nature ecology & evolution

ARTICLES

https://doi.org/10.1038/s41559-019-1059-z

Corrected: Publisher Correction

Predicting coexistence in experimental ecological communities

Daniel S. Maynard[1,2*], Zachary R. Miller[2] and Stefano Allesina[2,3]

The study of experimental communities is fundamental to the development of ecology. Yet, for most ecological systems, the number of experiments required to build, model or analyse the community vastly exceeds what is feasible using current methods. Here, we address this challenge by presenting a statistical approach that uses the results of a limited number of experiments to predict the outcomes (coexistence and species abundances) of all possible assemblages that can be formed from a given pool of species. Using three well-studied experimental systems—encompassing plants, protists, and algae with grazers—we show that this method predicts the results of unobserved experiments with high accuracy, while making no assumptions about the dynamics of the systems. These results demonstrate a fundamentally different way of building and quantifying experimental systems, requiring far fewer experiments than traditional study designs. By developing a scalable method for navigating large systems, this work provides an efficient approach to studying highly diverse experimental communities.

Ecologists have long used small experimental systems to explore and test fundamental ecological phenomena: Gause's experiments[1] illustrated and popularized the principle of competitive exclusion, chemostat-based systems highlighted the need for eco-evolutionary modelling[2], stage-structured populations grown in laboratory conditions offered a clear example of chaotic dynamics[3] and yeast communities have provided key insight into ecological tipping points[4].

Despite these advances, experimenting with more diverse ecological communities has proven exceedingly difficult for several reasons. First, a pool of n species can give rise to as many as 2^n-1 distinct species assemblages (that is, combinations of the presence and absence of species). Even for moderate n, these are too many assemblages to handle experimentally. Second, rarely will all possible assemblages lead to coexistence; often, large communities collapse to smaller subsets of species due to extinctions (for example, a two-species community collapsing to a monoculture). Finally, two replicates of the same assemblage might lead to different outcomes—due to multistability or stochasticity, for example—complicating efforts to characterize the dynamics of the system.

These challenges have been raised in numerous studies that experimented with speciose communities[5–8]. A common approach is to seed plots with single species (monocultures), pairs of species and several larger assemblages, often chosen at random[9]. Owing to the number of species considered, however, the space of possible assemblages is so large that only a small fraction of assemblages can be tested. For example, the Biodiversity II experiment included 18 different plant species, from which 168 of the $2^{18}-1=262{,}143$ possible species assemblages were planted, with their dynamics tracked for decades[10]. Ideally, one would use a principled method to choose which assemblages to test, to maximize the probability of coexistence across assemblages (and therefore not waste experiments), while ensuring that the selected subset of experiments allows for robust inference across the full set of unobserved assemblages.

The goal of this work is to address these challenges by developing a method in which a limited number of observations is used to predict coexistence for all 2^n-1 assemblages that can be formed from a pool of n species. For any combination of species in the pool, our method attempts to determine whether the species can coexist and, if so, at what abundances. As a starting point, a limited number of experiments should be conducted, using a variety of assemblages, and species abundances should be measured at a single (final) time point in each experimental community. These modest requirements make this method ideally suited for studying speciose communities. While this approach can be applied to any ecological system, it is likely to be most useful in experimental settings where stochasticity and exogenous factors can be minimized.

The approach is summarized in Box 1 and Fig. 1, and explained in detail in the Methods. By recording the abundances of species grown in different combinations (a snapshot of abundances that we refer to as a dynamical endpoint), we parameterize a matrix B that encodes the abundances of species in all possible combinations—including those not observed experimentally. Armed with this estimate, we can predict the outcomes of unobserved experiments: for each possible species assemblage, our method returns the probability that the species will coexist, and, if they do, predicts the abundance of each species in the assemblage. We test this approach using data from three published studies, showing that it can predict out-of-fit experimental outcomes with high accuracy. We then explore the flexibility of this method by simulating data from a variety of nonlinear, non-equilibrium and non-pairwise dynamical models of species interactions. Finally, we show that the number of experiments needed to fit this model scales linearly with the number of species, requiring on the order of n experiments to predict the outcomes of all 2^n-1 possible assemblages.

Results

Empirical systems. To demonstrate this method, we first analyse published data of three experimental systems. Although these studies were not specifically designed to test our method, each study reports the final abundances of all species in each experimental assemblage, and each contains a sufficient number of unique assemblages to benchmark the method by making out-of-fit predictions.

[1]Institute of Integrative Biology, ETH Zürich, Zürich, Switzerland. [2]Department of Ecology & Evolution, University of Chicago, Chicago, IL, USA. [3]Northwestern Institute on Complex Systems, Evanston, IL, USA. *e-mail: daniel.maynard@usys.ethz.ch

Box 1 | Approach

Goal and applications. Given a pool of n species, our method aims to predict the outcomes of all 2^n-1 possible assemblages. For each assemblage, we predict whether the species can coexist and, if so, at what abundances.

This method can aid experimentalists in a variety of settings. For example, when studying the relationship between diversity and productivity, one could conduct a pilot study using as few as $2n$ plots (for example, all monocultures and all leave-one-out communities), and then use the method to predict productivity in all other possible assemblages. Similarly, the method could be used to identify candidate assemblages with maximal productivity (for example, for the development of biofuels). The method can also be used as a road map to build large experimental communities in which all species coexist—a challenging feat when selecting assemblages at random[7,22,23].

Input. The input for the method is a list of empirical endpoints, each recording the abundance of species in one subcommunity at the end of the experiment. Given a sufficiently large and diverse set of endpoints (Methods), one can use these data to predict coexistence for all unobserved assemblages. In particular, one needs to start with a variety of assemblages that differ in species composition or initial densities. Each assemblage is allowed to follow its dynamical trajectory until a predetermined time or until the condition is reached (for example, optical density or chlorophyll content stabilizes when dealing with bacteria or algae), at which point the abundances of all extant species in the community are recorded. These endpoint measurements are then organized into a matrix, with columns corresponding to species and rows to communities (Fig. 1).

Model. To start with a simple, extensible model, we take the endpoint abundance of each species i in each endpoint assemblage k to be a linear function of the endpoint abundances of the other species present in the final assemblage:

$$z_i^{(k)} = \gamma_i + \sum_{j\neq i} \tau_{ij} z_j^{(k)} \quad \text{for all } z_i^{(k)} > 0 \tag{1}$$

where k is an index referring to the unique set of species present at non-zero abundance in the endpoint measurement; where γ_i models the abundance of i when growing alone; and τ_{ij} is the per capita effect that species j has, if present, on the endpoint abundance of i. Ecologically, we expect $\gamma_i > 0$ for producers and $\gamma_i \leq 0$ for consumers and predators. Equation (1) can be rewritten more concisely as (Methods):

$$-1 = \sum_j B_{ij} z_j^{(k)} \quad \text{for all } z_j^{(k)} > 0 \tag{2}$$

where $B_{ii} = -1/\gamma_i$ and $B_{ij} = \tau_{ij}/\gamma_i$ for $i \neq j$. These coefficients are then collected into the $n \times n$ matrix B. The empirically measured endpoint abundances for k are denoted $x^{(k)}$, and are assumed to be a noisy estimate of the 'true' endpoint abundance $z^{(k)}$. See the Methods for details.

Parameter estimation. In theory, estimating the entries of B amounts to performing n linear regressions (Fig. 1c), effectively fitting a collection of n hyperplanes through the measured endpoints[29]. However, this straightforward approach ignores the error structure inherent in the data (Supplementary Information). We thus develop and implement a Bayesian regression approach that accounts for measurement error in each $x_i^{(k)}$ (Methods). This yields posterior distributions for B and for the estimated endpoint abundances $\hat{z}_i^{(k)}$.

Prediction. Under the model given by equation (1), B encodes all possible endpoints of the system. With an estimate of B in hand, we can thus determine whether any particular set of species can coexist by taking a matrix inverse (Fig. 1d and Methods). For example, to predict the outcome of the unobserved combination of species 1 and 2 in Fig. 1d, we take the corresponding 2×2 submatrix of B, invert it and compute the negative row sums. The resulting values are the predicted endpoint abundances for this two-species community. If any abundances are negative, this indicates that these species cannot coexist. If all abundances are positive, then these species may coexist, although the endpoint might be dynamically unstable, and therefore unreachable experimentally (Methods). The probability of coexistence for unobserved or out-of-fit assemblages is calculated as the proportion of posterior B matrices that result in the coexistence of a given assemblage.

Assumptions. This method is best suited to settings where exogenous factors (for example, dispersal or immigration) and stochasticity can be minimized. It is therefore ideal for controlled laboratory or microcosm studies, but may require higher replication in field-based systems to overcome experimental noise.

Although our fitting procedure assumes that endpoints are related in a linear fashion (equation (1)), it does not assume that the community dynamics are linear or additive (Extended Data Figs. 1–6 and Supplementary Information). The method does, however, require that dynamics play out long enough to ensure that species have time to go extinct, thereby preventing 'false positive' coexistence (Extended Data Fig. 8 and Supplementary Information).

Finally, when endpoint abundances are not believed to be linearly related, equation (1) can be modified to include additional terms (for example, higher-order interactions), and the same fitting approach can be applied without requiring substantially more experiments.

First, we consider data from Kuebbing et al.[11], who conducted growth experiments using two phylogenetically paired sets of old-field plants. Both sets contained four species, drawn from the families Asteraceae, Fabaceae, Lamiaceae and Poaceae. In one set, all species were native to the study site; in the other, all species were non-native. For both the native and non-native pools, experimental communities were initialized with 14 out of 15 (2^4-1) possible assemblages, with each combination replicated 10 times. Dry-weight aboveground biomass for each species in each assemblage was recorded after 112 days of growth[11].

Second, we analyse data from Rakowski and Cardinale[12], who grew consumer-resource communities consisting of five species of green algae and two herbivorous species from the family Daphniidae (*Ceriodaphnia dubia* and *Daphnia pulex*). Here we focus on the four algal species that survived in a sufficient number of endpoints to test our method: *Chlorella sorokiniana*, *Scenedesmus acuminatus*, *Monoraphidium minutum* and *Monoraphidium arcuatum*. The algae were grown in all four-species combinations with high replication, and one of the two herbivores was added to two-thirds of the replicates. The communities were incubated for 28 days, during which time each assemblage collapsed to a subset of the initial pool, generating a variety of distinct endpoints.

Third, we turn to time-series data published by Pennekamp et al.[13], who studied six species of ciliates growing on a bacterized medium, capturing 53 of the 63 (2^6-1) possible assemblages. We omit one species that declined in abundance across the time series

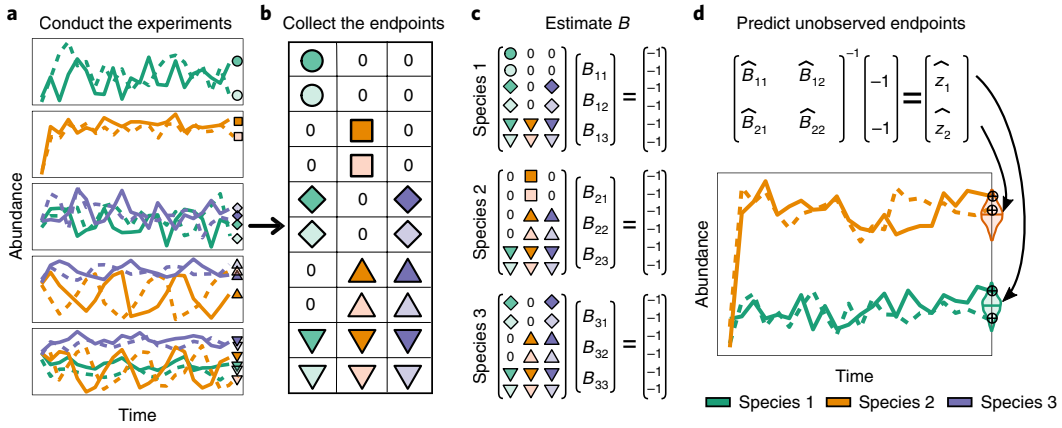

Fig. 1 | Predicting unobserved experimental outcomes. a, A set of experiments is conducted, with each assemblage varying in initial composition or initial abundances. Colours denote different species, and line styles denote each of two replicates. The experiments are run for a sufficient period, at which point a single snapshot is taken of each species' abundance in each replicate. **b**, These abundances are collected into a matrix, with columns corresponding to species and rows corresponding to endpoints. Here, shapes represent species' endpoint abundances for each of two replicates across five different assemblages. **c**, The matrix *B*, which defines the full set of endpoints, is estimated by fitting a separate hyperplane through each species' set of endpoints. **d**, To predict the endpoint of an unobserved assemblage—here, the combination of species 1 and 2—we take the corresponding submatrix of *B*, invert it and compute the negative row sums. To analyse the three empirical datasets, we implement a Bayesian approach to estimating *B*, resulting in a posterior distribution for the unobserved endpoint abundances (shown by the violin plots).

(*Tetrahymena thermophila*), leading to the five-species subsystem used here: *Colpidium striatum*, *Dexiostoma campylum*, *Loxocephalus* sp., *Paramecium caudatum* and *Spirostomum teres*. Each experimental assemblage was replicated twice, and each species' density in each assemblage was tracked for 57 days (ref.[14]). Experiments were repeated at six different temperatures ranging from 15 to 25 °C.

These three studies encompass ten distinct experimental systems: native and non-native plant communities, two different herbivore–algae subcommunities and six temperature conditions for protists. For each experiment, the species abundances for all assemblages were reported, providing a sufficient diversity of endpoints to fit our model and test the accuracy on out-of-fit assemblages using a jackknife leave-one-out approach (Methods and Supplementary Information). Specifically, for every system we omitted each of the k assemblages in turn, using the remaining $k-1$ assemblages to fit the model and predict the abundances of all species in the omitted community. We then compared our predictions with what was observed experimentally.

As the plant experiments included all but one of the possible combinations of species, we were also able to implement a more rigorous k-fold cross-validation approach. We used a subset of endpoints to parameterize the model and predict the endpoint abundances in the omitted communities; we repeated this procedure for all combinations of endpoints sufficient to fit the model, ranging from one to eight omitted communities (Supplementary Information).

Abundance patterns. As illustrated in Fig. 2, our method predicts the endpoints of each system with high accuracy. Across all assemblages, the medians of the posterior distributions generated by our method closely track the median observed abundances. Not only does our method achieve high accuracy for within-fit data (Fig. 2a,c,e); it also accurately estimates the out-of-fit endpoints and captures the observed variability in abundances (Fig. 2b,d,f).

By tallying the proportion of posterior draws of *B* that result in coexistence (Methods), we can also quantify the probability that an unobserved or out-of-fit assemblage will go extinct. For example, in both algal systems, our results suggest that the unobserved two-species community comprising *C. sorokiniana* and *M. minutum* has a high probability of extinction, regardless of the identity of the herbivore (Supplementary Figs. 5 and 6). Our findings also suggest that the four-species protist system comprising *C. striatum*, *D. campylum*, *P. caudatum* and *S. teres* has a non-zero risk of extinction at 19, 21 and 23 °C, driven in part by high variation and relatively low abundance of *D. campylum*, which causes the posterior abundance distribution to overlap zero (Supplementary Figs. 7–12). Verifying these predictions would require additional experimentation; yet such information can help identify which unobserved assemblages one might avoid in future experiments to minimize the risk of extinctions (see Study design, below).

For the plant data, a k-fold cross-validation approach demonstrates that only a few experimental endpoints are sufficient to predict abundances in all 15 assemblages with high accuracy (Supplementary Information). However, the choice of experiments used to fit the model (that is, the experimental design) is critical, as discussed in detail in the Study design section.

Finally, for the protist system, we have access to the full time series describing the dynamics, allowing us to compare the quality of fit of our method against traditional methods based on time-series data. The estimates obtained using only endpoint data compare favourably to those obtained using standard trajectory matching to fit a generalized Lotka–Volterra dynamical model to the full time series (Supplementary Information). Yet, by forgoing the need to model the dynamics, our method predicts endpoint abundances with a higher accuracy and requires only a fraction of the data.

Structure of *B*. Because our approach does not attempt to model the dynamics of these systems, *B* has no direct mechanistic interpretation—it simply encodes the patterns in the endpoint abundances (Methods). We also reiterate that the experimental studies used here were not designed to test our method, such that their use is intended only to illustrate this approach. Yet, by exploring the structure of *B*, we can demonstrate that our estimates are internally consistent and align with the known biology of the experimental systems, verifying that our method captures meaningful patterns in the data and is not overfitting the endpoints (Figs. 3 and 4; Supplementary Information).

For example, we find that the coefficients of *B*, which summarize the effect of one species' endpoint abundance on another's, are phylogenetically conserved between the native and non-native plant datasets (Fig. 3a). In the herbivore–algae system, the identity of the herbivore has little effect on the interactions among the algae,

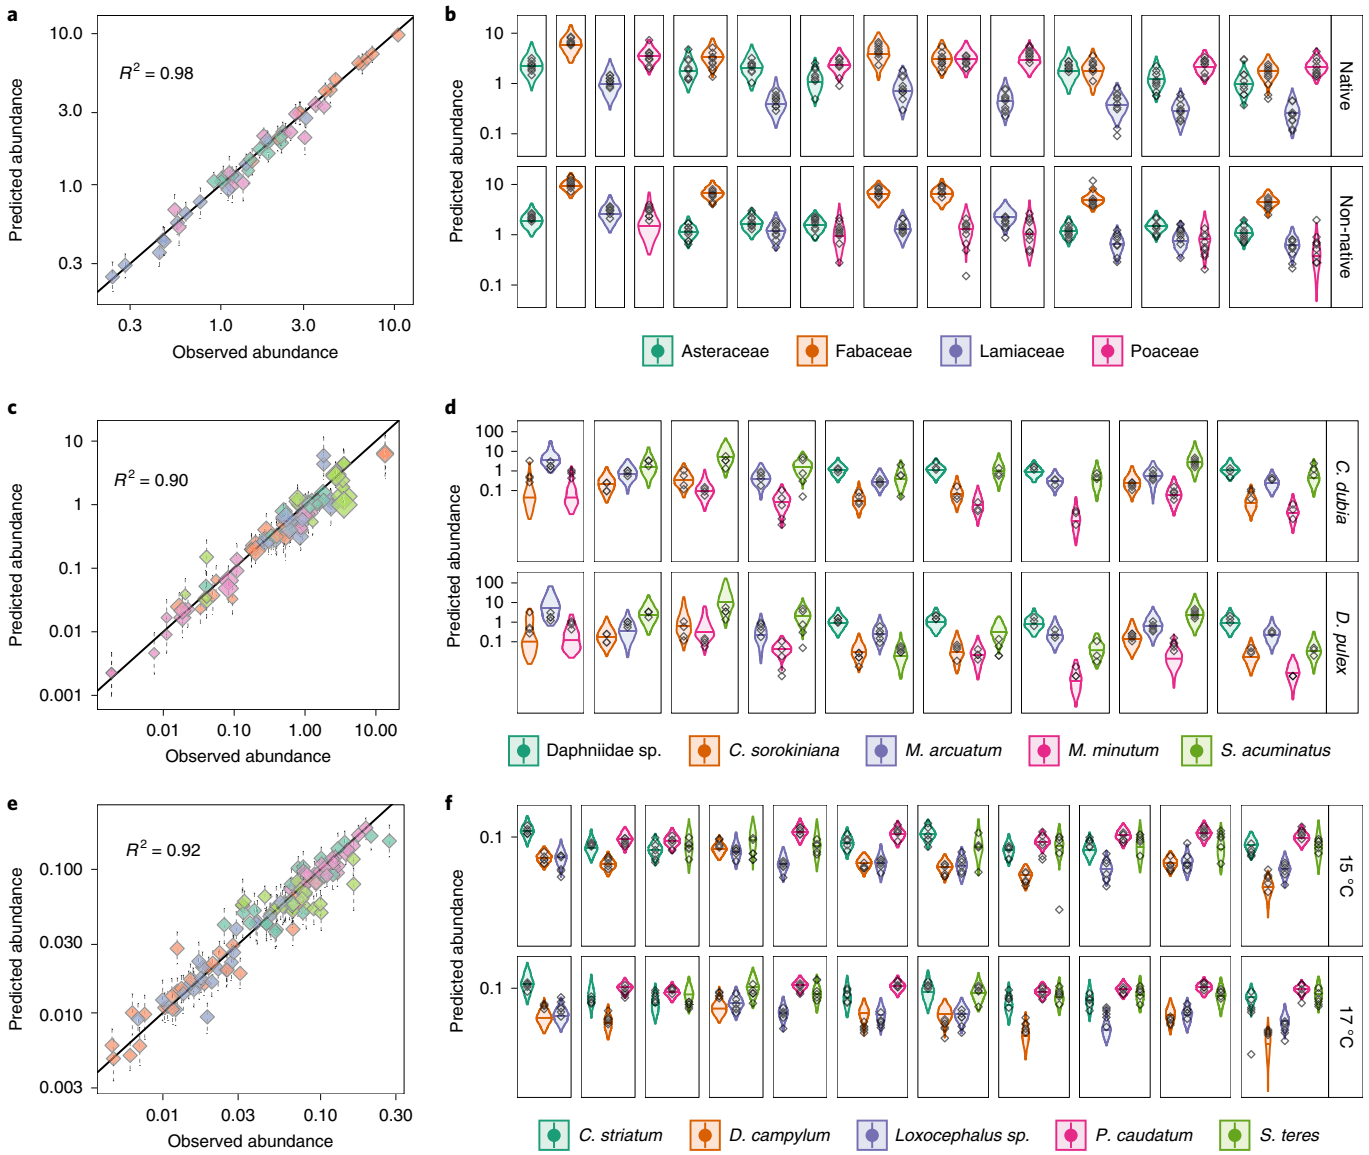

Fig. 2 | Results for three experimental systems. a,**b**, The old-field plant system comprising native and non-native communities. **c**,**d**, The herbivore–algae system with two different Daphniidae consumers. **e**,**f**, The ciliated protist system, assayed under six different temperature conditions, with 15 °C and 17 °C shown here. In **a**, **c** and **e**, the observed versus predicted abundances are shown for each experimental system, with vertical error bars denoting the 95% prediction interval of the posterior distribution; the sizes of the diamonds reflect the number of replicates per measurement. In **b**, **d** and **f**, the out-of-fit predicted abundances are shown for each assemblage, using a jackknife leave-one-out approach. Colours correspond to species, and each box corresponds to a different assemblage. The black diamonds show the measured abundance of each species in each replicate for each assemblage. The violin plots show the 95% prediction intervals for the posterior distribution of the abundance, with the horizontal bars indicating the median predicted abundance. Note that only assemblages containing more than two replicates and species are shown in **d** and **f**; see the Supplementary Information for the complete results.

suggesting weak or non-existent herbivore-mediated higher-order effects (Fig. 3b). In the protist system, the species-by-species coefficients exhibit consistent signs and magnitudes across temperatures, with many of these coefficients varying smoothly (Fig. 4).

These results demonstrate the type of ecological insight that can be gained, even without making mechanistic interpretations about the parameters. Moreover, the fact that these estimates of B are internally consistent—for example, varying smoothly across changing temperatures—highlights that this method is not overfitting the data, in which case we would expect our estimates to fluctuate dramatically with small changes in the system.

Simulations. To explore the robustness and generality of these findings, we simulated endpoint data from a variety of nonlinear, non-equilibrium and non-pairwise dynamical models, including Lotka–Volterra dynamics with limit cycles, competition with Allee effects[15], facultative mutualism with saturation[16], consumption with saturation[17] and competition with high-order interactions[18] (Supplementary Information).

In all cases, our approach accurately recovers the model endpoints (Extended Data Figs. 1–6), demonstrating that, despite the complex and nonlinear dynamics, the endpoint structures of these systems are approximately additive and linear (as in equation (1)). The method also successfully predicts which combinations of species will be unable to coexist. For example, in the model with mutualism and saturation, our method correctly identifies which assemblages cannot coexist (Extended Data Figs. 4 and 5), and, for the Allee effect model, it correctly identifies

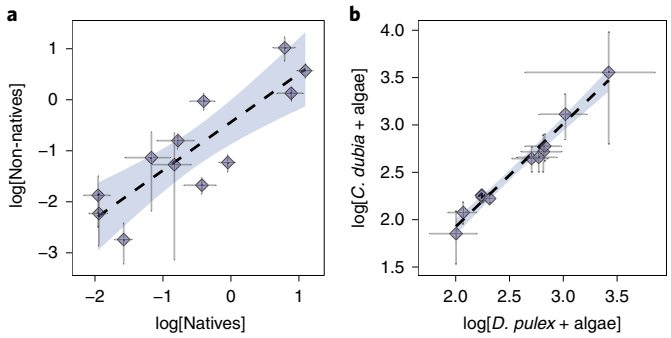

Fig. 3 | Comparison of *B* within the plant and herbivore–algae datasets. **a**, A plot of the log-transformed coefficients for the native versus non-native plant systems, paired by family into Asteraceae, Fabaceae, Lamiaceae and Poaceae, revealing that the effect of a species from one family on a species from another are conserved across the two systems. **b**, A plot of the log-transformed algae-by-algae coefficients for the herbivore–algae system, estimated for the communities containing *D. pulex* versus those containing *C. dubia*. The coefficients are strongly consistent across both herbivore systems, suggesting minimal herbivore-induced changes in algae-by-algae interactions (that is, negligible higher-order effects on the endpoints). The internal consistency in the estimated *B* for each system in **a** and **b** highlights that the method is not overfitting the data, and illustrates how this approach can offer basic ecological insight. The dashed lines give the mean regression trends, with the blue shaded regions showing 95% confidence bands. The horizontal and vertical bars show the posterior 95% prediction interval for each coefficient. The log transformations include the addition of an offset to ensure that all coefficients are positive; see the Supplementary Information for details.

the unstable fixed points even though they are never observed (Extended Data Fig. 3).

We intentionally selected dynamical systems that can exhibit oscillatory dynamics or alternative states—features that are difficult to reconcile with the assumption of a single endpoint (equation (1)). In these cases, our approach accurately estimates the average of the distribution of abundances (for example, the centroid of the oscillations; Extended Data Figs. 2 and 5). The usefulness of this type of estimate will depend on the context. In many ecological applications, we do not need to know the precise dynamics of a community; all we want to know is whether a given set of species can coexist (that is, if the centroid is far enough from zero to avoid stochastic extinctions), or what species' average abundances will be across a landscape (for example, for diversity-function or conservation questions). In such settings, this method can offer insight even if the centre of the oscillations is never observed experimentally.

Our method struggled most with systems containing strong higher-order interactions, where the effect of one species on another varies depending on the abundance of a third species (Extended Data Fig. 6). These nonlinear relationships, however, can easily be incorporated into our method by including additional terms in equation (1), such as interactions between pairs of species, $\sum_j \sum_m \tau_{ijm} z_j^{(k)} z_m^{(k)}$, exactly as in standard polynomial regression models. In this way, exploring more complex models does not require more complex experimental designs or substantially more data[19,20]. Thus, in settings where the baseline model performs poorly, different choices for equation (1) can be tested, allowing researchers to adjust the study design accordingly and benchmark the model formulations using out-of-fit or *k*-fold cross validation.

Study design. This approach suggests new study designs for parameterizing community models and building large experimental systems. Because multispecies endpoints contribute to multiple equations used to estimate *B*, diverse communities contain more information about *B* than do small communities (for example, the endpoint for the three-species communities in Fig. 1b (triangles) appears in all three systems of equations shown in Fig. 1c). Leveraging this fact, our method provides an approach to experimental design that scales linearly with the size of the species pool.

Rather than growing species in monocultures and in all pairwise combinations[7,8,21–23], one could first grow the full community of n species along with all leave-one-out communities comprising $n-1$ species each. Using this approach, it is possible to estimate *B* using only $n+1$ experiments (as opposed to $n^2/2$). This design, however, is not robust in practice: these large communities contain little information about species-poor assemblages, potentially leading to substantial prediction error for smaller assemblages. Yet, by also measuring a selection of species-poor assemblages—for example, the n monocultures—one can efficiently anchor the fitted hyperplanes, providing high quality of fit using only $2n+1$ experiments.

To examine the practical impact of experimental design, we fit the model for both plant systems using six assemblages to predict the remaining eight endpoints out-of-fit. There are 15 possible designs (combinations of assemblages) that make use of 6 endpoints and allow for the parameterization of the model; in Fig. 5 we show the best- and worst-fitting designs for the native plant pool. The prediction accuracy shows a wide range, highlighting the importance of selecting a robust design. However, the best six-assemblage design has a goodness of fit (R^2) on par with that obtained using the full dataset (Fig. 2), despite using less than half of the endpoints. The same qualitative outcome is found for the non-native plant community. Increasing the number of endpoints used to fit the native and non-native plant systems (from 6 to 13) reinforces this pattern: the best designs for any number of in-fit assemblages fare almost as well as the design using all experiments (Supplementary Information). However, across nearly 3,000 *k*-fold cross-validation assemblages, there is substantial variance in goodness of fit, with many designs performing well and a few performing very poorly.

Using simulated endpoint data, we show that designs using a mix of species-rich and species-poor communities fare better than those using only small or large communities to fit the data (Extended Data Fig. 7). In particular, the design using all monocultures and all pairs of species performs poorly relative to random designs with the same number of experiments. In contrast, the design using all monocultures and all leave-one-out communities is among the best designs for any number of experiments.

In practice, an experimental design may also fail because some assemblages collapse to smaller communities, leading to fewer unique assemblages than desired. To address this challenge, we propose a simple, iterative scheme for experimental design. First, one conducts a minimal set of experiments sufficient to obtain a draft estimate of *B*. This matrix is then used to predict the outcomes of all unperformed experiments. The assemblages with the highest inferred probability of coexistence are selected for a second round of experiments, increasing the chances that these new communities yield useful data. One might repeat this process several times, updating the estimate of *B* after each iteration. At each stage, it is possible to compute the out-of-fit accuracy of the previous *B*, giving a real-time measure of the model performance. If the quality of fit remains poor even after several iterations (or, for example, after annual sampling points in a biodiversity–ecosystem function experiment), the model assumptions can be adjusted, such as through the addition of higher-order interactions or quadratic terms in equation (1). By iteratively updating *B*, one maximizes the utility of each round of experiments, avoiding wasted experiments while ensuring adequate model fit. This approach can help experimentalists navigate the enormous space of possible assemblages, providing an efficient way to explore, build and quantify large experimental systems.

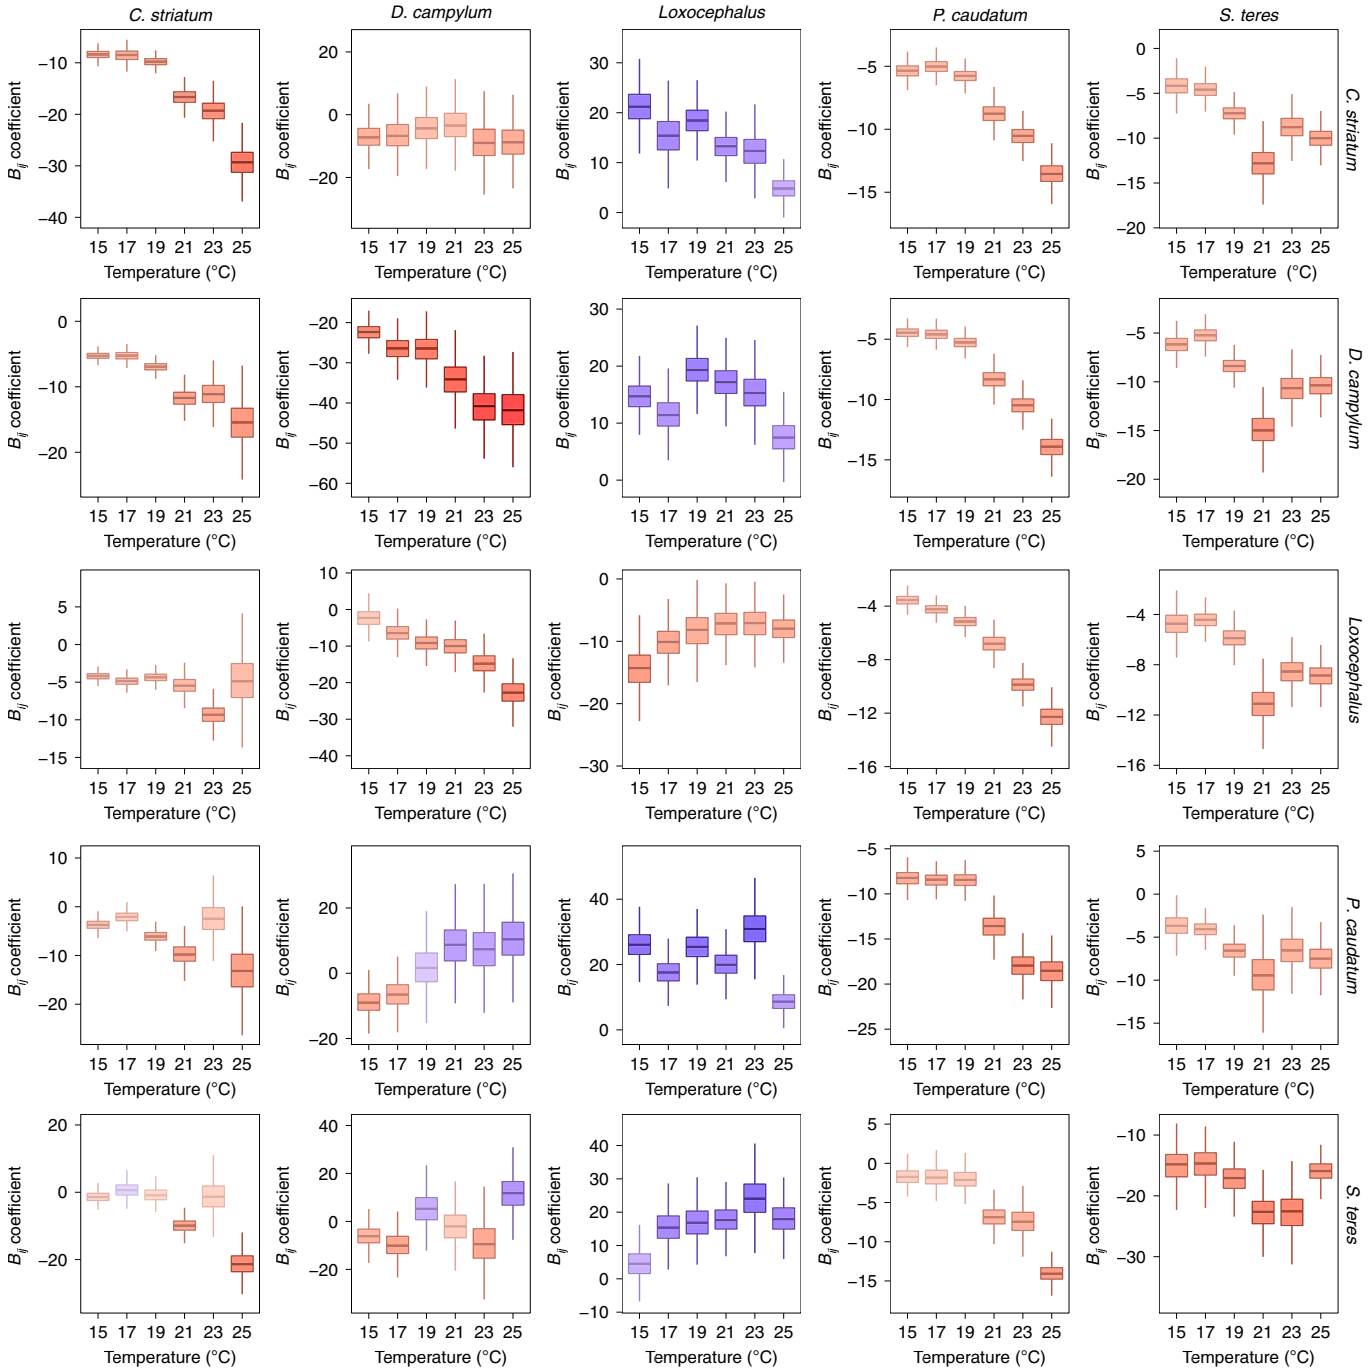

Fig. 4 | Comparison of *B* in the protist system across the six temperatures. Each panel depicts the change in the respective B_{ij} coefficient across temperature treatments, organized with B_{11} in the upper left and B_{55} in the lower right. The box plots in each panel report the medians and interquartile ranges for the posterior distribution of each coefficient at each temperature. The colours reflect the median coefficient value, with blue indicating positive effects and red indicating negative effects. Our method reveals generally consistent changes in these coefficients across the temperature gradient, further highlighting that this method is not overfitting the data and is identifying a biologically reasonable pattern. Some coefficients exhibit nonlinear fluctuations—such as the effect of *Loxocephalus* sp. on the endpoint abundance of *P. caudatum*—but these parameters generally exhibit consistent signs, similar magnitudes and overlapping prediction intervals, indicative of statistical noise. Nevertheless, because our method is statistical in nature, we caution against any mechanistic interpretation of these coefficients; they are depicted here to illustrate the consistency of the model fit across different experimental conditions.

Discussion and conclusions

From a small set of observed experimental assemblages, we show how to parameterize a model that can be used to predict the outcomes of all possible assemblages. We have successfully applied this approach to three independent experimental systems, obtaining high-quality predictions for out-of-fit data, despite the fact that our method completely neglects functional responses, behavioural changes, spatial effects, resource depletion and so on—all of which have been documented in these systems to varying degrees[24–28]. The simulation results further illustrate that nonlinear and non-equilibrium

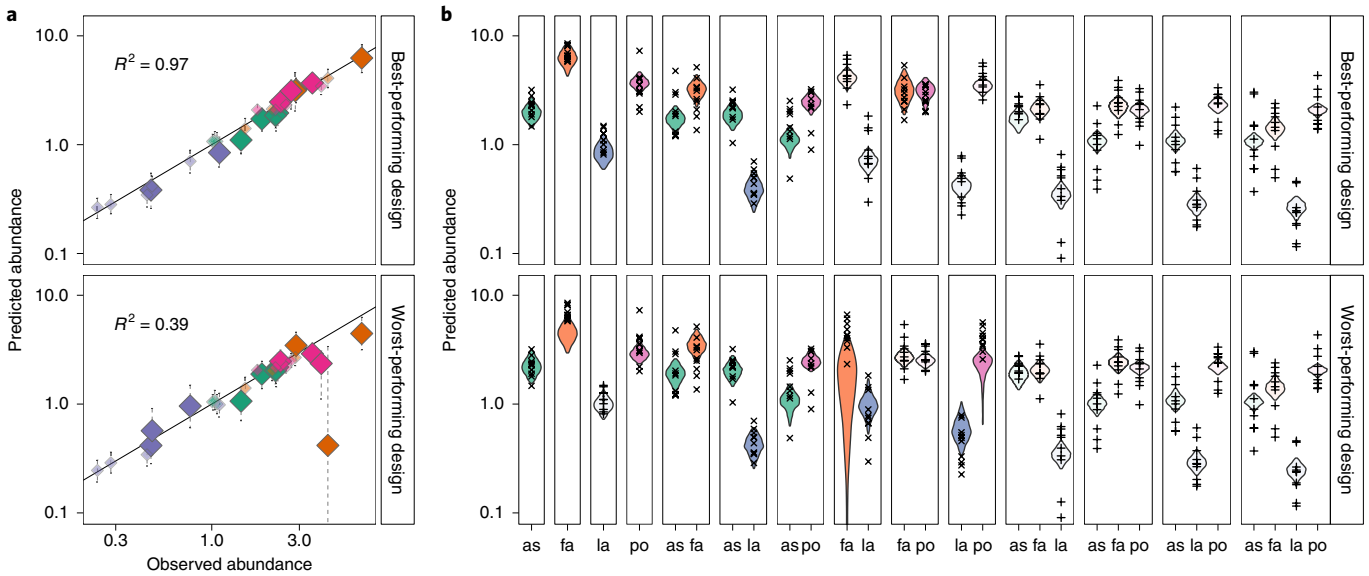

Fig. 5 | Predicting multiple endpoints out-of-fit. a,**b**, The best-performing (top) and worst-performing (bottom) experimental designs using six assemblages to fit the model and predict the remaining eight assemblages. **a**, The observed versus predicted abundances for each experimental system, with vertical error bars denoting the 95% prediction intervals of the posterior distributions. The points used to fit the model are in lighter shades and a smaller size; larger points with darker colours represent out-of-fit predictions. R^2 is calculated using only out-of-fit data. **b**, Predicted abundances for each assemblage. The unshaded violin plots show the distributions of the predictions for in-fit data used to fit the model (plus symbols); the violin plots with solid colours show the distributions of the out-of-fit data (cross symbols). The choice of which assemblages to use to fit the model (that is, the experimental design) is crucial, as made clear by the difference in quality of fit between the worst (bottom) and best (top) designs making use of six assemblages. Similar results are found for the non-native plants, and when more assemblages are used to fit the model (Supplementary Figs. 24–26). For larger communities, an efficient design is the choice of all monocultures, all leave-one-out communities and the full community (Extended Data Fig. 7). The *x*-axis labels denote different plant families: as, Asteraceae; fa, Fabaceae; la, Lamiaceae; po, Poaceae.

dynamics need not produce nonlinear endpoint structures. Together, these results demonstrate that a simple additive model and a small set of experiments can provide robust out-of-fit predictions for complex dynamical systems.

In applying this approach, an important practical consideration is deciding when to end an experiment and sample the species abundances. In some cases, there may be a biologically relevant time point (for example, the end of the growing season for annual plants). In other settings, one might use some inexpensive, non-invasive technique to monitor the state of the community (for example, optical density, chlorophyll fluorescence or total respiration) and track this measure until it converges on a stationary distribution, as was done here for the protist experiment (Supplementary Information). In general, our results suggest that one simply needs to ensure that the dynamics have passed through any transients, allowing species sufficient time to go extinct. In the protist system, for example, the model fit was poor between days 1 and 10, but increased sharply in quality at day 11 and remained good for all days between 11 and 30 (Extended Data Fig. 8 and Supplementary Information), highlighting that this method is not overly sensitive to the choice of when to end the experiment. In systems where the endpoints cannot be observed, or where there is no principled way to determine when to end the experiments—as in communities that cannot be easily manipulated or systems with dynamics that play out over very long periods—alternative methods might be more suitable.

A key benefit of this approach is that the minimum number of required experiments scales linearly with the size of the species pool, helping to overcome a central challenge in studying large experimental systems. In theory, this method becomes increasingly efficient as the size of the species pool increases, with the minimum sufficient proportion of experiments scaling as $(n+1)/(2^n-1)$. In practice, however, large communities may present computational challenges, as the number of parameters to be estimated still scales with n^2. In large experimental systems, it is also likely that some species will persist in fewer than n unique endpoints, preventing the estimation of their associated parameters. Invoking strict Bayesian priors or employing more advanced search algorithms may help overcome some of these limitations. Nevertheless, improving the computational efficiency of this approach is a challenge, and one that would help extend this method to highly speciose systems.

Because this method requires relatively few experiments, it can easily be integrated into traditional studies (for example, biodiversity–ecosystem function studies) without requiring a substantial overhaul of the experimental design. That is, we envision that this approach can be used to complement—rather than replace—current experimental methods. However, to ensure robust predictions, we emphasize that it may be necessary to use higher replication in field-based experiments to overcome stochasticity and experimental noise. Indeed, when external factors (including, especially, dispersal and recruitment) cannot be adequately controlled, this approach may not be appropriate. However, one advantage of our method is that it provides a straightforward way to assess performance using out-of-fit predictions. As this approach permits very efficient experimental designs, high replication may also be more feasible.

Central to our method is that it uses a single measurement for each assemblage. By focusing exclusively on the abundances of species present in a final snapshot, this approach requires drastically fewer measurements than time-series methods. Yet this method also completely ignores initial conditions: two experiments that start with different species compositions but collapse to the same assemblage are considered to be replicates of this endpoint. Extending the model to account for initial conditions is an important next step, as doing so should help improve fit and prediction by

incorporating knowledge about which assemblages have collapsed due to extinctions.

While many challenges remain in the study of speciose ecological communities, our framework provides a powerful approach for exploring these systems. We adopt a simplified, statistical method that robustly predicts the outcomes of unobserved experiments using relatively few observations. By forgoing dynamical modelling, we gain tractability, thereby providing a principled and efficient means of studying, navigating and building diverse experimental systems.

Methods

Experimental setting. We consider a pool of n species, which can give rise to as many as 2^n-1 distinct combinations of species' presence and absence. For any combination (henceforth, assemblage), we consider an experiment in which each of the selected species is inoculated at some initial density in controlled conditions, and the dynamics of the system (that is, inter- and intraspecific interactions) are allowed to play out. Once sufficient time has elapsed, the abundance of every species is measured. As explained in the main text, several methods could inform us of when to sample the abundances, and in the Supplementary Information we show the consequences of sampling during the transient phase. Here, we simply assume that the system has settled in some dynamical attractor, including the possibility of stochastic fluctuations around an equilibrium, deterministic chaos, periodic oscillations and so on. This experimental design permits destructive sampling methods (for example, high-throughput sequencing of the endophytic microbial assemblage in plants), and does not require any observation of the transient dynamics or initial conditions.

We refer to the set of abundance measurements for each assemblage as an endpoint of the dynamics, denoted by $x^{(k)}$. Note that there is not a bijective mapping between assemblages and endpoints in general. For example, a particular assemblage of species may not coexist, in which case the system will collapse to a subset of species and reach the same endpoint as if it had been seeded with only the subassemblage. Conversely, identical initial assemblages may result in different endpoints. Here, we distinguish between two cases: first, the endpoints might be sampled from the same attractor, but may differ because of cycling or sampling error; second, we might have true multistability[29]—the replicate systems have reached distinct attractors, possibly depending on the initial abundance of each species. In this work, we do not explicitly account for the latter case, assuming that replicate endpoints are drawn from the same stationary distribution (as such, multistability is conflated with sampling error). Instances of true multistability in ecological dynamics are well documented, but should be identifiable in data with sufficient replication[30–32].

Given a set of experimental endpoints, we attempt to predict which unobserved assemblages can coexist and at what endpoint abundances. To accomplish this, we assume that the endpoints are related by a simple linear model. As noted in Box 1, this model takes the form

$$z_i^{(k)} = \gamma_i + \sum_{j\neq i} \tau_{ij} z_j^{(k)} \quad k \in \{1, \ldots, 2^n-1\} \quad \forall i \text{ such that } z_i^{(k)} \neq 0 \tag{3}$$

where z_i denotes the average (across replicates) endpoint abundance of species i. In cases where the stationary distribution is not a fixed point, $\mathbf{z}^{(\mathbf{k})} = (z_1^{(k)}, z_2^{(k)}, \ldots)$ can be thought of as the centroid of the attractor for k. The coefficients γ and τ are statistically determined constants that relate the endpoint abundances to each other. Intuitively, γ_i models the average abundance of i when grown in isolation, and therefore we expect γ_i to be positive for producers—reflecting their carrying capacity—and zero (or negative) for consumers and predators. The coefficient τ_{ij} can be seen as the average per capita effect that members of j have on the endpoint abundance of i. We assume each τ_{ij} is constant across communities (that is, τ_{ij} has no dependence on k).

Manipulating equation (3) slightly, we can obtain

$$0 = 1 - \frac{1}{\gamma_i} z_i^{(k)} + \sum_{j\neq i} \frac{\tau_{ij}}{\gamma_i} z_j^{(k)} \tag{4}$$

and, letting $B_{ij} = \frac{\tau_{ij}}{\gamma_i}$ for $i \neq j$ and $B_{ii} = -\frac{1}{\gamma_i}$:

$$-1 = \sum_j B_{ij} z_j^{(k)} \tag{5}$$

Equation (5) can be written more compactly in matrix form as

$$-1 = B^{(k)} z^{(k)} \tag{6}$$

Here, if the kth endpoint contains w species, then the left side of equation (6) is a $w \times 1$ vector where every element in the vector is the number 1, and $B^{(k)}$ is the $w \times w$ submatrix obtained by selecting only the elements of B (the full $n \times n$ matrix of coefficients) whose rows and columns correspond to species that have non-zero density in endpoint k.

Having estimated B, it is possible to predict any of the 2^n-1 endpoints by solving equation (6) for a desired k. If the estimated solution, $\hat{z}^{(k)}$, contains any negative elements, we interpret this as an indication that the species in k cannot coexist. If every element of $\hat{z}^{(k)}$ is positive, these elements give the endpoint abundances at which the species may coexist. This condition (non-negativity, or feasibility) is necessary for the coexistence of k, but not sufficient, because the endpoint (or the attractor associated with it) may not be attractive or stable.

Inferring B from a set of endpoints. We are primarily concerned with the practical challenge of inferring B from endpoint data. As a first approach, we form equations relating the elements of B and the endpoint abundances, and use these to solve for B on a row-by-row (species-by-species) basis. Although this 'naive' procedure ignores the complex error structure of the model (Supplementary Information), it is illustrative of the general approach, helping to provide intuition.

To implement this first approach, we introduce the matrices E_i, which contain the observed $x^{(k)}$ for any $k \in \{1, \ldots, 2^n-1\}$ such that $x_i^{(k)} \neq 0$. That is, each E_i contains all the endpoints in which i is present. Every E_i has n columns, corresponding to the n species, and we fill in a zero wherever a species is not present in an endpoint. The number of rows of E_i will be variable, depending on the dataset.

As a simple example, consider a case in which we have a pool of three species, and species 1 is present in five endpoints: two endpoints containing only species 1 (monocultures), two endpoints in which species 1 and 2 coexist and one in which species 1 and 3 coexist. The structure of matrix E_1 would be:

$$E_1 = \begin{pmatrix} x_1^{(1)} & 0 & 0 \\ x_1^{(1\prime)} & 0 & 0 \\ x_1^{(2)} & x_2^{(2)} & 0 \\ x_1^{(2\prime)} & x_2^{(2\prime)} & 0 \\ x_1^{(3)} & 0 & x_3^{(3)} \end{pmatrix} \tag{7}$$

where $x^{(k')}$ is a replicated endpoint containing the same set of species as $x^{(k)}$. Of course, the dataset may contain other endpoints in which i is not present; these endpoints will appear in E_2 or E_3, but not E_1. We also highlight the fact that some of these rows (endpoints) will be repeated in other E_i. For example, the third row of E_1 will also be present in E_2, and the last row of E_1 will also appear in E_3. This means that endpoints containing multiple species provide more information than those containing single species, allowing for an efficient experimental design (Extended Data Fig. 7 and Supplementary Information).

Having formed the matrices E_i, we can recover the ith row of B, denoted B_i, by solving

$$E_i B_i^t = -1 \tag{8}$$

where the right side of equation (8) is a vector of −1s with as many elements as there are endpoints in E_i. A general solution to this equation is given by $B_i^t = -E_i^+ 1$, where E_i^+ is the Moore–Penrose pseudoinverse of E_i. We write the pseudoinverse rather than the inverse because E_i need not be square. The use of the pseudoinverse is also convenient because, when analysing actual experimental data, replicate experiments may yield different abundances for the same endpoint, due to measurement errors or non-point attractors. With the notation introduced above, one can simply list all experimental endpoints, including replicates, in the corresponding E_i, and solve. In this case, there will not be an exact solution, because the system is overdetermined, but the Moore–Penrose pseudoinverse guarantees that the matrix recovered is a maximum likelihood (least-squares) estimate of B given the data.

An example may help clarify this approach. Consider a series of endpoints generated by adding a small amount of noise to solutions of equation (5), with

$$B = \begin{pmatrix} -2 & -\frac{3}{4} & -\frac{1}{4} \\ -\frac{3}{4} & -2 & -\frac{1}{4} \\ -\frac{1}{4} & -\frac{1}{2} & -2 \end{pmatrix} \tag{9}$$

Collecting the endpoints and constructing E_1 for these data gives

$$E_1 = \begin{pmatrix} 0.55 & 0 & 0 \\ 0.48 & 0 & 0 \\ 0.37 & 0.33 & 0 \\ 0.34 & 0.37 & 0 \\ 0.44 & 0 & 0.45 \\ 0.35 & 0.32 & 0.38 \\ 0.33 & 0.32 & 0.35 \end{pmatrix} \tag{10}$$

and we recover an estimate of B_1 by computing

$$\begin{pmatrix}\tilde{B}_{11}\\ \tilde{B}_{12}\\ \tilde{B}_{13}\end{pmatrix} = \begin{pmatrix}0.55 & 0 & 0\\ 0.48 & 0 & 0\\ 0.37 & 0.33 & 0\\ 0.34 & 0.37 & 0\\ 0.44 & 0 & 0.45\\ 0.35 & 0.32 & 0.38\\ 0.33 & 0.32 & 0.35\end{pmatrix}^{+}\begin{pmatrix}-1\\ -1\\ -1\\ -1\\ -1\\ -1\\ -1\end{pmatrix}$$

$$= \begin{pmatrix}0.904 & 0.789 & 0.213 & 0.116 & 0.287 & -0.175 & -0.179\\ -0.658 & -0.574 & 0.832 & 1.022 & -0.879 & 0.519 & 0.566\\ -0.532 & -0.465 & -0.617 & -0.619 & 1.123 & 0.718 & 0.634\end{pmatrix}\begin{pmatrix}-1\\ -1\\ -1\\ -1\\ -1\\ -1\\ -1\end{pmatrix}$$

$$= \begin{pmatrix}-1.95\\ -0.82\\ -0.24\end{pmatrix}$$

which closely approximates the true coefficients (−2, −0.74 and −0.25). An identical procedure yields estimates for the remaining rows of B.

Requirements. To solve for the coefficients in a given row, B_i, it is necessary that the (minimal) structural rank of E_i be n, meaning that the rank must remain n even if all of the non-zero values of E_i were made identical[33]. This condition is equivalent to simultaneously imposing three biologically meaningful requirements: (1) each species must be present in at least n distinct endpoints, not counting replicates; (2) each species must co-occur with each other species in at least one endpoint (that is, for every pair of species i and j, there must be some endpoint where i and j co-occur, possibly along with other species); and (3) for each i there must exist a perfect matching between the n species and the endpoints in which they co-occur with i. Put another way, for a focal species i, each endpoint can only count once towards the second condition.

These conditions place obvious limitations on the datasets and ecological systems for which our method is applicable. Coexistence among species must be reasonably widespread for the first and second conditions to hold. In particular, if a given pair of species i and j never co-occur, it is impossible to estimate the B_{ij} relating their endpoint abundances. Systems with substantial trophic structure are unlikely to satisfy these conditions. For example, in a linear food chain, the top consumer occurs in only a single endpoint (the full assemblage), violating the first condition. Similarly, systems with strong competitive hierarchies will usually violate these conditions. In general, we envision our method applied to systems where most species are able to persist in isolation, and where the majority of interactions are relatively weak (for example, many plant and microbial communities).

Finally, we note that our model closely resembles a linear regression on the endpoint abundances. In fact, exactly as in linear regression, if we assume (1) that the endpoints are measured without error and (2) that the values $y_i = E_i B_i^t + 1$ are independently, normally distributed with mean 0 and variance σ^2, then by taking the pseudoinverse we minimize $\left\|E_i B_i^t + 1\right\|_2$ and therefore the variance σ^2. The first of these assumptions is clearly incompatible with the earlier assumption that replicate endpoints are noisy samples from the same attractor. This issue is partially reconciled by imposing the structural rank condition explained above; however, the error structure of our model is still quite distinct from that of a typical linear regression. The fact that these assumptions are not fully appropriate for our setting motivates the development of more sophisticated approaches, which are explained below and employed in practice.

Accounting for the error structure. The naive regression approach illustrated above is straightforward, but has several drawbacks. In particular, empirical data will always have some degree of error in the measurements of the densities; moreover, each multispecies endpoint is present in multiple matrices (in the example above, $x^{(3)}$ would be reported in E_1 and E_3), and therefore these equations are coupled. Lastly, the maximum likelihood approach above attempts to find the best-fitting B that yields an approximate solution, that is, such that residuals are independent and normally distributed. In doing so, it can allow species' true endpoints to be quite far from the observed values, such that species may be observed to be present in an endpoint but have a predicted endpoint abundance that is negative.

Thus, although the method outlined here can be naively solved using simple linear regression, there is no guarantee that the result is accurate, and it can allow for results that are inconsistent with the biology of the system (for example, species "coexisting" at negative abundances). We therefore must use a method that allows us to find a matrix B such that the corresponding set of true endpoints $z^{(k)} = -(B^{(k)})^{-1}1$ are as close as possible to the observed $x^{(k)}$ across all replicates and communities k.

We present two complementary approaches for estimating B, both of which account for the complex error structure and prevent species from coexisting at negative abundances.

First, as detailed in the Supplementary Information, one could use a sum-of-squares approach, which uses numerical optimization to minimize the deviation between the observed and predicted endpoints to get a single estimate of B. The benefit of this approach is that it correctly handles the error structure by explicitly incorporating measurement error. It can also be computationally faster than the Bayesian approach, which is detailed next. A drawback to the sum-of-squares approach, however, is that it may struggle to find a global maximum, particularly if the likelihood surface for B is relatively flat or has many local maxima—a fact that is made more complex by the need to invert B (or its submatrices) to calculate the endpoint abundances. This method also does not provide a measure of the uncertainty or standard error surrounding the coefficients, complicating model-selection approaches and preventing one from estimating confidence intervals for the resulting abundance predictions $\hat{z}^{(k)}$. Nevertheless, because of the computational efficiency of this method, we employ it for the analysis of experimental designs for the plant systems and for the simulations, where we are primarily interested in goodness of fit of the median value rather than measures of uncertainty.

To address the limitations of the sum-of-squares approach, the most rigorous method for handling the complex error structure—and the one used to fit the three empirical datasets—is a Bayesian Markov chain Monte Carlo (MCMC) approach, as detailed next. This method allows for probabilistic inference about coexistence, while appropriately handling the complex error structure.

Measurement error model—a Bayesian approach. The Bayesian MCMC approach is similar to the sum-of-squares approach, but it yields posterior distributions for the elements of B and for each $z^{(k)}$. It allows for standard Bayesian model-selection approaches, and can more easily handle relatively flat likelihood surfaces, provided the priors are chosen appropriately.

To implement this method, we assume there is some underlying B that gives rise to the true endpoints, given by $z^{(k)} = -(B^{(k)})^{-1}1$. The observed $x^{(k)}$ are viewed as random variables, sampled with error from some distribution centred at $z^{(k)}$ with a vector of standard deviations $\sigma^{(k)}$.

The basics of this approach are as follows:

1. Assign prior distributions for the coefficients of $B = \{B_{ij}\}$ and $\sigma = (\sigma_1, \ldots, \sigma_n)'$, with the hyperparameters for these distributions encoded in the vector $\boldsymbol{\alpha}$. For the sake of generality, we assume a species-specific standard deviation, but one could assume a constant standard deviation across all species. Alternatively, one could make more complex assumptions about both B and σ (for example, that they vary smoothly across environments or that they are phylogenetically correlated across species).
2. Sample $\hat{B}$ and $\hat{\sigma}$ from these distributions, and predict the equilibrium abundance for each observed k by calculating $\hat{z}^{(k)} = -(\hat{B}^{(k)})^{-1}1$.
3. If any species is predicted to have a negative abundance for an endpoint where it was observed to be present, replace it with an arbitrarily small positive value (for example, 10^{-20}) to ensure this outcome is assigned a near-zero probability of occurrence under a log-normal error structure. Without this step, the model would produce unbiological results, whereby species coexist at negative abundances. This step can be modified based on the specific error structure (for example, if assuming the errors are normally distributed rather than log-normal).
4. Calculate the logarithm of the posterior probability for $\hat{B}$ and $\hat{\sigma}$ by summing the log-probabilities across all endpoints and replicates, following Bayes' theorem:

$$\log P(\hat{B}, \hat{\sigma} | x^{(1)}, \ldots, x^{(k)}, \alpha) \propto \sum_k \log P(x^{(k)} | z^{(k)}, \hat{\sigma}) + \log P(\hat{B}, \hat{\sigma} | \alpha)$$

   The specific error structure for the endpoints is encoded in the first probability term. Thus, to include a log-normal error structure for $x^{(k)}$, as we do for the datasets below, we set $P(x^{(k)} | z^{(k)}, \hat{\sigma})$ to be the density function of a log-normal distribution with parameters $z^{(k)}$ and $\hat{\sigma}$.
5. To obtain a posterior distribution for B, repeat this process via MCMC sampling.

For the datasets analysed here, we use the Stan programming language (rstan v.2.19.2) to implement the MCMC sampling, called via the stan function in R (v.3.6.1) with the default No-U-Turn variant of the Hamiltonian Monte Carlo algorithm[34,35]. For each dataset, we ran four separate MCMC chains for 50,000 iterations each, with a warm-up of 20,000 iterations and thinning every 40 iterations. Posterior plots were investigated to ensure proper mixing of the chains, and the adaptive-delta parameter was set to 0.85 to minimize divergent transitions. See the data-specific sections (Supplementary Information) for the exact priors and fitting details for each of the three systems.

Probabilistic inference. The results of this method are posterior distributions for $\hat{B}$, $\hat{\sigma}$ and $\hat{z}^{(k)}$, which provide estimates of certainty for each value. We can use these posteriors to conduct basic probabilistic inference. For example, for each unobserved or out-of-fit assemblage, we can infer the probability of coexistence by calculating the proportion of posterior Bs that resulted in all species in the assemblage having positive abundance. Note that this probability cannot be

estimated for in-fit data because our method requires that observed assemblages be predicted to coexist (see Step 3, above). For each experimental system, for example, we sampled 1,000 bootstrap estimates from the posterior for *B*, and calculated the probability of coexistence for each assemblage (Supplementary Information). With stricter assumptions about the dynamics of the system, one can also infer the probability of local and global stability of the fixed points, for both the observed and unobserved or out-of-fit assemblages (Extended Data Figs. 1–6).

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The data and code needed to replicate the central findings of this work are available at https://git.io/fjvON

Received: 30 April 2019; Accepted: 13 November 2019;
Published online: 16 December 2019

References

1. Gause, G. F. *The Struggle for Existence* (Williams & Wilkins Company, 1934).
2. Yoshida, T., Jones, L. E., Ellner, S. P., Fussmann, G. F. & Hairston, N. G. Jr. Rapid evolution drives ecological dynamics in a predator–prey system. *Nature* **424**, 303–306 (2003).
3. Costantino, R. F., Desharnais, R., Cushing, J. M. & Dennis, B. Chaotic dynamics in an insect population. *Science* **275**, 389–391 (1997).
4. Dai, L., Vorselen, D., Korolev, K. S. & Gore, J. Generic indicators for loss of resilience before a tipping point leading to population collapse. *Science* **336**, 1175–1177 (2012).
5. Cadotte, M. W. Experimental evidence that evolutionarily diverse assemblages result in higher productivity. *Proc. Natl Acad. Sci. USA* **110**, 8996–9000 (2013).
6. Heemsbergen, D. A. et al. Biodiversity effects on soil processes explained by interspecific functional dissimilarity. *Science* **306**, 1019–1020 (2004).
7. Friedman, J., Higgins, L. M. & Gore, J. Community structure follows simple assembly rules in microbial microcosms. *Nat. Ecol. Evol.* **1**, 0109 (2017).
8. Dormann, C. F. & Roxburgh, S. H. Experimental evidence rejects pairwise modelling approach to coexistence in plant communities. *Proc. R. Soc. B* **272**, 1279–1285 (2005).
9. Bell, T., Newman, J. A., Silverman, B. W., Turner, S. L. & Lilley, A. K. The contribution of species richness and composition to bacterial services. *Nature* **436**, 1157–1160 (2005).
10. Tilman, D. et al. Diversity and productivity in a long-term grassland experiment. *Science* **294**, 843–845 (2001).
11. Kuebbing, S. E., Classen, A. T., Sanders, N. J. & Simberloff, D. Above- and below-ground effects of plant diversity depend on species origin: an experimental test with multiple invaders. *New Phytol.* **208**, 727–735 (2015).
12. Rakowski, C. & Cardinale, B. J. Herbivores control effects of algal species richness on community biomass and stability in a laboratory microcosm experiment. *Oikos* **125**, 1627–1635 (2016).
13. Pennekamp, F. et al. Biodiversity increases and decreases ecosystem stability. *Nature* **563**, 109–112 (2018).
14. Pennekamp, F. et al. Dynamic species classification of microorganisms across time, abiotic and biotic environments—a sliding window approach. *PLoS ONE* **12**, e0176682 (2017).
15. Sun, G. Q. Mathematical modeling of population dynamics with Allee effect. *Nonlinear Dyn.* **85**, 1–12 (2016).
16. Holland, J. N., Okuyama, T. & DeAngelis, D. L. Comment on 'Asymmetric coevolutionary networks facilitate biodiversity maintenance'. *Science* **313**, 1887–1887 (2006).
17. Fussmann, G. F. & Heber, G. Food web complexity and chaotic population dynamics. *Ecol. Lett.* **5**, 394–401 (2002).
18. Bairey, E., Kelsic, E. D. & Kishony, R. High-order species interactions shape ecosystem diversity. *Nat. Commun.* **7**, 12285 (2016).
19. Mayfield, M. M. & Stouffer, D. B. Higher-order interactions capture unexplained complexity in diverse communities. *Nat. Ecol. Evol.* **1**, 0062 (2017).
20. Levine, J. M., Bascompte, J., Adler, P. B. & Allesina, S. Beyond pairwise mechanisms of species coexistence in complex communities. *Nature* **546**, 56–64 (2017).
21. Carrara, F., Giometto, A., Seymour, M., Rinaldo, A. & Altermatt, F. Inferring species interactions in ecological communities: a comparison of methods at different levels of complexity. *Methods Ecol. Evol.* **6**, 895–906 (2015).
22. Maynard, D. S. et al. Diversity begets diversity in competition for space. *Nat. Ecol. Evol.* **1**, 0156 (2017).
23. Kraft, N. J. B., Godoy, O. & Levine, J. M. Plant functional traits and the multidimensional nature of species coexistence. *Proc. Natl Acad. Sci. USA* **112**, 797–802 (2015).
24. Fox, J. W. The dynamics of top-down and bottom-up effects in food webs of varying prey diversity, composition, and productivity. *Oikos* **116**, 189–200 (2007).
25. Crawford, K. M. & Knight, T. M. Competition overwhelms the positive plant–soil feedback generated by an invasive plant. *Oecologia* **183**, 211–220 (2017).
26. Beveridge, O. S., Petchey, O. L. & Humphries, S. Direct and indirect effects of temperature on the population dynamics and ecosystem functioning of aquatic microbial ecosystems. *J. Anim. Ecol.* **79**, 1324–1331 (2010).
27. Golberg, D. E. Neighborhood competition in an old-field plant community. *Ecology* **68**, 1211–1223 (1987).
28. Sarnelle, O. & Wilson, A. E. Type III functional response in *Daphnia*. *Ecology* **89**, 1723–1732 (2008).
29. Xiao, Y. et al. Mapping the ecological networks of microbial communities. *Nat. Commun.* **8**, 2042 (2017).
30. Beisner, B., Haydon, D. & Cuddington, K. Alternative stable states in ecology. *Front. Ecol. Environ.* **1**, 376–382 (2003).
31. Scheffer, M. & Carpenter, S. R. Catastrophic regime shifts in ecosystems: linking theory to observation. *Trends Ecol. Evol.* **18**, 648–656 (2003).
32. Seekell, D. A., Cline, T. J., Carpenter, S. R. & Pace, M. L. Evidence of alternate attractors from a whole-ecosystem regime shift experiment. *Theor. Ecol.* **6**, 385–394 (2013).
33. Steffen, T. *Control Reconfiguration of Dynamical Systems: Linear Approaches and Structural Tests* (Springer Science & Business Media, 2005).
34. Hoffman, M. D. & Gelman, A. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593–1623 (2014).
35. Carpenter, B. et al. Stan: a probabilistic programming language. *J. Stat. Softw.* **76**, 1–32 (2017).

Acknowledgements

We thank S. Kuebbing, C. Rakowski, F. Pennekamp and O. Petchey for making their data available and for comments on earlier drafts of this manuscript. We thank C. Serván, M. Pascual, J. Bergelson, E. Baskerville, G. Barabás and E. Friedlander for assistance and suggestions throughout this study.

Author contributions

D.S.M., Z.R.M. and S.A. conceived this study and developed the methods. D.S.M. collected and analysed the data, and wrote the supplementary data analysis. S.A. and Z.R.M. wrote the Supplementary Information and Methods and implemented the simulations. All authors contributed to the writing of the manuscript and assisted with revisions.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41559-019-1059-z.

Supplementary information is available for this paper at https://doi.org/10.1038/s41559-019-1059-z.

Correspondence and requests for materials should be addressed to D.S.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

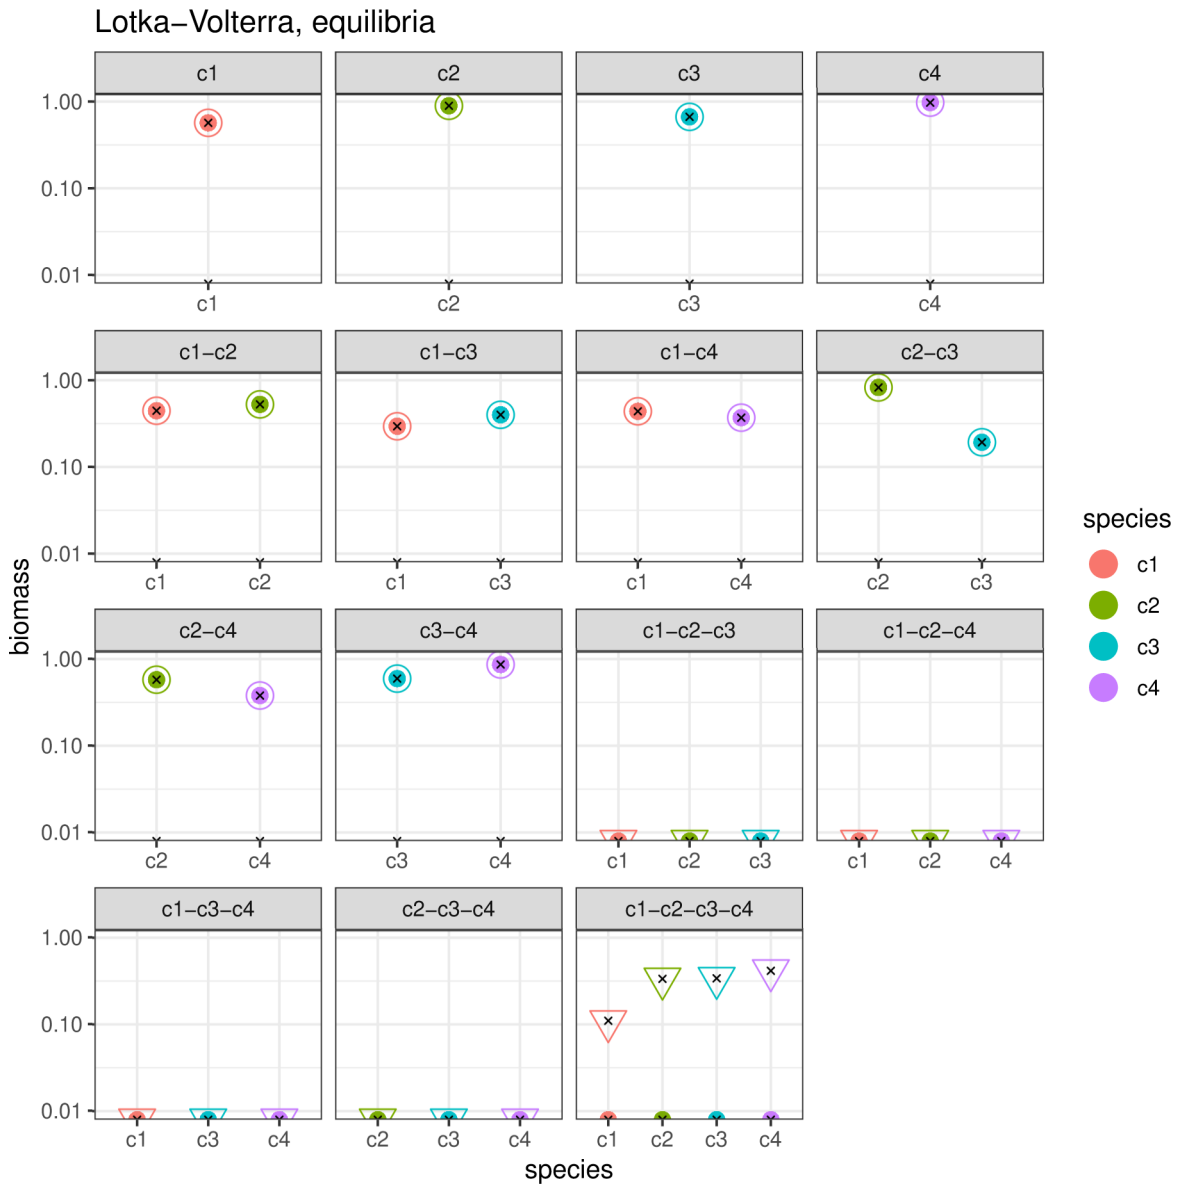

Extended Data Fig. 1 | Simulation results for a Generalized Lotka-Volterra (GLV) competitive system. For each combination of the four competitors (c1 to c4, colours), we ran 96 simulations starting from different initial conditions. Each panel shows the location of the simulations' endpoints (solid circles), as well as the true location of the equilibria for the system (crosses, computed analytically). Experiments resulting in a lack of coexistence are represented as half-points at the bottom of the graph. The predictions obtained using our method are reported using open symbols. There are two cases: we use circles for predictions of (locally) stable endpoints, and triangles for unstable ones; the stability is calculated under the assumption of Lotka-Volterra dynamics and equal growth rates. For this system, as expected, we recover a perfect fit for the positive densities. We can also correctly predict the lack of coexistence among triplets. We predict perfectly the position of the four-species equilibrium, and correctly classify it as unstable, despite having used growth rates that differ substantially from each other.

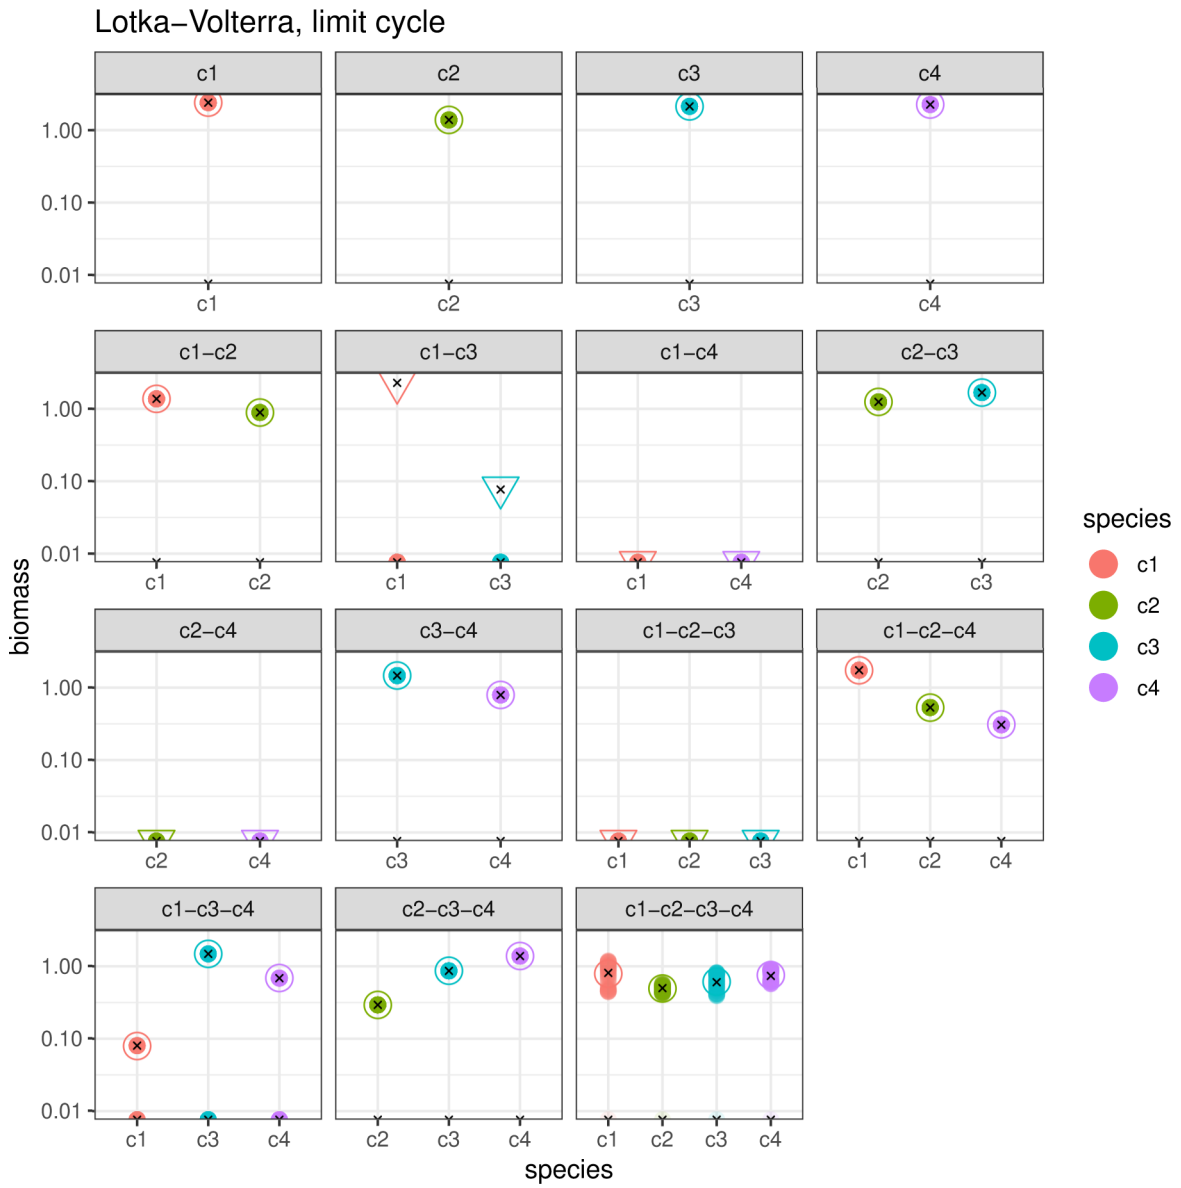

Extended Data Fig. 2 | Simulation results for a GLV competitive system exhibiting a limit cycle. Colours and symbols are as in Fig. 17. Notable features of this system are: competitors might (c1-c2, c2-c3) or might not (c1-c3, c1-c4) coexist in pairs. However, in one case a feasible but unstable equilibrium exists (c1-c3, correctly predicted by our method), while in the other there is no feasible equilibrium (c1-c4, also correctly predicted). The system including c1-c3-c4 shows dependence on initial conditions (some trajectories collapse to another system, while others converge to equilibrium), signaling a locally (but not globally) stable equilibrium. The method correctly identifies the position of the 4-species equilibrium surrounded by the limit cycle. However, it suggests stability for the equilibrium, while it must be unstable to give rise to the stable limit cycle. The misclassification stems from the fact that in the calculation of stability, we consider growth rates to be equal (because we cannot infer growth rates from endpoints), while this is not the case here.

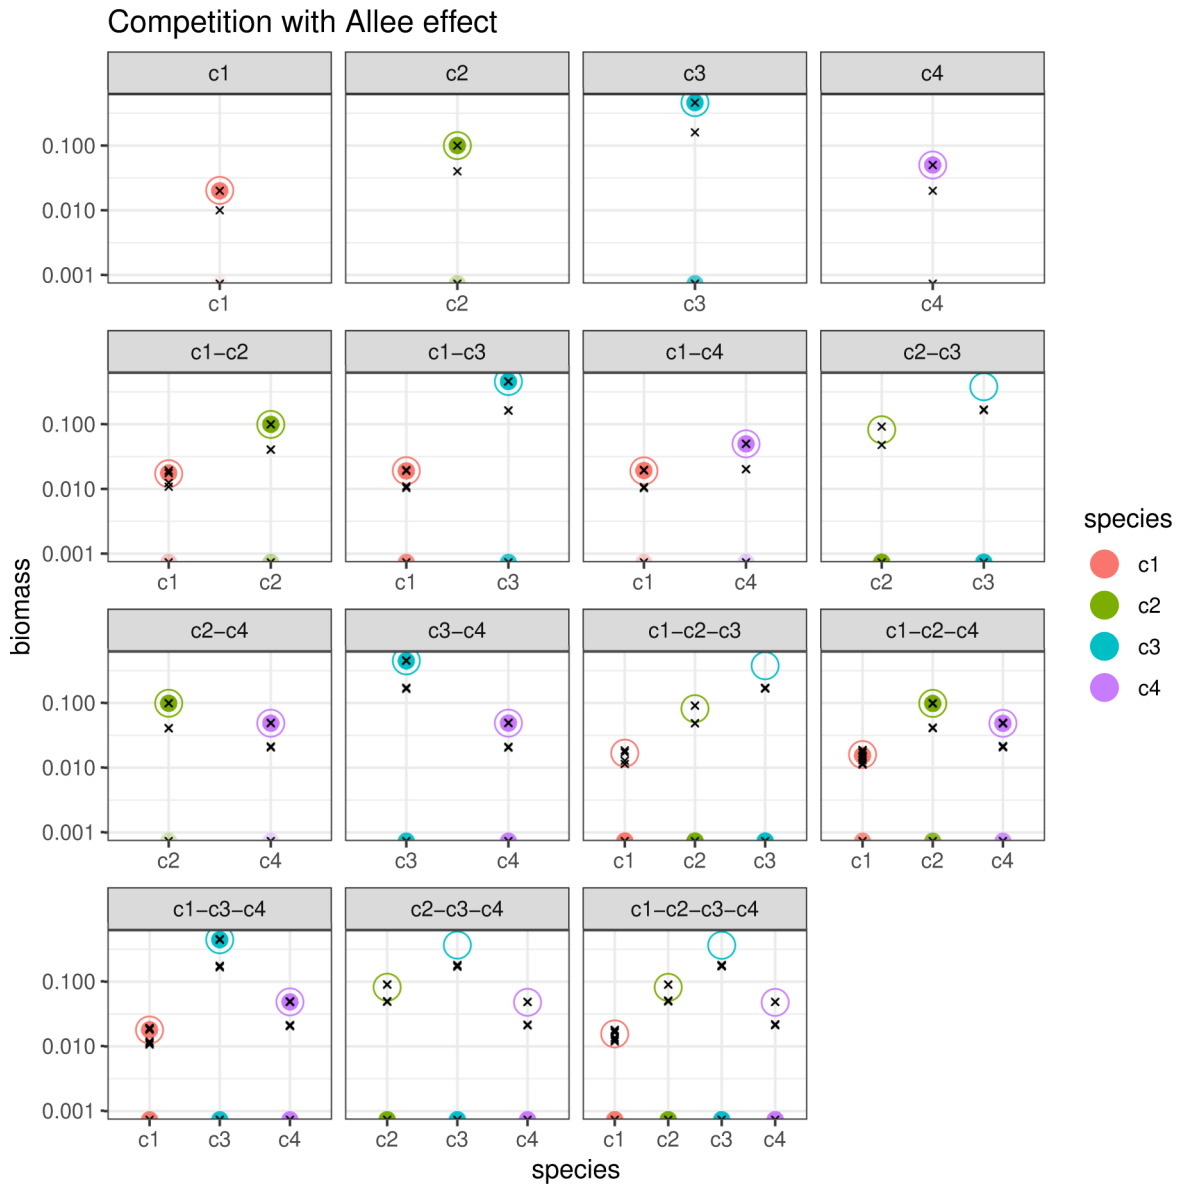

Extended Data Fig. 3 | Simulation results for competition with Allee effects, in which competitors cannot grow when rare. This system exhibits multistability in all cases (half-points at the bottom of each graph signal trajectories that resulted in extinctions). Despite the fact that the model contains cubic terms (while our method can deal only with quadratic terms), the in-fit is excellent, in all cases fitting the location of the endpoints perfectly. Experiments in which the species do not coexist are however misclassified—while there exist equilibria close to the prediction, they are unstable, rather than stable as predicted by our method.

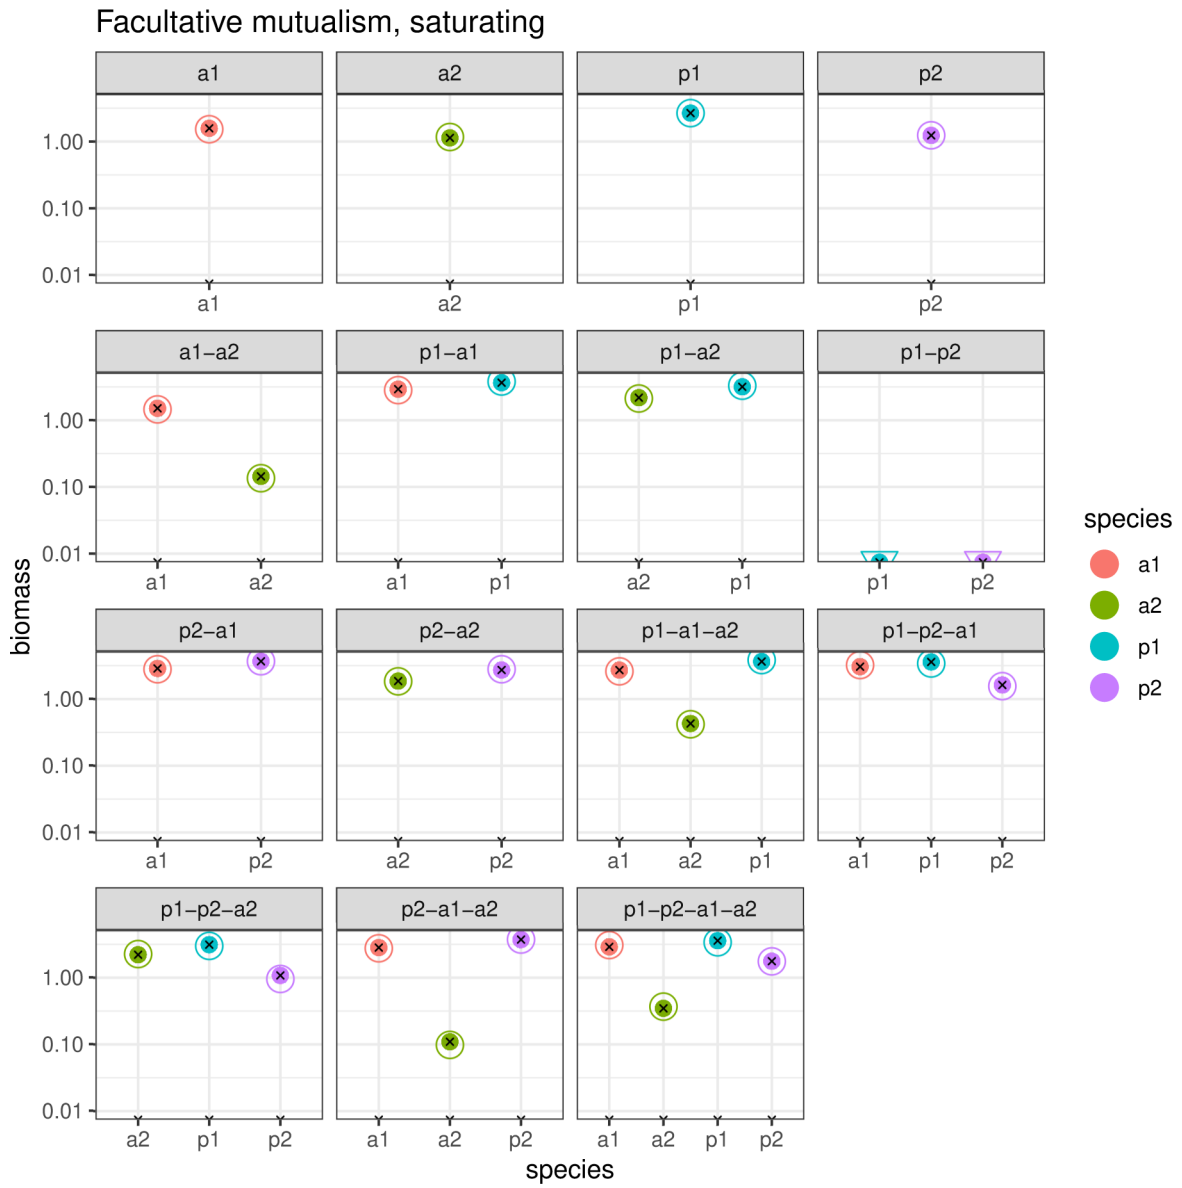

Extended Data Fig. 4 | Simulation results for a system of facultative mutualism between two classes of competitors. Plants are denoted by *p*, and animals are denoted by *a*. Despite the non-linear functional response, the method predicts the location of the endpoints (all characterized by equilibrium dynamics) almost perfectly. The method also correctly predicts the lack of coexistence between the two plants (panel p1-p2).

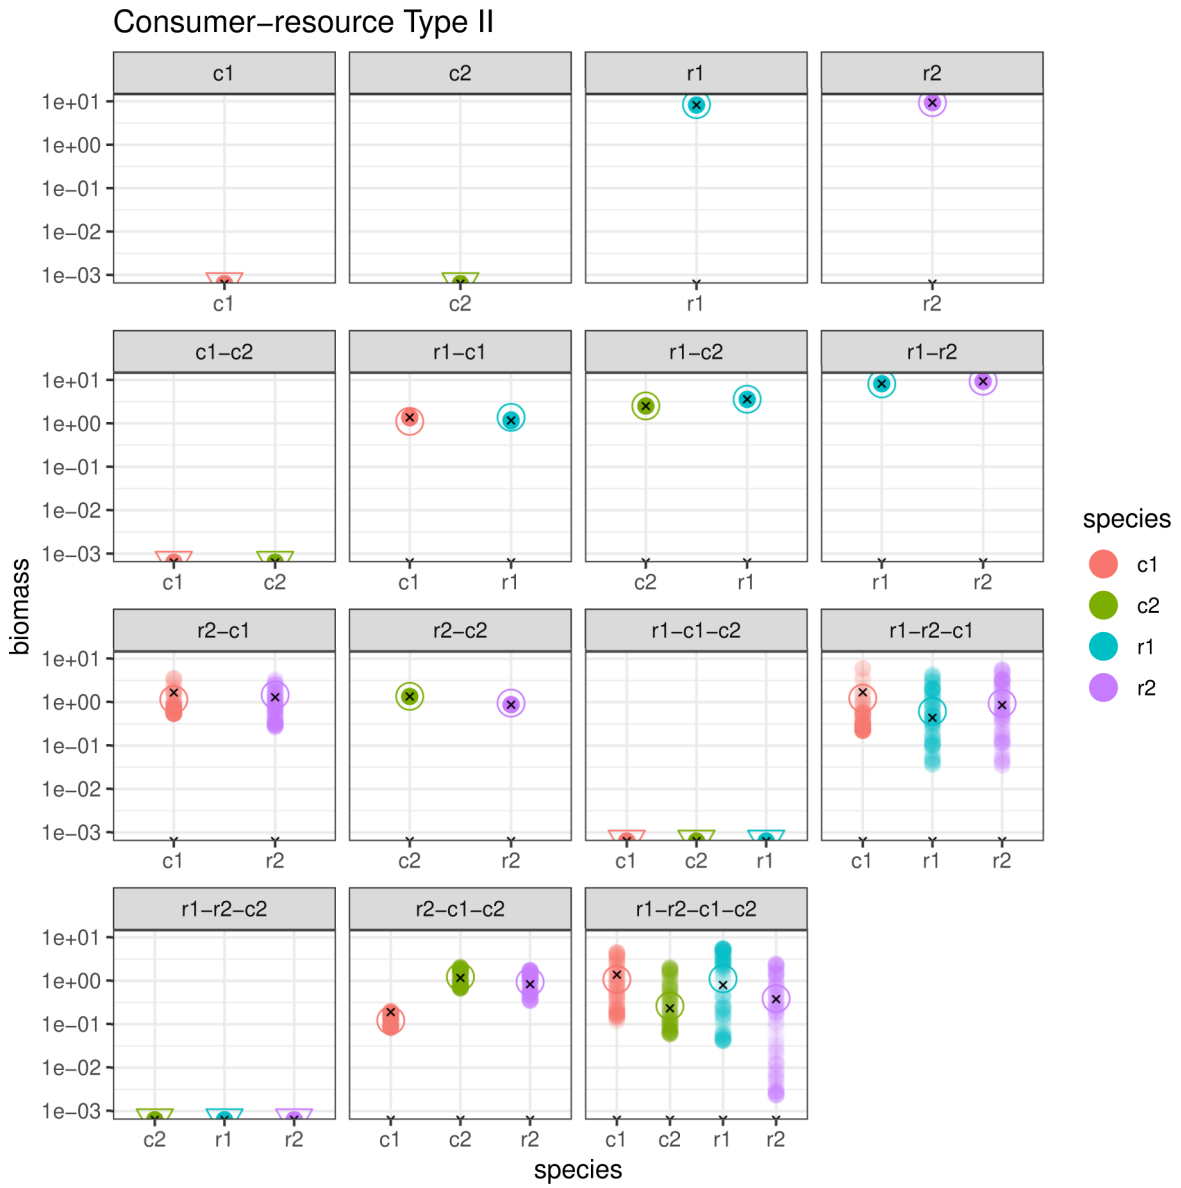

Extended Data Fig. 5 | Simulation results for a consumer-resource system. Two resources and two consumers are simulated in all possible combinations, giving rise to cases of coexistence at equilibrium, stable limit cycles, or extinctions. In all cases, the proposed method predicts the location of the equilibria of the nonlinear system quite perfectly, making the correct inference for all cases in which species cannot coexist.

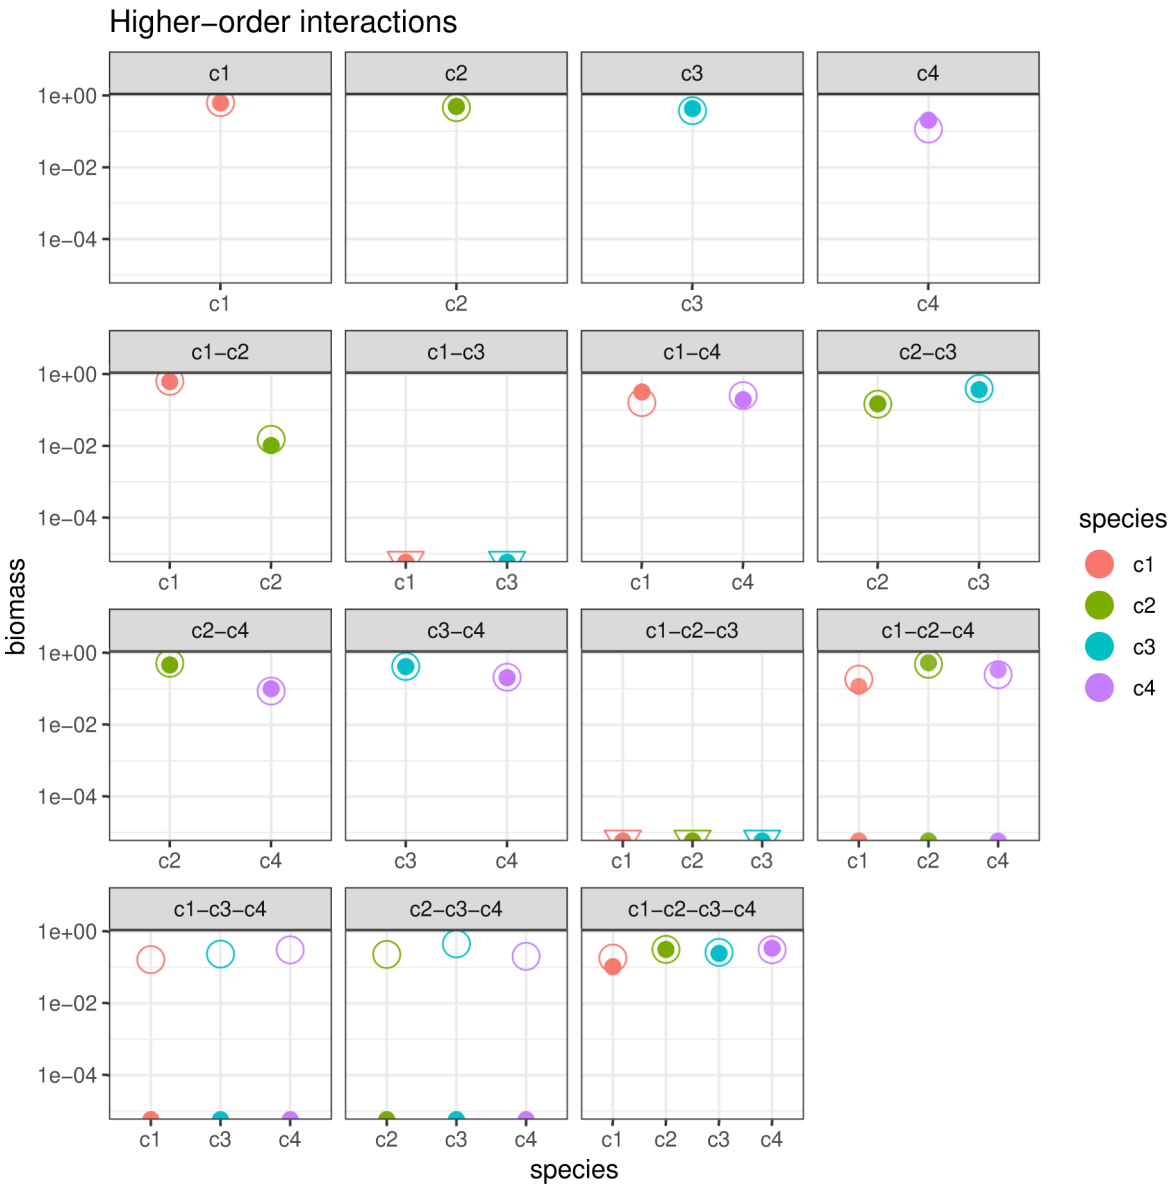

Extended Data Fig. 6 | Simulation results for a system characterized by higher-order interactions. Despite the strong effect of HOIs, the recovered solution is close to all endpoints. Moreover, the method correctly predicts that coexistence between c1 and c2, or c1, c2 and c3 is precluded. The method however predicts coexistence between two triplets, despite simulations showing that either no feasible equilibrium exists, or it is unstable.

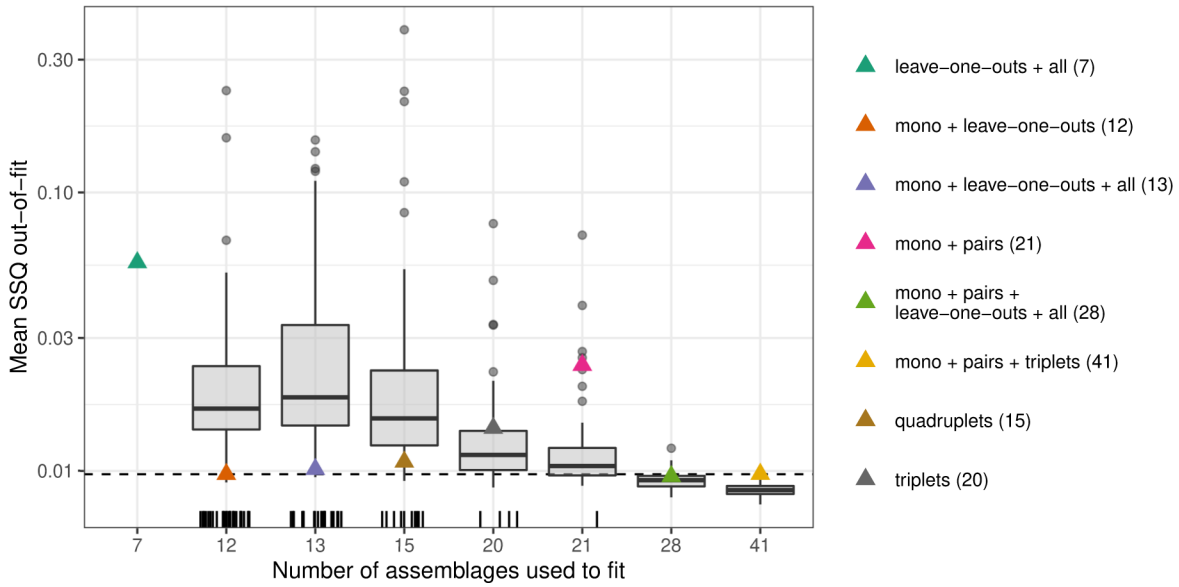

Extended Data Fig. 7 | Quality of fit for different experimental designs. We simulated a 6-species GLV model, in which all 63 possible assemblages lead to coexistence. We measured abundances at these endpoints by adding noise, and producing five 'replicates'. For each design, we use the specified number of assemblages to fit the model, and predict out-of-fit the abundance of all species at all other endpoints. Designs that produce qualitatively wrong predictions (that is, predicting a lack of coexistence for assemblages that do in fact coexist) are represented as vertical bars at the bottom of each boxplot. The horizontal dashed line marks the performance of the monoculture + leave-one-outs design, which fares among the best despite using only 12 assemblages to predict the remaining 51.

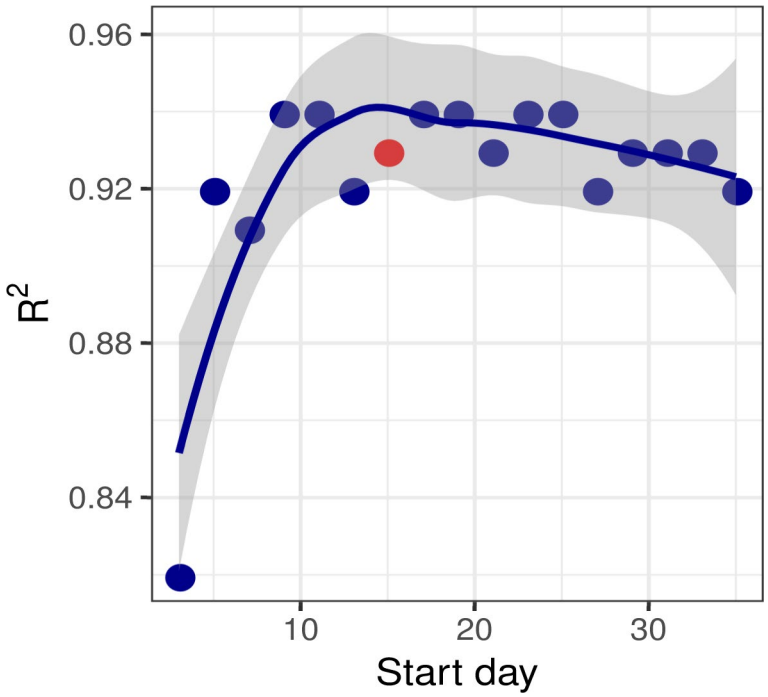

Extended Data Fig. 8 | The quality-of-fit for the protist system at 15 °C as a function of the sampling day. Rather than sample the community at days 15-17, as in the main text, we 'ended' the experiment at the indicated day, ±2 days, and fit the model with the corresponding endpoints. These results demonstrate that there is a clear initial period where the model fits poorly due to transient dynamics; followed by a stable period between days 10 and 20 where the approach performs well; followed by a period where the quality of fit starts to deteriorate as the species decline in abundances. The point in red denotes the point used in the main analysis, independently identified by quantifying when total biomass of the community stabilized.

nature research

Corresponding author(s): Daniel S. Maynard

Last updated by author(s): Oct 14, 2019

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒ A description of all covariates tested
☐	☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☐	☒ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☐	☒ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	No custom software was used to collect data for this analysis.
Data analysis	Analyses were conducted using R v3.6.1 and Stan v2.19.2.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The data needed to replicate the central findings of this work are available at: https://git.io/fjvON

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☐ Behavioural & social sciences ☒ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	We present a statistical method that uses a small number of experiments to predict the outcomes of all possible assemblages. We use three previously published, independent datasets to illustrate this method.
Research sample	The experimental datasets were obtained by searching Data Dryad, Figshare, and published papers for studies that reported the final abundances of each species in each community, and which contained enough endpoints to fit the model.
Sampling strategy	This is not relevant to this study, as we only use previously published data to illustrate a statistical method.
Data collection	Data were collected from the supplemental materials accompanying the published papers.
Timing and spatial scale	This is not relevant to this study, as we only use previously published data to illustrate a statistical method.
Data exclusions	As detailed in the main text, for each of the algae and protist systems we excluded one species that did not have enough endpoints to fit the model.
Reproducibility	The models were validated using a jackknife leave-out-out approach.
Randomization	This is not relevant to this study, as we only use previously published data to illustrate a statistical method.
Blinding	This is not relevant to this study, as we only use previously published data to illustrate a statistical method.

Did the study involve field work? ☐ Yes ☒ No

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging