction ϵ of them also result in a gain of one consumer). We originally set ϵ to a small fraction, reflecting a need for many resource units to be consumed before each reproduction event, but this had no effect on outcomes other than to slow down the dynamics. Thus we set $\epsilon = 1$ for expediency. η_i is the per capita mortality rate for the i-th consumer.

Every consumer goes extinct over long enough time scales, but in any real ecological system, speciation will tend to maintain species richness. Here we represent this balance in a simple way: every time a species with abundance $N_i = 1$ is selected for a death event (extinction), we keep it at $N_i = 1$, while logging the time it took for the death event to occur. This is akin to point speciation [35], although for simplicity we introduce the 'new' species with the same resource profile as the 'extinct' species. Alternatively, this process could be interpreted as very rare immigration events for that same species. While speciation and/or immigration can be modeled at increasing levels of realism (e.g. [36, 37]), our results are primarily based on the dynamics of extant species, and hence are largely insensitive to the details of this process. We confirmed this robustness for the specialists scenario under an alternative speciation/immigration scheme (Fig 10 in S1 Appendix).

Fig 1. Model scheme. The left panel is a schematic of the models in Eqs 1 and 2. Resources flow into the system at fixed rate ρ_i, while consumers die at rates $\eta_i N_i$. The arrows connecting resources to consumers represent non-zero consumption coefficients C_{ij}. In this example, all consumers are capable of consuming resource 2, so there is community-wide competition for it. At the same time, consumer 1 is the only consumer that can utilize resource 1 and, similarly, consumer 3 is the only consumer that can deplete resource 3. The right panels are examples of the C matrices resulting from our two different parametrizations. In the generalist case, each entry is an independent sample from the same probability distribution with positive support. In the specialist case, consumer i can consumer resource i more quickly than any other resource. Other than this special resource, consumers consume every other resource at the same rate.

In parallel, we consider a deterministic version of the same model:

$$\dot{R}_k = \rho_k - R_k \sum_j C_{kj} N_j$$
$$\dot{N}_i = \epsilon N_i \sum_j C_{ji} R_j - \eta_i N_i \; . \tag{2}$$

Using this deterministic model, we can calculate consumer and resource abundances at equilibrium. We will focus on coexistence solutions, i.e. equilibria of Eq 2 where all resources and consumers have positive equilibrium abundances [38]. While ecological disturbance or other perturbations can lead to important transient or cyclic behavior [39, 40], stable equilibria attract all nearby transient states, and therefore have a special role in determining the behavior of a dynamical system. Accordingly, they have been studied extensively in community ecology [41–43]. Let's denote the abundances by vectors $\vec{R} = \begin{bmatrix} R_1 & \dots & R_K \end{bmatrix}^T$ and $\vec{N} = \begin{bmatrix} N_1 & \dots & N_S \end{bmatrix}^T$ and the inflow and mortality rates by vectors $\vec{\rho} = \begin{bmatrix} \rho_1 & \dots & \rho_K \end{bmatrix}^T$ and $\vec{\eta} = \begin{bmatrix} \eta_1 & \dots & \eta_S \end{bmatrix}^T$. Then, given an invertible matrix C, we can choose positive inflow and mortality vectors $\vec{\rho}$ and $\vec{\eta}$ that lead to any

positive abundance vectors ($\vec{R}^*$ and $\vec{N}^*$) as equilibrium solutions of Eq 2. For simplicity, we will choose $\vec{\rho}$ and $\vec{\eta}$ so that all resources converge to a single value r and all consumers converge to a value n. The values of the fixed points of Eq 2 correspond approximately to the average values of the abundances in Eq 1, so the stationary distribution for each consumer and resource will have a mean of n and r respectively.

In both Eq 1 and Eq 2, consumer i is determined by its consumption preferences (the i-th column of C) and mortality rate η_i. Fixing a given set of equilibrium abundances determines $\vec{\eta}$, leaving the principal object of biological interest as the set of consumption preferences in C. For example, if all entries in C are identical, we recover a neutral limit where all consumer species have the same resource preferences. In this study we consider two different ways to move away from this neutral limit, shown in Fig 1. The first, which we call the specialist scenario, is when C has all diagonal entries equal to one value C_d, while all off-diagonal entries are equal to another value C_o, where $C_d > C_o$. In this case, each consumer has a unique preferred resource that they consume more quickly than all others. Other than this special resource, each species consumes all other resources at the same rate. If $C_d = C_o$, then all species are equivalent and we expect the stochastic version of this model to behave neutrally. In our second parametrization, called the generalist scenario, we sample each entry of C independently from a uniform distribution with mean μ and variance σ^2. When $\sigma^2 = 0$, the entries of C are simply μ and we expect the system to behave neutrally. As such, the magnitude of the ratio C_d/C_o in the specialists scenario, and the coefficient of variation $CV = \sigma/\mu$ of the matrix in the generalists scenario quantify the system's departure from neutral consumption preferences. Note that when $CV = 0$ in the generalist scenario and $C_d/C_o = 1$ in the specialist scenario, the species abundance distribution (SAD) can be shown analytically to be a log-series (see S1 Appendix section 2).

In the next section, we predict the threshold values of these two measures of non-neutrality, below which we predict the community dynamics will look neutral. In other words, we predict how far from neutrality the model can be and still display neutral outcomes.

Neutral-Niche Threshold

Species abundances in our stochastic process are constantly being perturbed away from the deterministic equilibrium by demographic noise. At the same time, consumer-resource feedbacks stabilize those abundances. The consumer community may therefore display neutral- or niche-like dynamics depending on the balance of influence between these forces. Because it is difficult to calculate the rates at which the abundances return to equilibrium in the stochastic model, we use the dynamics of the deterministic model near equilibrium as a proxy for the stabilizing force that the consumers experience in the stochastic system. In S1 Appendix, we show that the timescales of the dynamics of resources and consumers are naturally quite different, such that resources reach equilibrium and respond to perturbations much faster than consumers. As a result, resource abundances quickly converge to Poisson distributions (see Figs 3, 4 in S1 Appendix). This time-scale separation recapitulates the classic expectation of fast resource dynamics relative to consumer dynamics, but does not alone tell us whether or not the consumer system will be neutral-like. We thus turn our attention to consumer behavior for the remainder of this paper.

In order to identify a transition between niche and neutral dynamics, we consider a neutral community where the average species abundance is n. Let $\bar{\eta}$ be the average consumer mortality rate (such that $1/\bar{\eta}$ is the expected lifespan of an individual) and let T_n be the expected time to extinction (measured in individual lifespans) for a species undergoing drift from its mean abundance. We then use $T_n/\bar{\eta}$ as a characteristic timescale for drift in our model—i.e. $T_n/\bar{\eta}$ is the timescale over which drift takes a species with

average abundance all the way to extinction. On the other hand, we use $1/|\lambda_+|$ as the timescale for the stabilizing mechanism, where λ_+ is the most negative eigenvalue corresponding to consumer dynamics in the deterministic model (i.e. λ_+ determines the fastest timescale among the consumer dynamics). If the characteristic time for consumer abundances to deterministically return to equilibrium is much longer than the timescale for drift to cause large abundance fluctuations (ie. if $1/|\lambda_+| \gg T_n/\bar{\eta}$), then drift should dominate consumer dynamics. This inequality also suggests a threshold at which drift ceases to be a good description of the abundance patterns, namely when $1/|\lambda_+| \sim T_n/\bar{\eta}$.

We should note here that our estimate of the characteristic timescale of the stabilizing mechanism assumes that the linearized system is a good approximation to the actual dynamics throughout a large range of abundances. In fact, the linearized system is only guaranteed to be a good description in the neighborhood of the fixed point, but we hypothesize that it yields a good order-of-magnitude estimate for this transition more generally (for another example of this approach, see [44]).

In the S1 Appendix, we derive λ_+ in both of our parametrizations of C, and from it obtain analytical expressions for when drift should be a good description of abundance patterns at equilibrium as a function of the other parameters in the model. For the generalist parametrization, we find that the threshold for the coefficient of variation $CV = \sigma/\mu$ is

$$CV \sim \sqrt{\frac{S}{4n}\left(1 - \frac{\epsilon r}{n^2}\right)} \tag{3}$$

so the square of the threshold CV value should inversely depend on the mean consumer abundance n, if we disregard the second term in the parenthesis, which is small for our parameter choices. For the specialist parametrization, we derive a threshold scaling for the ratio between diagonal and off-diagonal elements of the consumption matrix, and find that

$$\frac{C_d}{C_o} \sim \frac{1 + \sqrt{\gamma}\,(S-1)}{1 - \sqrt{\gamma}} \tag{4}$$

where $\gamma = \frac{1}{n}\left(1 - \frac{\epsilon r}{n^2}\right)$. So, the specialist threshold is linear with the number of consumers in the community.

Testing for Neutrality

We now test our predicted threshold by comparing our observations to neutral predictions using both static and dynamical properties of the community. Neutrality can be generally understood as symmetry at the species level, which corresponds to indifference in the fate of individuals upon changing species labels among populations [45]. This can be achieved in several ways [45–47]. Here we take the definition of [12], whereby neutrality occurs when all individuals have the same probability of death and birth events regardless of species identity and abundance. This definition serves our purpose of showing that non-neutral species may display neutral behavior at the community level when niche differentiation is not sufficiently large.

Under neutrality, species abundances follow a distribution which, in high-diversity communities, converges to Fisher's log-series [48, 49]. The probability that a species has abundance k in this case is

$$P_k = \alpha \frac{p^k}{k}, \tag{5}$$

where $\alpha = -1/\log(1 - p)$, and the parameter p can be estimated from the mean abundance n by solving the equation $\alpha p/(1 - p) = n$. We fit the log-series distribution to our simulated communities using the discrete Cramér-von Mises goodness-of-fit test [50], and consider the fit successful if the p-value exceeds 0.05.

To test our analytical predictions from the previous section, we define the neutrality threshold as the point of departure from neutral resource preferences at which the probability of a successful log-series fit drops below 50% (other choices for this probability cutoff do not change results qualitatively, see S1 Appendix). We determine this probability by fitting the log-series distribution onto ensembles of simulated communities, and running a logistic regression of the successful and rejected fits against the neutrality index of the different ensembles. In the generalist scenario, the neutrality index is the coefficient of variation in the consumption matrix, while in the specialist scenario it is the ratio between on- and off-diagonal entries.

We next define a test for whether long-timescale neutral dynamics fail to hold when the degree of non-neutrality of the consumer model passes through our predicted transition. Given a species current abundance k, its life expectancy (i.e. time to extinction) T_k under neutral dynamics has been shown to be [51]

$$T_k = \frac{1}{1-p} \left(p^{-k} B(p; 1+k, 0) + H_k + \log(1-p) \right), \tag{6}$$

where H_k is the k-th harmonic number [52], $B(z; a, b)$ is the incomplete beta function, and p is the log-series parameter (related to the speciation rate in neutral metacommunity models with speciation events). The extinction time in Eq 6 is given in generation units, with a generation time defined as the inverse of the mean mortality rate of consumers, $1/\bar{\eta}$, i.e. the expected lifetime of an individual. In our numerical simulations, we obtain life expectancy estimates as follows. Over the course of one simulation, after a burn-in period to ensure stationarity, we record all species abundances every 30 minutes of simulation time, corresponding to roughly 54 (212) generations in the generalists (specialists) scenario, for a total of 30,000 generations of simulation runtime in either scenario. An extinction event occurs when a species with abundance 1 is chosen for a death event. Because in our model the state $N_i = 0$ is not allowed, multiple extinction events may occur to the same species. Also, multiple abundance recordings may occur between successive extinction events, such that when an extinction eventually occurs, it will be linked to all previously recorded abundances since the last extinction. The time intervals between abundance recordings were sufficiently large to ensure independent data points, and the total runtime was long enough to observe extinctions in species with a wide range of initial abundances.

Finally, we test for neutral behavior at shorter timescales by comparing temporal fluctuations in species abundance against neutral predictions. For a subcritical stochastic birth-death process representing neutral dynamics, [51] derived an expression for the probability that a species has abundance n at time t given initial abundance n_0 ([51], Eq. 2a). We use that expression to calculate the expected variance across histories of the stochastic process of the quantity $D(t) = \frac{n(t) - n_0}{\sqrt{n_0}}$ [24], against which we compare our numerical observations. We use an ensemble of 100 histories, sampling the communities every generation for up to 100 generations. A species is considered to have zero abundance upon extinction, and we calculate the ensemble variance of $D(t)$ across species with the same n_0, and then average across all species with $n_0 \geq 20$ (the cutoff was necessary because low-n_0 species displayed noisy behavior at longer timescales, and would require a larger ensemble).

Results

In the generalist scenario, when the coefficient of variation (CV) in the consumption matrix is low (reflecting that consumers have similar preferences for all resources), the log-series distribution fits the SAD (Fig 2A). However, as the CV of the consumption

matrix increases, the probability that the SAD conforms to the LS distribution declines (Fig 2B). For a community with 50 species and 50 resources, where the average species abundance is 245 and the average resource abundance is 100, a logistic regression indicates that the probability of a successful log-series fit drops below 50% when the CV of the consumption matrix is higher than 0.23. Defining this as the threshold for neutral-like patterns in the species abundance distribution, we note a power law between $CV^{\text{threshold}}$ and the mean species abundance n (Fig 2C), with an exponent close to our predicted value of -0.5 from Eq 3. A similar agreement occurs when we draw the consumption rates in the C matrix from a normal distribution rather than a uniform distribution (Fig 13 in S1 Appendix), as expected since the random-matrix theory we used to derive Eqns 3 and 4 is distribution-independent.

Fig 2. SAD results. Species abundance distribution (SAD) results for the generalists scenarios (A-C) and specialist scenarios (D-F). **A**: In the generalist scenario, the log-series distribution (black curve) fits the SAD when the coefficient of variation (CV) in the consumption matrix is sufficiently low (red points; Cramér-von Mises goodness-of-fit test p-value 0.78), but is rejected when the CV is sufficiently high (blue points; CvM test p-value 0.02). **B**: Probability that the logs-series distribution fits the SAD decreases with the CV of the consumption matrix. Points and error bars show the mean and standard error of the count of successful fits, out of an ensemble of 106 communities. Blue curve shows logistic regression. The threshold CV, defined as the point where the probability falls below 50%, is $CV^{\text{threshold}} = 0.23$. **C**: $\log(CV^{\text{threshold}})$ has a linear relationship with $\log(n)$, with slope -0.61 ± 0.09. This indicates a power law between $CV^{\text{threshold}}$ and n, with an exponent close to our analytic prediction of -0.5. Error bars show uncertainty propagated from the standard errors of the fitted parameters in the respective logistic regressions. Bands show the 95% CI of the linear regression. **D**: In the specialists scenario, communities with low C_d/C_o ratio (red points; CvM test p-value 0.267) conform to the log-series distribution, while communities with sufficiently high C_d/C_o (blue points; CvM test p-value < 0.001) reject the LS distribution. **E**: The probability of the LS distribution fitting the SAD decreases as C_d/C_o increases, with the threshold at $C_d^{\text{threshold}} = 4.0$ (we set $C_o = 1$). Each data point summarizes an ensemble of 143 communities. **F**: $C_d^{\text{threshold}}$ increases linearly with the number of species in the community, in qualitative agreement with Eq 4. Parameters: A-C: $K = S = 50$, $n = 500$, $r = 100$. D-F: $K = S = 50$, $n = r = 100$.

Results were analogous in the specialists scenario. Communities with low C_d/C_o ratio, indicating a small-magnitude difference between preferences for main and secondary resources, conform to the log-series distribution, while communities with sufficiently high C_d/C_o reject it (Fig 2D). Indeed, the probability that the log-series distribution fits the SAD decreases as C_d/C_o increases (Fig 2E). In communities with 50 species and 50 resources, with mean species and resource abundances at 100, the threshold ratio for a successful log-series fit is 4.0. This threshold increases linearly with the number of

species in the community (Fig 2F), as predicted by Eq 4.

Note we cannot use Eqn (4) to predict the slope in Fig 2F, as the formula for the threshold is defined up to a multiplicative constant. The reason for this constant is that we can estimate a characteristic timescale for relaxation after a perturbation under stabilized dynamics, but the exact time to relaxation depends on the size of the perturbation. We can always estimate the time it takes to get some arbitrary percentage of the way back to equilibrium–for example, the half-life of the decay back to the fixed point, but this introduces a multiplicative constant on the left side of our Eqns (3) and (4). The results shown herein are as quantitative as our formulas allow.

To compare the propensity for neutral-like abundance pattern across the generalist and specialist scenarios, we use a non-neutrality index applicable to both, based on the average cosine between the vectors representing the resource preference profiles of different species (i.e. the columns of the consumption matrix). The average cosine across all species pairs represents the average similarity between resource preference profiles. We therefore define the non-neutrality index as $NNI = 1 - \cos$. The limit of complete neutrality, i.e. identical resource preferences among all species, corresponds to a cosine of 1, and thus $NNI = 0$. The opposite limit of complete niche differentiation, i.e. where each species consumes a single resource with zero overlap, corresponds to a cosine of 0 ($NNI = 1$). Fig 3 shows that in both the generalist and specialist scenarios, the probability of neutral-like abundance distribution is close to 100% in the neutral limit, as expected, but remains positive as we deviate from neutrality. The probability of rejecting the log-series only falls below 50% when the NNI is as high as 0.15 ± 0.05 in the specialist scenario, and 0.22 ± 0.03 in the generalist scenario (Fig 3). For a given NNI, the abundance distribution in the generalist scenario typically appears more neutral-like than in the specialist scenario.

Fig 3. SAD results: generalists vs specialists. Probability of log-series SAD plotted against a non-neutrality index defined as NNI = 1 − cos, where cos is the cosine between vectors representing species resource preferences, averaged across all species pairs in the community. Complete neutrality would correspond to NNI = 0 (i.e. cos = 1), reflecting full overlap in resource preferences. For the same NNI value, communities in the generalist scenario are typically more likely to conform to the log-series distribution than communities in the specialist scenario. Parameters: $S = K = 50$, $n = r = 100$.

Extinction time results are shown in Fig 4. As expected, life expectancy increases with species abundance. In the generalist scenario, species extinction times in communities with low CV closely match predictions from the neutral model, whereas communities with sufficiently high CV depart from neutral predictions (Fig 4A). When there is a poor match, the neutral model underpredicts the extinction times, especially for species with high abundance (Fig 4A). Plotting observed extinction times against neutral predictions in communities with different CV reveals that higher CVs lead to increasingly poor matches to neutrality, especially for high-abundance species (Fig 4B). Interestingly, observed extinction times seem linearly related to predictions regardless of the CV. The slope of this relationship increases with the CV (Fig 4C), being close to 1 at low CV, indicating a perfect match to neutral predictions, and > 1 at higher CV. This indicates that species of high mean abundance have particularly long life expectancy beyond neutral expectations, suggesting that niche differentiation has a disproportionate stabilizing effect on common species.

Fig 4. Extinction time results. Extinction time results for the generalist (A-C) and specialist (D-F) scenarios. **A**: Points and error bars show average and standard errors of extinction times for species in logarithmically binned abundance categories. Curves show neutral predictions. In the generalist scenario, species extinction times match predictions from the neutral model in communities with low CV, but consistently exceed neutral predictions in communities with high CV, especially for high-abundance species. **B**: Plotting observed versus predicted extinction times in communities with different CV (colors) reveals that those with low CV conform closely to the neutral predictions (black line illustrates a perfect match), while higher CVs lead to increasingly poor matches to neutrality, especially for high-abundance species. Note that extinction times seem linearly related to predictions regardless of the CV. **C**: The slope of this relationship increases with the CV, being close to 1 (perfect match to neutral predictions) at low CV and > 1 at higher CV. **D-F**: Results in the specialist scenarios are analogous to the generalist scenarios, except that extinction times of high-abundance species saturate. Parameters: A-C: $K = S = 50$, $n = 500$, $r = 100$. D-F: $K = S = 50$, $n = r = 100$. Summary statistics were obtained from ca. 5,000 to 20,000 data points for each abundance bin in the generalists scenario, and 1,000 to 7,000 data points in the specialists scenario.

The specialist scenarios showed analogous results to the generalist scenarios regarding matches to neutrality or lack thereof, except that extinction times of high-abundance species saturate. This could be because the restoring force in this scenario increases very quickly for large deviations from equilibrium, so a species that fluctuates to high abundance almost immediately returns to its mean abundance, thus not significantly increasing its life expectancy. By contrast, the generalist case lacks a strong stabilizing force, so excursions towards high abundance tend to substantially increase time to extinction.

Non-neutral community dynamics also displayed neutral-like behavior at shorter timescales (Fig 5). Temporal fluctuations in species abundances in the specialist scenario were indistinguishable from neutrality for up to 10 generations, even when the non-neutrality index was as high as 0.5. At longer timescales, niche stabilization tended to

reduce the ensemble variance of $D(t) = \frac{n(t)-n_0}{\sqrt{n_0}}$ relative to neutrality, as expected. However, it was only at maximum non-neutrality (NNI = 1), corresponding to fully specialized and therefore non-interacting species, that var(D) was immediately distinguishable from neutral, plateauing within a few generations.

Fig 5. Shorter timescale results. Variance in species abundances over time, for different parametrizations of the specialist scenario. Vertical axis plots the variance across histories of the stochastic process of $D(t) = \frac{n(t)-n_0}{\sqrt{n_0}}$, which is then averaged across species with different initial abundances n_0 (bands show standard error of the mean). Colors show parametrizations with increasing non-neutrality index (NNI), with lines showing the loess regression with smoothing parameter set to 1. Black line shows neutral prediction. Inset highlights similarity of all curves except NNI = 1 at timescales up to 10 generations.

Notably, the neutral-like behavior observed here applies to communities whose resource preference profiles would be easily distinguishable from true neutrality, if they were directly measured. In other words, our predictions and results are not solely for "small" departures from neutrality, and yet we still observe neutral-like static and temporal patterns. To quantify this, we consider a neutral community of same richness and size as in Fig 3. If all consumption events under true neutral dynamics are observed over the course of an average species lifetime (in this case, approximately 250 generations), the 95% confidence interval for the average pairwise cosine is estimated to lie between 0.992 and 0.996 (see Fig 9 in S1 Appendix), well above the threshold for 50% rejection of the log-series abundance distribution. This means communities that are unmistakably non-neutral under direct observations of consumption events would still look neutral

under summary distributions like the SAD or extinction times. Of course, in most if not all cases it is unfeasible to directly observe consumption events, especially over such long timescales. This impracticality is one reason why ecologists have indeed tested theories of community assembly using summarizing distributions.

Discussion

We showed that communities of competing species may appear neutral despite differences in resource preferences that confer niche stabilization, and this behavior is quite predictable from summarized properties of the consumer-resource network. We provided quantitative underpinnings for why non-neutral dynamics may lead to neutral outcomes, namely because the niche stabilization provided by species differences in resource use profiles may not be strong enough to overcome the influence of demographic stochasticity in generating patterns at the community scale. Formally, this is reflected in the small eigenvalues of the Jacobian of the deterministic counterpart of our stochastic resource-consumer model.

Our simulation results qualitatively supported our hypothesis that when the relaxation time to equilibrium in the stabilized resource-consumer system is comparable to the timescales of drift to extinction under neutral dynamics, one cannot easily reject the log-series abundance distribution, and species life expectancies match neutral predictions. Furthermore, at generational timescales, even full knowledge of species abundances over time may not suffice to reject neutrality. This has implications for inferring niche differentiation in real life systems, such as distinct metabolic syndromes in microbial communities, where the differences between consumers may not be large enough to overcome the influence of drift.

Our study demonstrates the disconnect between neutrality at the pattern level—i.e. whether the distribution and temporal variation in species abundances are distinguishable from neutral predictions, and neutrality at the process level—i.e. whether species differ in resource preferences. In doing so, it builds on the point made by others that neutral SADs do not imply neutral ecology [1, 2]. Furthermore, our connecting emergent neutrality to stabilization timescales in a stochastic world provides a mechanistic explanation for the success of neutrality as a macroecological theory.

Earlier investigations of neutral-like outcomes in non-neutral communities have focused on phenomenological models of species interactions such as Lotka-Volterra competition [32, 33]. Our approach differs from these studies in two main ways. First, these approaches compare neutral dynamics to niche dynamics combined with exogenous immigration. That makes it challenging to isolate the stabilizing effects due to niche differentiation from those arising from immigration. By contrast, we model an essentially closed system and thus compare drift dynamics to stabilization driven almost purely by niche differentiation (strictly speaking, immigration is present but very low, only affecting the abundances of the rarest species). Second, in these stochastic Lotka-Volterra approaches, the corresponding deterministic dynamics do not completely determine the stochastic transition rates, leaving the degree of demographic noise as a free parameter that can be tuned to adjust the balance of noise relative to deterministic ecological selection. In our stochastic consumer-resource model, the level of noise is set endogenously by the structure of the Gillespie simulation. Our finding that neutral dynamics are still possible demonstrates that drift is a natural outcome of self-contained niche-differentiated ecological dynamics without tuning the degree of noise. While the previous approaches analytically solve simplified versions of stochastic Lotka-Volterra dynamics, we derive an analytical prediction for when and why the full stochastic consumer-resource model appears neutral.

In this context of consumer-resource dynamics, neutral behavior has been previously

shown to arise from consumers competing for a small number of resources [34]. In that study, consumers satisfy fine-tuned metabolic tradeoffs that allow an arbitrarily large number of species to coexist on a neutrally stable manifold of fixed points, on which drift fully drives the stationary distribution. However, if consumers violate the precise metabolic tradeoff constraints, then only a small number of consumers coexist at one fixed point, and the stationary abundance distribution of these consumers is no longer guaranteed to appear neutral. In contrast, the numbers of consumers and resources are equal in our modeling framework, so the consumers coexist at a unique fixed point. In our stochastic model, it is the timescale of the stabilizing dynamics, which is determined by the degree of heterogeneity in resource preferences, that gives rise to neutral dynamics.

The neutrality threshold herein defined is a useful construct for testing our analytical predictions, not a methodological prescription for field ecologists. Indeed, the threshold's numerical value depends on the choice of cutoff for the probability of rejecting the log-series. In reality, this probability increases gradually rather than abruptly with increasing departures from neutral resource preferences. Furthermore, one's ability to reject neutral pattern in a community of interest depends not only on similarities between neutral and non-neutral outcomes, but also on the statistical power of goodness-of-fit tests. As such, in order to test our analytical predictions, the threshold's numerical value is less meaningful than how it scales with changes in parameters such as community size and richness. We verified that this scaling is robust to different cutoff choices, and also to different tests and methodological approaches for rejecting neutrality (see Fig 8 and Fig 11 in S1 Appendix).

In nature, the consumption matrix could be far more complex than the two cases considered here, reflecting diverse metabolic strategies of consumers from different environments and evolutionary histories. For example, groups of the columns of C could be highly correlated when the corresponding groups of consumers perform similar functional roles in the community. The community as a whole could also exhibit a stronger preference for certain resources over others, introducing correlations between the rows of C. We ignored these complications for the sake of analytical tractability, but our results are likely generalizable: it is the bipartite structure of consumer-resource models, and not the specific structure of C, that separates the spectrum of our deterministic model into consumer and resource bulks. Because it is this separation that causes long excursions from the mean abundance in the stochastic model, our results suggest that a transition from neutral to niche dynamics will be a feature of consumer-resource models with more complex consumption preferences than considered here.

The fact that abundances in specialist communities are easier to distinguish from neutrality than in generalist communities with similar niche overlap suggests that specializing towards a single resource has a stronger niche-like impact on abundance distributions than unstructured variation in resource preferences. It follows that departures from neutrality are not all equal with regards to impact on abundances, leaving open the possibility that some special structure in the consumption matrix may very quickly lead to detectable differences from neutral pattern. One method to test this possibility would be to infer the consumption matrix by fitting our consumer-resource model to experimental abundance measurements from monoculture experiments on each resource. Then, different consumption matrix structures could be investigated by choosing which consumers to include in community experiments. Although they did not infer a consumption network, [53]'s microbial community assembly experiments found that the resulting community structure was highly variable at the species level, while highly predictable at the family level. Our theoretical results suggest a possible interpretation of these findings. When viewed at the species level, consumers do not differ enough to strongly affect the outcome of community assembly, suggesting that these species differences are below our analytical threshold and drift is the primary driver of abundance dynamics.

On the other hand, consumer abundances converge to nearly deterministic outcomes at the family level, suggesting that between-family differences in resource profiles are well above our threshold.

Supporting Information

S1 Appendix. Characterizing the spectrum of the deterministic model; Master equations for our stochastic model; Discussion of abundance distributions in the neutral and niche limits; Fitting outcomes using different cutoff values;

References

1. McGill BJ. Strong and weak tests of macroecological theory. Oikos. 2003;102(3):679–685.

2. McGill BJ, Etienne RS, Gray JS, Alonso D, Anderson MJ, Benecha HK, et al. Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. Ecology letters. 2007;10(10):995–1015. doi:10.1111/j.1461-0248.2007.01094.x.

3. Harte J. Maximum entropy and ecology: a theory of abundance, distribution, and energetics. OUP Oxford; 2011.

4. Harte J. Tail of death and resurrection. Nature. 2003;424(6952):1006–1007. doi:10.1038/4241006a.

5. Gotelli NJ, McGill BJ. Null versus neutral models: what's the difference? Ecography. 2006;29(5):793–800.

6. Alonso D, Etienne RS, McKane AJ. The merits of neutral theory. Trends in Ecology & Evolution. 2006;21(8):451–7. doi:10.1016/j.tree.2006.03.019.

7. Clark JS. Beyond neutral science. Trends in ecology & evolution. 2009;24(1):8–15. doi:10.1016/j.tree.2008.09.004.

8. Rosindell J, Hunter MO, Etienne RS. The unified neutral theory of biodiversity and biogeography at age ten. Trends in ecology & evolution. 2011;26(7):340–8. doi:10.1016/j.tree.2011.03.024.

9. Clark JS. The coherence problem with the Unified Neutral Theory of Biodiversity. Trends in Ecology & Evolution. 2012;27(4). doi:10.1016/j.tree.2012.02.001.

10. Rosindell J, Hubbell SP, He F, Harmon LJ, Etienne RS. The case for ecological neutral theory. Trends in Ecology & Evolution. 2012;27(4):204–209. doi:10.1016/j.tree.2012.01.004.

11. Caswell HAL. Community Structure: A Neutral Model Analysis. Ecological Monographs. 1976;46(3):327–354.

12. Hubbell SP. The Unified Neutral Theory of Biodiversity and Biogeography. Princeton, NJ: Princeton University Press; 2001.

13. Chave J. Neutral theory and community ecology. Ecology Letters. 2004;7(3):241–253. doi:10.1111/j.1461-0248.2003.00566.x.

14. Wootton JT. Field parameterization and experimental test of the neutral theory of biodiversity. Nature. 2005;433:309–312.

15. McGill BJ, Maurer BA, Weiser MD. Empirical evaluation of neutral theory. Ecology. 2006;87(6):1411–1423.

16. Volkov I, Banavar JR, Hubbell SP, Maritan A. Patterns of relative species abundance in rainforests and coral reefs. Nature. 2007;450(6197):45–49. doi:10.1038/nature01883.

17. Bell G. Neutral macroecology. Science (New York, NY). 2001;293(5539):2413–8. doi:10.1126/science.293.5539.2413.

18. Volkov I, Banavar JR, Hunter MO, Maritan A. Neutral theory and relative species abundance in ecology. Nature. 2003;424(6952):1035–7. doi:10.1038/nature01883.

19. Clark JS, McLachlan JS. Stability of forest biodiversity. Nature. 2003;423(6940):635–638. doi:10.1038/nature01632.

20. Volkov I, Banavar JR, Maritan A, Hubbell SP. The Stability of Forest Biodiversity - Reply. Nature. 2004;427(6976):696–697.

21. Pigolotti S, Flammini A, Marsili M, Maritan A. Species lifetime distribution for simple models of ecologies. Proceedings of the National Academy of Sciences of the United States of America. 2005;102(44):15747–15751. doi:10.1073/pnas.0502648102.

22. Ostling AM. Large-scale spatial synchrony and the stability of forest biodiversity revisited. Journal of Plant Ecology. 2012;5(1):52–63. doi:10.1093/jpe/rtr035.

23. Wang S, Chen A, Fang J, Pacala SW. Why abundant tropical tree species are phylogenetically old. Proceedings of the National Academy of Sciences. 2013;110(40):16039–16043. doi:10.1073/pnas.1314992110.

24. Kalyuzhny M, Seri E, Chocron R, Flather CH, Kadmon R, Shnerb NM. Niche versus Neutrality: A Dynamical Analysis. The American Naturalist. 2014;184(4):439–446. doi:10.1086/677930.

25. O'Dwyer JP, Kembel SW, Sharpton TJ. Backbones of evolutionary history test biodiversity theory for microbes. Proceedings of the National Academy of Sciences. 2015;112(27):201419341. doi:10.1073/pnas.1419341112.

26. Fung T, O'Dwyer JP, Kassim AR, Fletcher CD, Chisholm RA. Reproducing static and dynamic biodiversity patterns in tropical forests : the critical role of environmental variance. Ecology. 2016;97(5):1207–1217.

27. D'Andrea R, Ostling A. Biodiversity maintenance may be lower under partial niche differentiation than under neutrality. Ecology. 2017;98(12):3211–3218. doi:10.1002/ecy.2020.

28. Chave J, Muller-Landau HC, Levin S. Comparing classical community models: theoretical consequences for patterns of diversity. The American Naturalist. 2002;159(1):1–23. doi:10.1086/324112.

29. Chisholm RA, Pacala SW. Niche and neutral models predict asymptotically equivalent species abundance distributions in high-diversity ecological communities. Proceedings of the National Academy of Sciences of the United States of America. 2010;107(36):15821–5. doi:10.1073/pnas.1009387107.

30. Al Hammal O, Alonso D, Etienne RS, Cornell SJ. When Can Species Abundance Data Reveal Non-neutrality? PLOS Computational Biology. 2015;11(3):e1004134. doi:10.1371/journal.pcbi.1004134.

31. Rael RC, D'Andrea R, Barabás G, Ostling A. Emergent niche structuring leads to increased differences from neutrality in species abundance distributions. Ecology. 2018;doi:10.1002/ecy.2238.

32. Haegeman B, Loreau M. A mathematical synthesis of niche and neutral theories in community ecology. Journal of theoretical biology. 2011;269(1):150–65. doi:10.1016/j.jtbi.2010.10.006.

33. Fisher CK, Mehta P. The transition between the niche and neutral regimes in ecology. Proceedings of the National Academy of Sciences. 2014;111(36):13111–13116. doi:10.1073/pnas.1405637111.

34. Posfai A, Taillefumier T, Wingreen NS. Metabolic Trade-Offs Promote Diversity in a Model Ecosystem. Physical Review Letters. 2017;118(January):1–5. doi:10.1103/PhysRevLett.118.028103.

35. Etienne RS, Alonso D, McKane AJ. The zero-sum assumption in neutral biodiversity theory. Journal of theoretical biology. 2007;248(3):522–36. doi:10.1016/j.jtbi.2007.06.010.

36. Schreiber SJ, Tobiason GA. The evolution of resource use. Journal of Mathematical Biology. 2003;47:56–78. doi:10.1007/978-0-387-70984-0.

37. Rosindell J, Cornell SJ, Hubbell SP, Etienne RS. Protracted speciation revitalizes the neutral theory of biodiversity. Ecology Letters. 2010;13(6):716–727. doi:10.1111/j.1461-0248.2010.01463.x.

38. Saavedra S, Rohr RP, Bascompte J, Godoy O, Kraft NJB, Levine JM. A structural approach for understanding multispecies coexistence. Ecological Monographs. 2017;87(3):470–486.

39. Scheffer M, van Nes EH. Self-organized similarity, the evolutionary emergence of groups of similar species. Proceedings of the National Academy of Sciences of the United States of America. 2006;103(16):6230–6235. doi:10.1073/pnas.0508024103.

40. Jansen VAA, Sigmund K. Shaken not stirred: On permanence in ecological communities. Theoretical Population Biology. 1998;54(3):195–201. doi:10.1006/tpbi.1998.1384.

41. Armstrong RA, McGehee R. Coexistence of Species Competing for Shared Resources. Theoretical Population Biology. 1976;9:317–328.

42. Armstrong RA, McGehee R. Competitive Exclusion. The American Naturalist. 1980;115(2):151–170.

43. Barabás G, Pásztor L, Meszéna G, Ostling AM. Sensitivity analysis of coexistence in ecological communities: theory and application. Ecology Letters. 2014;17:1479–1494. doi:10.1111/ele.12350.

44. Fort H, Scheffer M, van Nes EH. The paradox of the clumps mathematically explained. Theoretical Ecology. 2009;2(3):171–176. doi:10.1007/s12080-009-0040-x.

45. Volkov I, Banavar JR, He F, Hubbell SP, Maritan A. Density dependence explains tree species abundance and diversity in tropical forests. Nature. 2005;438(7068):658–661. doi:10.1038/nature04030.

46. Vallade M, Houchmandzadeh B. Analytical solution of a neutral model of biodiversity. Physical review E, Statistical, nonlinear, and soft matter physics. 2003;68(6 Pt 1):061902. doi:10.1103/PhysRevE.68.061902.

47. Azaele S, Suweis S, Grilli J, Volkov I, Banavar JR, Maritan A. Statistical mechanics of ecological systems: Neutral theory and beyond. Reviews of Modern Physics. 2016;88(3):1–31. doi:10.1103/RevModPhys.88.035003.

48. Fisher RA, Corbet AS, Williams CB. The Relation Between the Number of Species and the Number of Individuals in a Random Sample of an Animal Population. Journal of Animal Ecology. 1943;12(1):42–58.

49. Alonso D, McKane AJ. Sampling Hubbell's neutral theory of biodiversity. Ecology Letters. 2004;7(10):901–910. doi:10.1111/j.1461-0248.2004.00640.x.

50. Arnold TB, Emerson JW. cvm.test: Discrete Cramer von Mises Goodness-of-Fit Tests; 2010. Available from: `https://r-forge.r-project.org/projects/ks-test/`.

51. Chisholm RA, O'Dwyer JP. Species ages in neutral biodiversity models. Theoretical Population Biology. 2014;93:85–94. doi:10.1016/j.tpb.2014.02.002.

52. Abramowitz M, Stegun IA. Handbook of mathematical functions: with formulas, graphs, and mathematical tables. vol. 55. Courier Corporation; 1965.

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