Predicting global community properties from uncertain estimates of interaction strengths

György Barabás & Stefano Allesina

Journal of the Royal Society Interface (2015), 12:20150218

Abstract

The community matrix measures the direct effect of species on each other in an ecological community. It can be used to determine whether a system is stable (returns to equilibrium after small perturbations of the population abundances), reactive (perturbations are initially amplified before damping out), and to determine the response of any individual species to perturbations of environmental parameters. However, several studies show that small errors in estimating the entries of the community matrix translate into large errors in predicting individual species responses. Here we ask if there are properties of complex communities one can still predict using only a crude, order-of-magnitude estimate of the community matrix entries. Using empirical data, randomly generated community matrices, and those generated by the Allometric Trophic Network model, we show that the stability and reactivity properties of systems can be predicted with good accuracy. We also provide theoretical insight into when and why our crude approximations are expected to yield an accurate description of communities. Our results indicate that even rough estimates of interaction strengths can be useful for assessing global properties of large systems.

Keywords: community matrix, eigenvalues, pseudospectrum, random matrix, robustness, stability

1 Introduction

Ecological communities can be modeled through a set of deterministic differential equations keeping track of population growth as a function of the (biotic and abiotic) environment. One central question in the study of communities is their stability (May 1973): when perturbing the population abundances slightly, does the community tend to return to its original state? This question naturally follows from the fact that, in nature, populations undergo constant perturbations, which they have to withstand to avoid extinction. Mathematically, a community equilibrium is locally stable if the Jacobian matrix, evaluated at that equilibrium, has eigenvalues with all negative real parts.

The Jacobian evaluated at an equilibrium point is called the community matrix (Levins 1968, May 1973), whose (i, j)th entry measures the change in the total population growth rate of species i in response to a (small) change in species j's abundance, in units of inverse time. This matrix has many useful properties in addition to determining local stability. For instance, a stable equilibrium is reactive (Neubert and Caswell 1997, Tang and Allesina 2014) if perturbations, before damping, are initially amplified in a transient manner. Reactivity is measured by the leading eigenvalue of the Hermitian part of the community matrix, with positive values signaling reactive systems. In addition, the inverse community matrix can be used to determine the response of any species in the community to perturbations of the environment (i.e., not the abundances) via a community-wide sensitivity formula (Levins 1974, Bender et al. 1984, Yodzis 1988, Meszéna et al. 2006, Aufderheide et al. 2013, Barabás et al. 2014).

However, studies have revealed (Yodzis 1988, 2000, Dambacher et al. 2002, Novak et al. 2011) that even small uncertainties in estimating the entries of the community matrix translate into large errors of prediction. The problem is that small perturbations to the matrix can have large effects on the inverse matrix, to the point where even the directionality of the species' responses to environmental perturbations are predicted erroneously. Since the eigenvalues of the inverse matrix are the inverses of the eigenvalues, we do not expect the inverse matrix to be ill-behaved as long as all eigenvalues are far from the origin of the complex plane. But problems arise when some eigenvalues are close to zero: then the slightest error in measurement may lead to qualitatively different outcomes. Suppose the leading eigenvalue of a community matrix is (-0.01 ± 0.02)/year; its inverse is then either smaller than -33.3 years, or larger than 100 years. A slight measurement error can make the difference in whether the inverse matrix is deemed to have a large negative or a large positive eigenvalue. We do not know for certain whether community matrices of large natural systems possess eigenvalues that are close to zero, but a heuristic argument can be made that this is indeed the case based on the shape of the species-abundance curve (McGill et al. 2007), which shows that rare species are always overrepresented in natural communities. Since even small perturbations of the abundances could knock such rare species to extinction, the system must lie close to a transcritical bifurcation, meaning that some eigenvalues are necessarily very close to zero.

In large systems, it is already logistically impossible to measure every pairwise interaction (a community of 100 species with connectance 0.1 has 1000), let alone doing so with high accuracy. This fact, combined with the above argument, seems to imply the futility of relying on the "inverse problem" to obtain species' responses to environmental perturbations. What one *can* do, however—and this is the focus of our work—is to use imperfect information about the system to estimate properties which do not rely on inverting the community matrix, such as stability and reactivity. By avoiding the matrix inversion, small errors no longer translate into large ones, and so even crude estimates may provide useful information about a system.

Here we ask how well one can approximate the eigenvalue distribution of community matrices based on only an order-of-magnitude knowledge of interaction coefficients. Even though accurate measurement of all matrix entries is impossible, ecologists with extensive field experience can frequently rely on their intuition to classify interactions as "strong" or "weak". We assume that the magnitude of the strongest interactions are known (reasoning that, since they are strong, they

are the most likely to be noticed and possibly the easiest to measure), and coarse-grain each matrix entry into bins, based on their relative magnitudes in comparison to the strongest interactions. We show that the eigenvalue structure of complex community matrices can be captured well using this procedure, and therefore such qualitative information can be used to approximate system properties not relying on the inverse.

Below, after describing how we construct approximate matrices using imperfect data, we show how this procedure works on empirical datasets. We then apply the procedure to matrices that are randomly generated, as well as those generated by the Allometric Trophic Network model (Berlow et al. 2009). We then go on to give a theoretical justification to our method, building on the theory of random matrices (Bai and Silverstein 2009) and pseudospectra (Trefethen and Embree 2005). We end by discussing the limitations of our approach and its relevance for the classic stability-complexity debate in community ecology.

2 Constructing approximate matrices

We are given a community matrix A, and we would like to know its eigenvalues, but information on A's entries is limited. Quantitatively, we assume we know only the magnitudes of the largest positive and negative entry (denoted by p and n, respectively), and the zero entries of A, i.e., we know which interactions are absent. Apart from this quantitative information, we assume a qualitative knowledge of all other matrix entries: based on expert opinion or other indirect information, we know whether a given entry is strong or weak compared to n and p. Based on this, we assign numerical bins into which the entries of A will be lumped. Let B denote this binned approximation to A. We ask how well the spectrum of B approximates that of A.

We thus need to specify a binning procedure. First, we choose a number of bins. We then assign numerical values to these bins. Finally, each entry of B is set to the value of the bin closest to the value of the corresponding entry in A. For a given binning, we use the notation

$$(x_1, x_2, \ldots, x_k),$$

meaning that the first bin goes from x_1 to halfway between x_1 and x_2, the second goes from halfway between x_1 and x_2 to halfway between x_2 and x_3, and so on, until the last bin going from halfway between x_{k-1} and x_k to x_k. As an example, consider the matrix

$$A = \begin{pmatrix} 7.8 & 6.7 & 3.7 & -1.2 \\ -7.5 & 2.6 & -7.4 & 0 \\ -10.0 & -6.9 & 0.4 & 5.8 \\ 0 & 0 & 10.0 & -8.7 \end{pmatrix}.$$

Its strongest negative entry is $n = -10$; its strongest positive entry is $p = 10$. If we now decide to construct B using three bins with values $(n,\ 0,\ p)$, we get

$$A = \begin{pmatrix} 7.8 & 6.7 & 3.7 & -1.2 \\ -7.5 & 2.6 & -7.4 & 0 \\ -10.0 & -6.9 & 0.4 & 5.8 \\ 0 & 0 & 10.0 & -8.7 \end{pmatrix} \quad \Rightarrow \quad B = \begin{pmatrix} 10 & 10 & 0 & 0 \\ -10 & 0 & -10 & 0 \\ -10 & -10 & 0 & 10 \\ 0 & 0 & 10 & -10 \end{pmatrix}.$$

In principle, the choice for the number of bins and their values is arbitrary. Here we consider the following, more specific procedure. Let the number of bins be $k \geq 3$ an odd integer. Let us specify a binning resolution constant b whose powers help define the bins; in effect, b fixes the definition of an order of magnitude. The k bins are then given by

$$(n,\ nb^{-1},\ nb^{-2},\ \ldots,\ nb^{-(k-1)/2},\ 0,\ pb^{-(k-1)/2},\ \ldots,\ pb^{-2},\ pb^{-1}, p).$$

Using our previous example for A, we can bin A with $k = 5$ and binning resolution $b = 10$. Since $n = -10$ and $p = 10$, the bins are $(-10,\ -1,\ 0,\ 1,\ 10)$:

$$A = \begin{pmatrix} 7.8 & 6.7 & 3.7 & -1.2 \\ -7.5 & 2.6 & -7.4 & 0 \\ -10.0 & -6.9 & 0.4 & 5.8 \\ 0 & 0 & 10.0 & -8.7 \end{pmatrix} \quad \Rightarrow \quad B = \begin{pmatrix} 10 & 10 & 1 & -1 \\ -10 & 1 & -10 & 0 \\ -10 & -10 & 0 & 10 \\ 0 & 0 & 10 & -10 \end{pmatrix}.$$

Our scheme for binning matrix entries involves exponentially shrinking bin sizes. Any number of other schemes may be implemented—e.g., linear binning, where adjacent bins are equally spaced. The reason for choosing the exponential scheme is that there is both theoretical (Tang et al. 2014) and empirical (Woodward et al. 2005) evidence that the interaction strengths in large ecological networks follow a distribution close to lognormal. Therefore the exponential binning strategy is expected to yield better resolution of the underlying data than a linear one. This is not to say that different binning schemes are not better suited to different problems. However, regardless of the problem at hand, one should make sure that the procedure does not depend on the particular choice of units the matrix entries are expressed in, which means that the procedure should only be sensitive to the relative magnitudes of the matrix entries, not their actual numerical values which are unit-dependent. The exponential scheme obviously satisfies this requirement.

3 Exploratory analysis: some empirical datasets

We show the eigenvalues (red dots) of community matrices from nine different empirical interaction webs in Fig. 1, parameterized using allometric scaling relationships (Tang et al. 2014). These matrices were binned with $b = 4$, $k = 7$; the eigenvalues of the binned matrices (blue circles) are also shown. The number of species in each web is indicated in the panel titles. To interpret the eigenvalue distributions correctly, note the scale discrepancy between the real and imaginary axes.

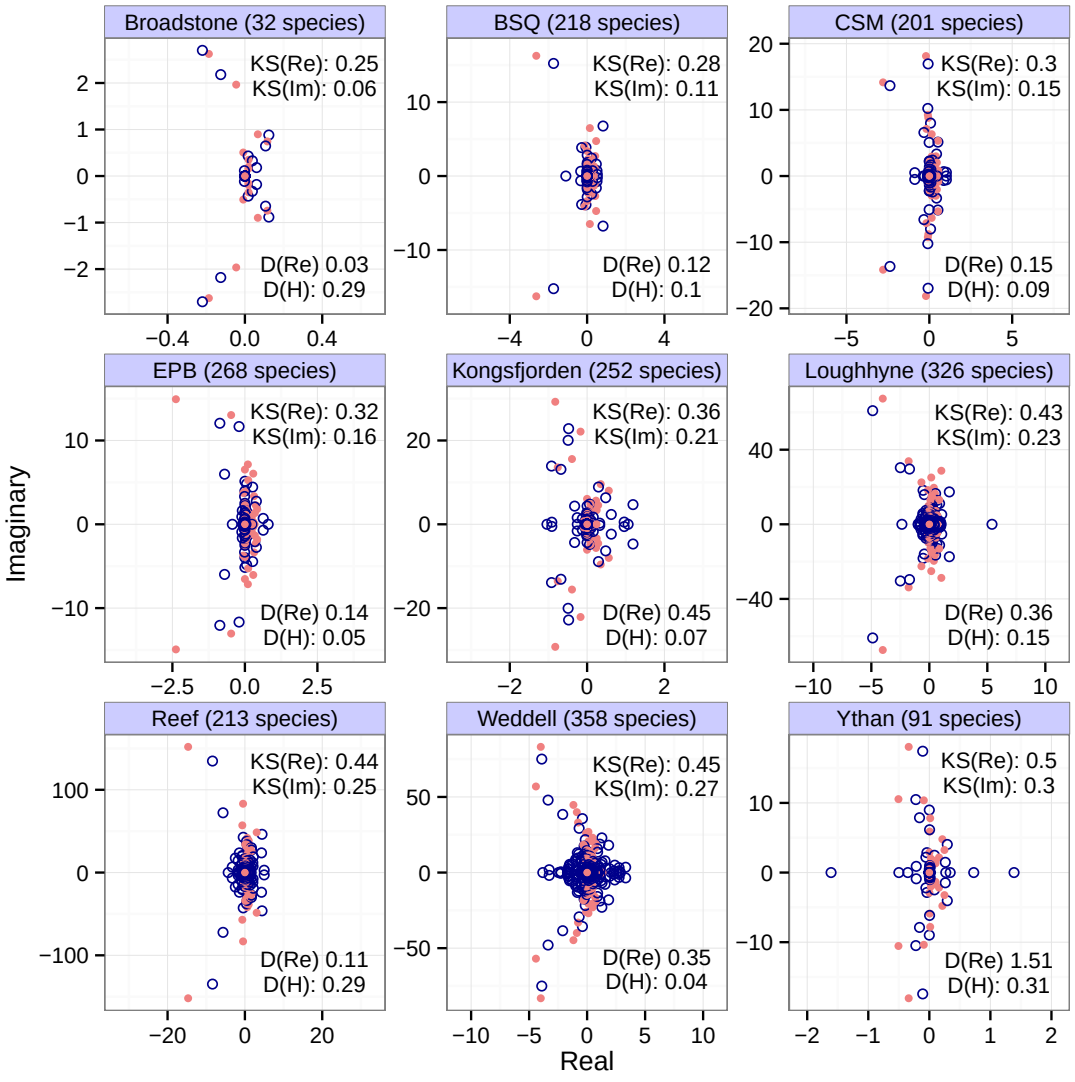

Figure 1: Eigenvalues of community matrices (red dots) and their binned counterparts (blue circles) from nine different empirical interaction webs. Datasets used (left-right, top-bottom): Broadstone Stream, Baja San Quintin, Carpinteria Salt Marsh, Estero de Punta Banda, Kongs Fjorden, Lough Hyne, Caribbean Reef, Weddell Sea, and Ythan Estuary (see Tang et al. 2014 and references therein). The Kolmogorov-Smirnov distance between the original and binned eigenvalue distributions are given in each panel (top right), as well as the relative difference in the leading eigenvaules of the matrices and their Hermitian parts (bottom right).

The spectra of the binned matrices capture general features of the original ones, such as a larger bulk of eigenvalues near the origin of the complex plane, and semicircular arcs composed of a handful of eigenvalues protruding from this bulk towards the left half plane. More quantitatively, one may consider the Kolmogorov-Smirnov distance (a number between 0 and 1, corresponding to the maximum vertical distance between two empirical cumulative distribution functions) between the original and binned matrices' eigenvalue distributions as a measure of how well the distributions approximate each other. Since the Kolmogorov-Smirnov distance is defined for univariate samples, we consider the distance between the real and the imaginary parts of the eigenvalues separately. Their numerical values are shown in the panel insets of Fig. 1 (top two lines).

From the point of view of local stability, the leading eigenvalue (that with the largest real part) is of crucial importance: the sign of the real part of this eigenvalue determines whether the system is stable (negative) or unstable (positive). We consider the difference between the leading eigenvalues of the binned and original matrices to see how well it is captured. However, the raw difference itself is not informative, since the numerical value of this eigenvalue difference simply depends on the choice of units we measure the matrix entries in. A better question is how well the leading eigenvalue of the original matrix is approximated compared to the total spread of the eigenvalues; that is, can we say that the leading eigenvalue of the binned matrix is close to that of the original one, compared to the total range of the real parts of all the eigenvalues of the original matrix? These values are included in the panels under "D(Re)". We can see that, for the datasets presented, the binning approximation always errs on the conservative side, overestimating the actual leading eigenvalue. This conservatism is due to the fact that our method of matrix binning include both largest entries of the matrices as bins, and also the zero bin. The binning procedure will lump several entries into these extreme bin values. Therefore the variance of the entries of the binned matrix will in general exceed that of the original matrix, which leads to a higher variance in the eigenvalue distribution as well.[1]

Note that none of the matrices in Fig. 1 are stable. This is because the empirical studies which they are based on only document trophic links, with no data on self-interactions. Due to this lack of information, we set all diagonal entries to zero. However, since the sum of the diagonal entries of a matrix is also the sum of its eigenvalues, such matrices cannot be stable. Their instability is therefore likely an artifact of our ignorance rather than an actual phenomenon. However, since we are interested in the ability of binned matrices to approximate spectra (whether they are stable or not), this lack of information is not a problem for us here.

Just as with stability, we can also examine reactivity, measured by the leading eigenvalue of A's Hermitian part $H(A) = (A + A^\top)/2$; $A^\top$ is the transpose of A. We therefore calculate $H(A)$ and $H(B)$, and ask how well the leading eigenvalue of the latter approximates that of the former in comparison with the total spread of the eigenvalues of $H(A)$. The results, for the empirical webs of Fig. 1, are shown under "D(H)".

[1] Let A be a normalized matrix, with entries having zero mean and unit variance. Denote its eigenvalues by λ_A and their variance by $\mathrm{Var}(\lambda_A)$. Now let $B = \alpha A$ with $\alpha > 0$ a scaling constant. B's entries then have zero mean and variance α^2, its eigenvalues are $\lambda_B = \alpha\lambda_A$, and their variance is $\mathrm{Var}(\lambda_B) = \mathrm{Var}(\alpha\lambda_A) = \alpha^2\mathrm{Var}(\lambda_A)$. In other words, increasing/decreasing the variance of the matrix entries increases/decreases the variance of the eigenvalues as well.

Let us highlight some of the general conclusions from these motivating examples. First, the eigenvalue distributions of the original and binned matrices roughly overlap: their means and variances are very similar in both the real and imaginary directions. Second, some of the finer structure of the eigenvalues distributions is also captured: more eigenvalues near the origin, and a few ones protruding in arcs toward the left half plane. Third, despite these features, the match between the actual and predicted values for stability and reactivity is not always very close—cautioning that the approximation could fail to capture specific numerical estimates for these properties. The method therefore reveals the general features of the system, and gives a rough idea of the quantitative details.

4 Randomly generated interaction matrices

Here we test the binning procedure on randomly generated interaction matrices. For each matrix we first fixed the number of species S and the connectance C. We then generated two uniform random numbers between 0 and 1 to determine the types of all nonzero interactions. The first number is the fraction of trophic interactions, the second the fraction of mutualistic ones out of those that were nontrophic (the rest of the interactions were designated competitive). For each of the three interaction types there was an associated probability distribution from which the matrix entries pertaining to the interaction type in question were drawn. The procedure for generating the probability distribution was the same in all cases: first, the shape of the probability distribution was determined: either lognormal or Gamma. (We did this to check the robustness of our results to the choice of the underlying distribution. The results are indeed robust; see Supporting Information, Figs. S23-S26). Second, the mean and standard deviation of the distribution were uniformly sampled: for the mean, from $[0.1,\ 10]$, and for the standard deviation, from $[1,\ 10]$. For the trophic interactions, a random conversion efficiency was uniformly sampled from $[0.05,\ 0.2]$ to take into account the limited energy flow between trophic levels. To set the diagonal of the matrix, we sampled each diagonal entry from either a lognormal or a Gamma distribution, times -1 to keep the diagonal entries negative. The mean and standard deviation of the distribution were randomly drawn as in the offdiagonal case. See the Supporting Information for a more detailed breakdown of our method for generating these matrices.

We repeated the parameterization for four different values of the species richness S (50, 100, 250, and 500) and of the connectance C (0.1, 0.25, 0.5, and 1), in all possible combinations. Then, 300 replicates of all cases were generated. All resulting matrices were binned, based on all possible combinations of the following three variables. First, the number of bins was chosen to be either 3, 5, 7, or 9. Second, the binning resolution was set to 2, 4, 6, 10, and 14. Third, we took into account the effect of accidentally misclassifying matrix entries. In empirical situations, it seems likely that strong interactions may accidentally be deemed weak (or vice versa), given the insufficient, qualitative information one uses to generate approximate matrices. We therefore explored what happens when we deliberately misclassify some fraction of the matrix entries. In doing so, we assumed that zero interaction strengths do not get misclassified (absent interactions cannot ever

be observed, and so will not be accidentally classified as present), and that the bin category of a nonzero interaction can only move up or down one bin (a strong interaction may be misclassified as weak, but not as zero). We used three rates of misclassification: 0%, 10%, and 20%. Our results were not strongly affected by misclassification (Supporting Information, Figs. S1-S5). Here in the main text we always show results with a 10% misclassification rate.

First, we explored how accurately the leading eigenvalue is captured depending on the number of bins k, the binning resolution b, and the species richness S (Fig. 2). It is apparent that $b = 4, 6$ and $k = 7, 9$ yield the best results. For instance, 90% of all results with $b = 4$ and $k = 7$ have a difference between r_B and r_A less than 16% of σ_A, the total spread of the real parts of A's eigenvalues. Higher values of k leading to better predictions was expected—the more bins we use, the more accurate the predictions will be. Notice though that there are diminishing returns. In a sense this is fortunate, because using more than 7-9 bins is probably not feasible in practice (seven bins means that, apart from zero, we have a "strong", a "medium", and a "weak" category both for positive and negative interactions). The result that $b \approx 4$ is optimal concerns the "best" choice for the definition of an order of magnitude. This suggests one should consider interactions sufficiently different if they differ by a factor of about four. Importantly, this choice proves to be consistently the best through various models and various metrics considered (see the next section and the Supporting Information).

Instead of the relative, quantitative measure of how well the leading eigenvalue is approximated, we may also ask how often is it true that if a matrix A is stable, then its binned counterpart B is also stable. Of most interest are those matrices whose leading eigenvalues lie close to the imaginary axis, since in these cases a small perturbation to the spectrum may in principle change their stability properties. Once again, being "close" to the imaginary axis should not be measured on an absolute scale, since any given distance is unit-dependent, and the binning procedure is itself scale-invariant (Section 2). The relevant question is whether the leading eigenvalue is close to the imaginary axis compared to the total spread of all eigenvalues; i.e., whether $|r_A|/\sigma_A$ is small. The result (Supporting Information, Figs. S19-S20) depends on b and k; for example, when using only three bins, one is more likely to misjudge stability than not. However, for $b = 4$ and $k = 7$, of those results for which $|r_A|/\sigma_A < 0.05$, stability is accurately predicted in 90% of all cases. And, for $|r_A|/\sigma_A < 0.1$, this accuracy increases to 97%, after which it quickly approaches 100%, always increasing.

Similarly to the leading eigenvalue, one can also look at reactivity and how well it is predicted (Supporting Information, Figs. S3-S5). Just as before, $b = 4, 6$ and $k = 7, 9$ provide the best approximations, with more than 90% of all results falling within 23% of the total spread of the eigenvalues of $H(A)$, the Hermitian part of A (Supporting Information, Figs. S21-S22).

We also consider the effect of changing the connectance on the efficiency of binning—it turns out however that this effect is neither systematic nor very strong (Supporting Information, Figs. S11-S18).

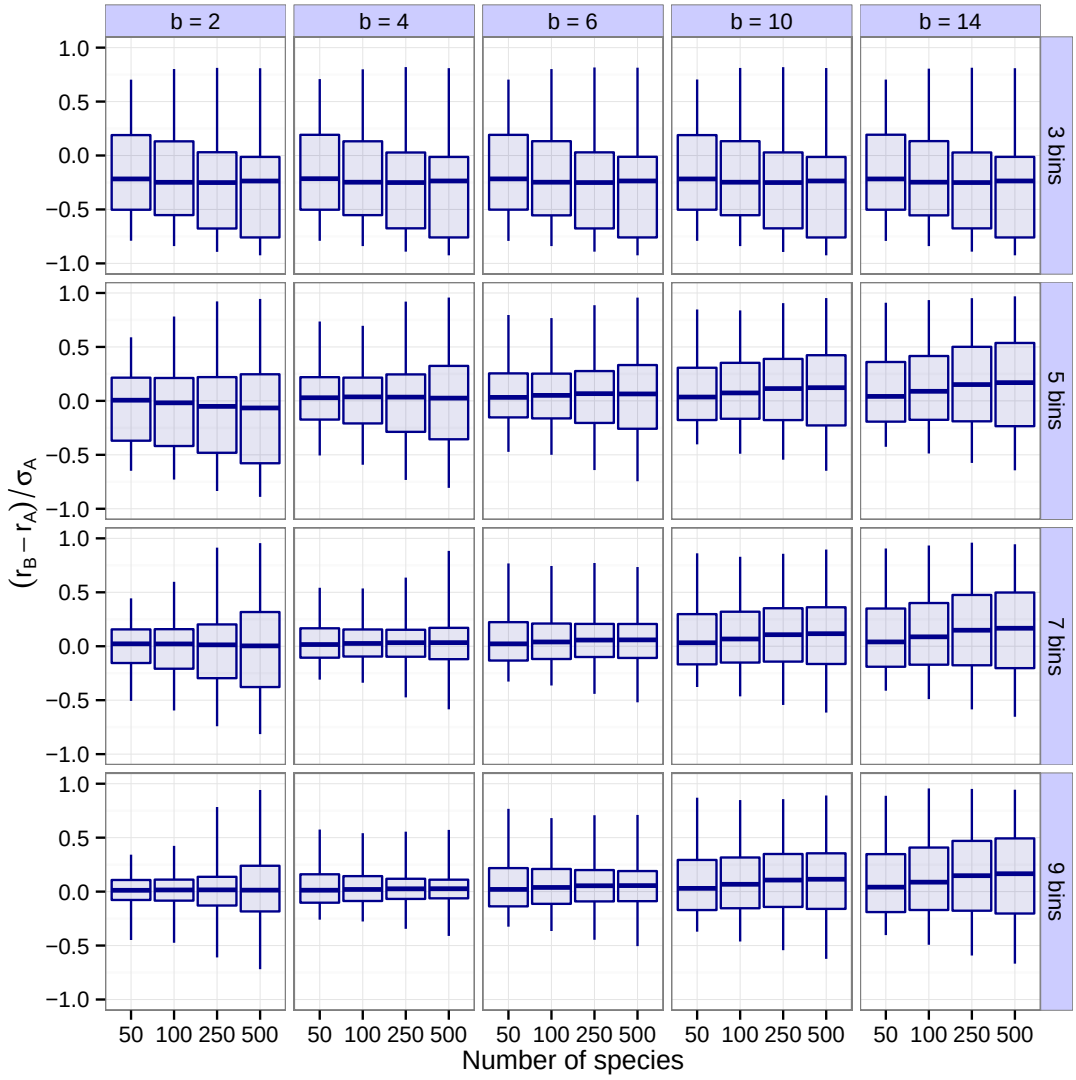

Figure 2: Box plots of how the leading eigenvalues of randomly generated community matrices A are captured by those of their binned counterparts B. Interpretation of the box plots: median (lines), 5% to 95% quantiles (boxes; note that they encompass 90% instead of the usual 50% of the data), and ranges (whiskers). Each matrix is binned with a 10% misclassification rate. Rows correspond to different values of the binning resolution; columns to different numbers of bins. The data in each panel are separated based on species richness. Panel ordinates show the difference between the leading eigenvalue r_A of the original and r_B of the binned matrices relative to σ_A, the total range of the real parts of the original matrix's eigenvalues.

5 The Allometric Trophic Network

The previous section explored the effects of matrix binning on randomly generated interaction matrices. Here we consider a mechanistic model of multispecies communities, the Allometric Trophic Network (Berlow et al. 2009). In this model there are a number of noninteracting abiotic resources (here we assume there are two), primary producers utilizing those resources, and consumers eating either the producers or other consumers. The feeding network is generated using the niche model (Williams and Martinez 2000). Consumers interact with their resources via generalized functional responses, which may include consumer interference. These interaction terms are functions of the organisms' average body masses, calculated based on species' trophic levels and simple allometric relationships.

We generated $10,000$ different communities using the Allometric Trophic Network model. In each simulation, we started out from a food web generated by the niche model, with 50 initial species whose abundances were uniformly distributed between 0.05 and 0.2. We followed the methodology described in Berlow et al. (2009) in every aspect to parameterize the model, except in choosing the Hill exponents for the trophic interactions: instead of randomly generating it for every interaction, we assumed it had the constant value of 2. This was done to make the system converge to a fixed point instead of a limit cycle or chaotic attractor, which is important because we are interested in predicting local asymptotic stability (indeed, in our simulations we only ever observed convergence to a fixed point). The model was run until equilibrium was reached, at which point we calculated the Jacobian to obtain the matrix A (see the Supporting Information for a detailed description of our methods). Fig. 3a shows a spectrum generated by the procedure, along with that of its binned counterpart ($b = 6$, $k = 7$). In the particular simulation shown, 24 species and both abiotic resources survived to stably coexist out of the initial 50 species and two resources. The number of species plus resources persisting at equilibrium was variable between runs, and approximately normally distributed with mean 13.6 and standard deviation 3.0.

The matrices were subsequently binned with the binning resolution b running through 2, 4, 6, 10, and 14; the number of bins k taking on the values 3, 5, 7, and 9, and the rate of misclassification being either 0%, 10%, or 20%. Since the number of species one ends up with is highly variable in this model, we do not factor the results based on the number of species. Except for the case with 3 bins, the leading eigenvalue is captured well (Fig. 4), with 90% of all other cases having a relative error less than 13%, and in some cases less than 6% (e.g., for $b = 4$, $k = 7$). Similarly, reactivity (Supporting Information, Fig. S8-S10) is captured with relative error less than 12% in 90% of all cases with $b = 4$, $k = 7$, and misclassification rate 10%.

6 Theoretical underpinning

Here we connect the matrix binning procedure with more rigorous mathematical concepts, to give a theoretical underpinning to why and when the method is supposed to work. We employ two arguments, one based on the theory of random matrices, the other on the concept of pseudospectra.

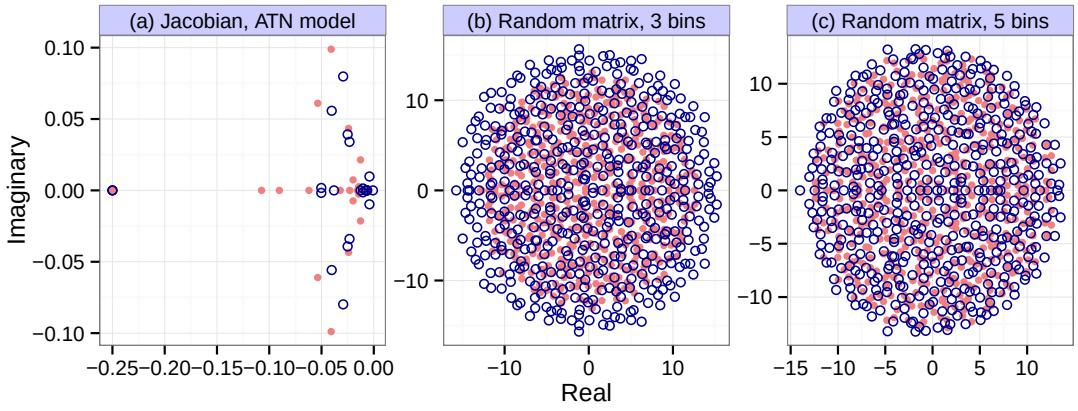

Figure 3: Spectra of various matrices (red dots) and their binned counterparts (blue circles). Panel (a): Jacobian matrix obtained at a stable equilibrium of one particular run of the Allometric Trophic Network model, binned with binning resolution $b = 6$ and $k = 7$ bins. Panel (b): Random matrix with independent uniformly distributed entries between -1 and 1. It is binned with three bins $(-1,\ 0,\ 1)$. Panel (c): the same random matrix, but binned with five bins $(-1,\ -0.5,\ 0,\ 0.5,\ 1)$.

6.1 Random matrices

Although empirical interaction webs are manifestly *not* random, an intuition for why the matrix binning procedure works may be gained by connecting it with random matrix theory (Bai and Silverstein 2009). There are several results in the theory of random matrices concerning the distribution of matrix eigenvalues in the complex plane, such as the circular (Girko 1984, Tao et al. 2010) and elliptic (Sommers et al. 1998) laws, which have found ecological applications as well (Allesina and Tang 2012, Tang et al. 2014). Here we only consider the simplest version of the circular law; more complexity can be incorporated analogously (see Supporting Information). Suppose the entries of the $S \times S$ matrix A are drawn independently from the same underlying probability distribution $p_A(x)$, which has mean zero and variance V_A. Then the law states that for S large, the eigenvalues are uniformly distributed in a circle of radius $\sqrt{SV_A}$ in the complex plane, centered at the origin.

Note that the circle's radius only depends on the variance of $p_A(x)$ but not its shape. This important property (Tao et al. 2010) means that two completely different underlying probability distributions will lead to the same eigenvalue distribution as long as their mean is zero and their variances are equal (in the limit of S going to infinity, the distributions would converge to be exactly the same; for S large but finite, there are slight but negligible differences). Even if the variances are not equal, the only difference between the eigenvalue distributions will be in the radii of the circles within which the eigenvalues are found.

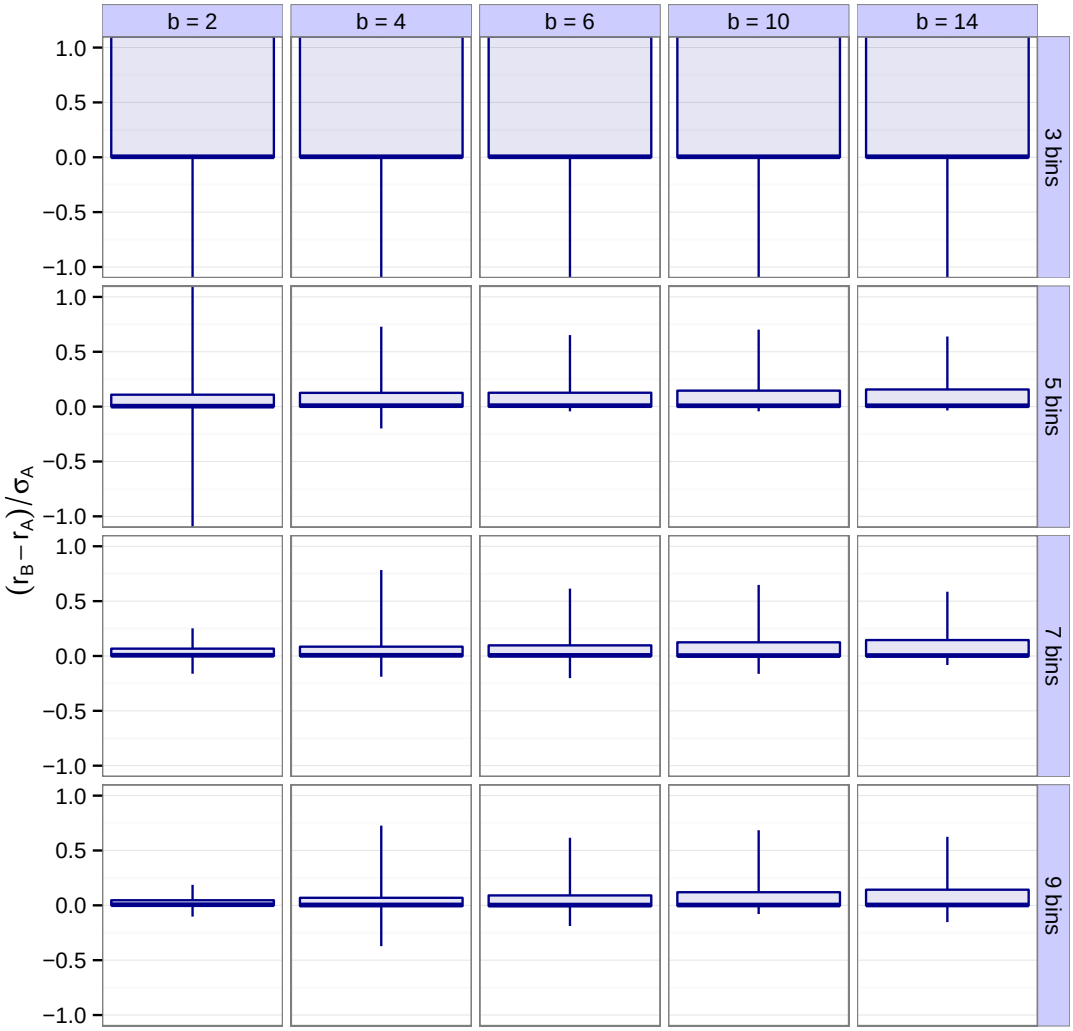

Figure 4: Box plots of how the leading eigenvalues of community matrices A are captured by their binned counterparts B, where the As are generated by the Allometric Trophic Network model. The figure is organized just like Fig. 2, except the data in the panels are not separated based on species richness.

The key idea concerning matrix binning is as follows. Consider a random matrix A whose entries are drawn from some distribution $p_A(x)$. We create its binned counterpart B. But the binned matrix B is just another random matrix with a different underlying probability distribution: we essentially replace the original $p_A(x)$ with a discrete distribution $p_B(x)$, one which we can calculate from $p_A(x)$ and the bin positions. We will then know the eigenvalue distribution of B as well, since that only depends on $p_B(x)$'s variance V_B. The spectra of A and B may then be compared analytically.

As an example, let the entries of A come from the uniform distribution $p_A(x) = \mathcal{U}[-1,1]$, which has variance $V_A = 1/3$. Let us bin A with three bins $(-1,\ 0,\ 1)$. The probability, on average, of any one entry being lumped into the -1 or 1 bins is $1/4$, while the probability of being lumped into the 0 bin is $1/2$, defining the discrete distribution $p_B(x)$. This distribution has variance $V_B = 1/2$. The eigenvalues of A are therefore uniformly distributed in a circle of radius $r_A = \sqrt{S/3}$, and those of B in a circle of radius $r_B = \sqrt{S/2}$ (Fig. 3b). If we now ask how well the leading eigenvalue is approximated, we first note that they are simply given by r_A and r_B (since the eigenvalues fall in a circle). We therefore take the ratio of the two radii to assess the goodness of the approximation: $r_B/r_A = \sqrt{3/2} \approx 1.22$. The binned matrix overestimates the leading eigenvalue by this factor.

One could repeat the analysis with a more refined binning scheme, for instance with $k = 5$ and $b = 2$. We then have five bins $(-1,\ -0.5,\ 0,\ 0.5,\ 1)$ instead of the original three, leading to $V_B = 3/8$ (see Supporting Information). Then the ratio of the circles' radii (and that of the leading eigenvalues) is $r_B/r_A = \sqrt{V_B/V_A} = \sqrt{9/8} \approx 1.06$, a near-perfect match brought about by the refinement of the binning resolution (Fig. 3c).

In summary, certain classes of random matrices allow for a simple analytical evaluation of the effects of matrix binning. Although real-world matrices are not going to conform to true random matrices exactly, there are nevertheless good reasons to make use of them anyway. Most importantly, random matrix theory serves as a theoretically well-understood benchmark, a reference model, which, when not fitting real-world data, reveals properties of the empirical system that are causing the departure from the random expectation, thus facilitating a better understanding of the system. Quite apart from this justification, random matrix theory actually does have some success in interpreting empirical data as well (Tang et al. 2014), suggesting its use may prove more than purely theoretical.

6.2 Pseudospectra

The spectrum of a matrix A is the set of complex numbers that are eigenvalues of A. In contrast, its ε-*pseudospectrum* (Trefethen and Embree 2005) is the set of complex numbers that are eigenvalues of all possible perturbed matrices $A + P$, with $\|P\| < \varepsilon$ (the matrix norm $\|P\|$ is defined as the square root of the largest eigenvalue of P^*P, where P^* is the conjugate transpose of P). Whereas the spectrum is composed of discrete points in the complex plane, the ε-pseudospectrum comprises of the union of regions of various sizes around the original eigenvalues. See the Supporting Information for an algorithm to compute pseudospectral regions.

An important result (Trefethen and Embree 2005, Theorem 2.2) states that the ε-pseudospectrum of normal matrices (i.e., matrices A for which $A^*A = AA^*$) is the union of circular disks of radius

ε around A's unperturbed eigenvalues. Moreover, such matrices have the smallest possible pseudospectra. Any deviation from normality will increase the size of this set, with strongly nonnormal matrices potentially having very large pseudospectral regions even for small values of ε.

Pseudospectra provide a rigorous and general measure of the effect of perturbations on the eigenvalues of matrices. Importantly, the binning procedure can be thought of as applying a certain perturbation P to the underlying community matrix A, with $A+P=B$, where B is the binned matrix. The pseudospectrum reveals how sensitively the eigenvalues respond to the perturbation induced by binning. In Fig. 5 the blue regions show the ε-pseudospectrum for each of the nine empirical webs of Fig. 1, with $\varepsilon = \|P\|$ being the appropriate perturbation norm for each web; the red dots are the unperturbed eigenvalues of A.

Pseudospectra measure the union of the effects of *all* possible perturbations of a given norm, which is why the blue regions are much wider than the positions of the binned eigenvalues on Fig. 1 would warrant them to be. Importantly, since normal matrices are the least sensitive to perturbations, the comparison of the actual pseudospectrum with the smaller one that would have been obtained if the matrix had been normal carries useful information: pseudospectral regions much larger than the one obtained under the assumption of normality signal a matrix whose spectrum is oversensitive to perturbing its entries. The red regions in Fig. 5 were computed as the union of disks of radius $\varepsilon = \|P\|$, i.e., it is what the pseudospectrum would look like if the matrices were in fact normal. Since the blue regions barely exceed the red ones, the empirical matrices are "almost normal" and therefore their spectra are not overly sensitive to perturbations of their entries.

Calculating pseudospectra is straightforward but computationally expensive. In the Supporting Information we therefore introduce a much simpler metric, the *scaled departure from normality* $\mathrm{depn}(A)$, which characterizes matrix sensitivity with a single number such that $\mathrm{depn}(A) \leq 1$ guarantees low sensitivity to perturbations. The value of $\mathrm{depn}(A)$ for each empirical matrix is reported in the panel titles of Fig. 5. All are significantly lower than one, implying their spectra are not sensitive to perturbations—in line with what we see on their pseudospectra.

7 Discussion

Our results show that certain global community properties, such as stability (May 1973) and reactivity (Neubert and Caswell 1997, Tang and Allesina 2014), can be predicted using crude, order-of-magnitude estimates of community matrix entries. These properties depend, broadly speaking, on the distribution of eigenvalues (stability depends only on the leading eigenvalue; reactivity on their whole ensemble), which was reasonably captured by the crude approximation. We gave theoretical justification to when and why this would be so, one based on the theory of random matrices, the other on the concept of pseudospectra. To check the robustness of our method, we applied it to three very different scenarios: empirical interaction webs parameterized via allometric relationships (Tang et al. 2014), randomly generated matrices, and those generated by the Allometric Trophic Network model (Berlow et al. 2009). Though the degree to which the method produced accurate results was situation-dependent, on the whole, it was able to make reliable predictions in each of these cases.

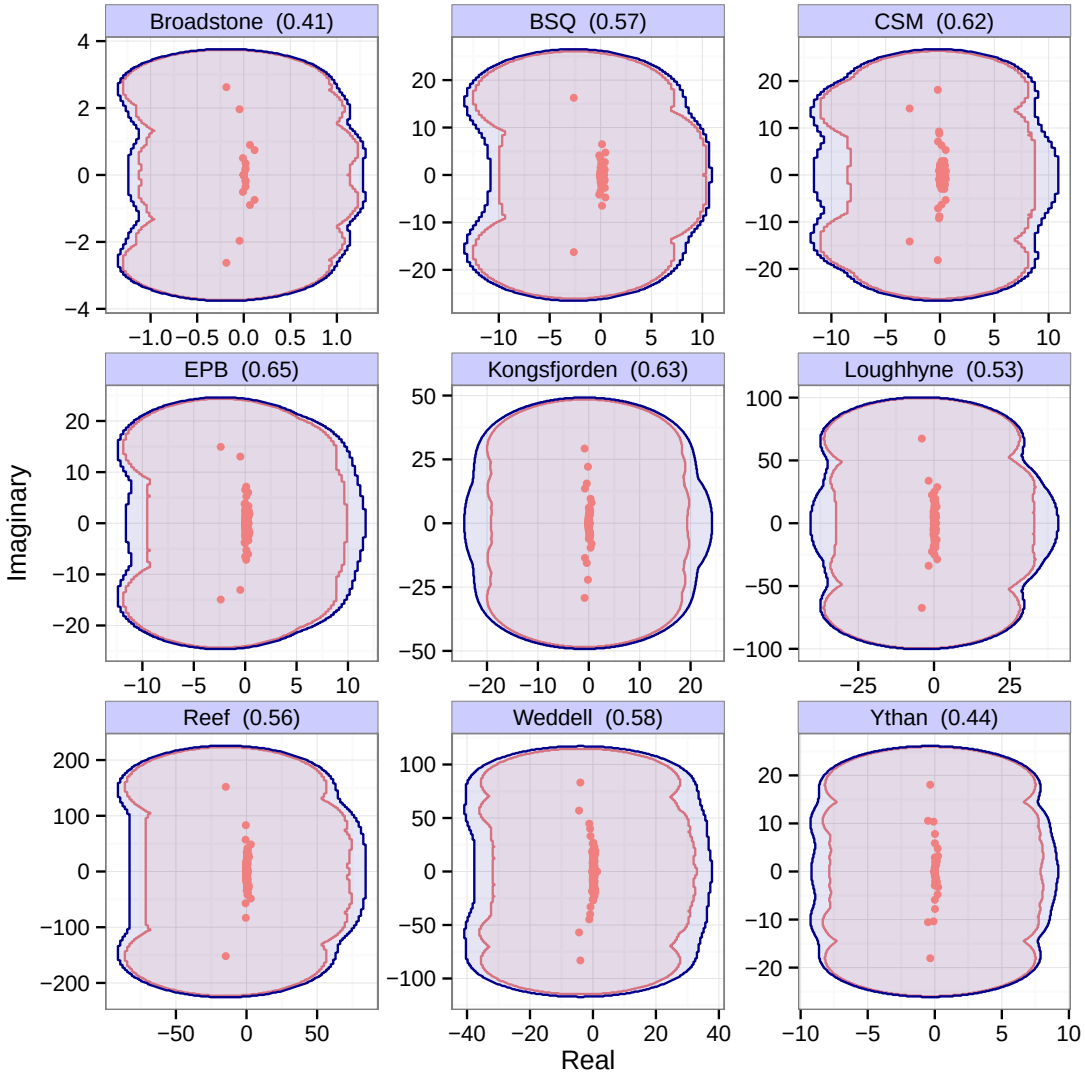

Figure 5: Pseudospectra of the empirical matrices of Fig. 1. The red dots show the original eigenvalues. The blue regions show the ε-pseudospectra with ε equal to the norm $\|B-A\|$, where B is the binned and A is the original matrix. The red regions are what the same pseudospectra would look like if the matrices were normal ($A^*A = AA^*$). The numbers in the panel titles are the given matrix's scaled departure from normality; see the Supporting Information for calculating these values.

This predictability appears to be at variance with earlier work emphasizing that even small errors in measuring the entries of the community matrix translate into large errors of prediction (Yodzis 1988, 2000, Dambacher et al. 2002, Montoya et al. 2009, Novak et al. 2011, Aufderheide et al. 2013). If the spectrum of the community matrix is well approximated, why would this be the case? The answer, we believe, is that the response of species to press perturbations depends on the inverse spectrum. Due to the preponderance of rare species in natural communities, such systems are necessarily close to a transcritical bifurcation point (small perturbations of the abundances may drive rare species extinct), implying that some eigenvalues are close to zero, making the inverse overly sensitive to measurement errors. One may try to ignore the rare species from a community model, reasoning that—since they are rare—their impact on the community is slight. But unfortunately, due to the general shape of the species-abundance curve (McGill et al. 2007), there is no natural cutoff point for doing that. Therefore, one should try to approximate properties that do not depend on inverting the community matrix.

One property we have not yet mentioned is feasibility, i.e., whether a community equilibrium has all-positive species abundances. The stability or reactivity of unfeasible equilibria is of no relevance. The reason we did not consider feasibility separately is that we treat the community matrix as the linearization of some arbitrary nonlinear dynamics around some equilibrium state we observe in nature. Its feasibility is therefore already guaranteed by the fact that we are observing the system. Apart from this, note that feasibility is a property that depends on the inverse problem. For instance, in a simple Lotka–Volterra model given by $\mathrm{d}n/\mathrm{d}t = n \circ (b + An)$ (where n is the vector of densities, b the vector of intrinsic growth rates, A the matrix of interaction coefficients, and $\circ$ denotes the Hadamard or element-by-element product), the equilibrium densities are given by $n = -A^{-1}b$. Therefore, as discussed, our method is ill-suited for determining feasibility to begin with.

The accuracy of prediction was dependent on the number of bins k the entries were classified into (fewer bins meant larger errors), and on the choice of the binning resolution b. In practice, a k of seven to nine is probably the largest feasible number, since beyond this it becomes increasingly difficult to assign weak interactions to correct bins. For the binning resolution b, we found that values beyond 10 gave significantly worse results; $b \approx 4$ was usually optimal, but the sensitivity of the results was not very great, and $b = 2$ and $b = 6$ performed similarly (Figs. 2, 4, S4, S9). Interestingly, we have found $b \approx 4$ to perform the best regardless of whether we looked at the empirical data, the randomly generated webs, or the Allometric Trophic Networks, and regardless of whether we estimated stability or reactivity.

It may seem counterintuitive that the smallest value of b is not the most accurate in recovering matrix properties. After all, the finer the resolution, the more accurate the binning. The reason is that, because the number of bins is finite, this finer resolution only applies to larger matrix entries but not necessarily to smaller ones. Consider the following example. A 100×100 matrix is generated by uniformly sampling all but two of its entries from $[-2,\ 2]$, and then setting the remaining two entries to -8 and 8, respectively. If we bin this matrix with $b = 2$ and $k = 5$, the bins are $(-8,\ -4,\ 0,\ 4,\ 8)$. This means that all but the two outliers will be classified in the binned matrix as zero—a crude estimate if there ever was one. However, for $b = 4$ the bins become $(-8,\ -2,\ 0,\ 2,\ 8)$, resolving the underlying data much better. In the end, the best binning is, of course, achieved for $b \to 0$ and

$k \to \infty$. Since this is not feasible in practice, one has to find the best compromise between a b that is not too large but cannot be too small, and a k that is not too small but cannot be too large.

We also checked what happens if, in estimating the strongest interactions p and n, we use 10 empirically measured results and take their average. This way, a single very strong interaction that dominates the system will not artificially distort the binning. However, the results proved insensitive to doing this. The implication is that one should concentrate expensive and time-consuming empirical effort wisely: very accurate measurement of a couple of interaction coefficients does not improve predictive power much, while qualitative knowledge of many matrix entries does.

It is easy to think of scenarios in which the binning procedure produces grossly inaccurate results. For instance, we have seen that matrices with $\mathrm{depn}(A) \gg 1$ will have sensitive spectra, therefore even relatively slight perturbations of their entries (e.g. introduced by binning) may lead to large changes in their eigenvalue structure. We emphasize however that our metric for the scaled departure from normality is merely a proxy for matrix sensitivity: the complete picture is gained by looking at the full pseudospectrum.

But the binning procedure may be inaccurate even when the degree of nonnormality is low. If the interaction strengths span too many orders of magnitude or are dominated by a small number of very large matrix entries, then even a reasonably fine-grained binning structure may classify all but the handful of very large entries as zero, leading to an eigenvalue distribution wildly different from the actual one. The reason why low nonnormality does not matter in this case is that the perturbation induced by binning is in fact very large. In practice however, whenever a few interactions dominate the system, instead of binning, one should concentrate just on those very large entries to gain insight into its workings. In other words, such systems naturally require different methods of analysis than the one presented here, which works better for complex interaction structures where many interactions together shape the properties of the system.

Our results also shed light on the classic stability-complexity debate from a slightly different angle. The original "conventional wisdom" was that more complex systems (i.e., ones with more species, higher variance in the strength of interactions, and higher connectance) would be more stable in the face of perturbations (Begon et al. 2005, p. 586). This view was challenged by the classic result of May (1972) which argued that, all else being equal, more complex systems have a lower probability of being stable: the eigenvalues of the community matrices of more complex systems are less likely to reside in the left half of the complex plane. Whether May's argument actually poses a true challenge to the conventional wisdom has been heavily debated (McCann 2000). Our findings contribute to this debate by showing that large, complex systems, if not necessarily more stable, are very *robust* against perturbations of the community matrix entries. The binning of a matrix can be thought of as a structural perturbation of the system: we are altering the matrix entries, changing the strength of interactions between species. Since, as we have seen, the binning of large complex interaction matrices has only a small effect on the spectrum, the perturbation induced by binning does not have a large effect on the system's large-scale properties. For instance, the stability and reactivity properties were unchanged. This means that if the system was actually stable, it was likely to stay that way, and conversely, unstable systems remained unstable after the perturbation induced by binning. In fact, the robustness interpretation of the stability-complexity debate is closer

to its original formulation, where it was argued that the more pathways there are in an interaction web, the less it matters if one of those links is lost, since other pathways will compensate for the loss (MacArthur 1955, Elton 1958).

In summary, the results from approximating system properties using semiquantitative information point in the direction of using such approximations in practice, where obtaining precise quantitative information is hard, expensive, and time-consuming. The presented results show that such crude parameterization of complex systems may still reveal important global system properties of interest.

Acknowledgements

We thank A. Golubski, J. Grilli, M. J. Michalska-Smith, and E. L. Sander for discussions, and B. Althouse and four anonymous reviewers for their helpful comments. This work was supported by NSF #1148867.

References

Allesina, S., Tang, S., 2012. Stability criteria for complex ecosystems. Nature 483, 205–208.

Aufderheide, H., Rudolf, L., Gross, T., Lafferty, K. D., 2013. How to predict community responses to perturbations in the face of imperfect knowledge and network complexity. Proceedings of the Royal Society of London Series B 280, 20132355.

Bai, Z., Silverstein, J. W., 2009. Spectral analysis of large dimensional random matrices. Springer.

Barabás, G., Pásztor, L., Meszéna, G., Ostling, A., 2014. Sensitivity analysis of coexistence in ecological communities: theory and application. Ecology Letters 17, 1479–1494.

Begon, M., Townsend, C. R., Harper, J. L., 2005. Ecology: From Individuals to Ecosystems. Fourth edition. Blackwell Science Publisher, London.

Bender, E. A., Case, T. J., Gilpin, M. E., 1984. Perturbation experiments in community ecology: Theory and practice. Ecology 65, 1–13.

Berlow, E. L., Dunne, J. A., Martinez, N. D., Stark, P. B., Williams, R. J., Brose, U., 2009. Simple prediction of interaction strengths in complex food webs. Proceedings of the National Academy of Sciences USA 106, 187–191.

Dambacher, J. M., Li, H. W., Rossignol, P. A., 2002. Relevance of community structure in assessing indeterminacy of ecological predictions. Ecology 83, 1372–1385.

Elton, C. S., 1958. The Ecology of Invasions by Animals and Plants. Methuen, London.

Girko, V. L., 1984. The circle law. Theory of Probability and its Applications 29, 694–706.

Levins, R., 1968. Evolution in changing environments. Princeton University Press, Princeton.

Levins, R., 1974. Qualitative analysis of partially specified systems. Ann. NY Acad. Sci. 231, 123–138.

MacArthur, R. H., 1955. Fluctuations of animal populations and a measure of community stability. Ecology 36, 533–536.

May, R. M., 1972. Will a large complex system be stable? Nature 238, 413–414.

May, R. M., 1973. Stability and Complexity in Model Ecosystems. Princeton University Press, Princeton.

McCann, K. S., 2000. The diversity-stability debate. Nature 405, 228–233.

McGill, B. J., Etienne, R. S., Gray, J. S., Alonso, D., Anderson, M. J., Benecha, H. K., Dornelas, M., Enquist, B. J., Green, J. L., He, F., Hurlbert, A. H., Magurran, A. E., Marquet, P. A., Maurer, B. A., Ostling, A., Soykan, C. U., Ugland, K. I., White, E. P., 2007. Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. Ecology Letters 10, 995–1015.

Meszéna, G., Gyllenberg, M., Pásztor, L., Metz, J. A. J., 2006. Competitive exclusion and limiting similarity: a unified theory. Theoretical Population Biology 69, 68–87.

Montoya, J. M., Woodward, G., Emmerson, M. C., Solé, R. V., 2009. Press perturbations and indirect effects in real food webs. Ecology 90, 2426–2433.

Neubert, M. G., Caswell, H., 1997. Alternatives to resilience for measuring the responses of ecological systems to perturbations. Ecology 78, 653–665.

Novak, M., Wootton, J. T., Doak, D. F., Emmerson, M., Estes, J. A., Tinker, M. T., 2011. Predicting community responses to perturbations in the face of imperfect knowledge and network complexity. Ecology 92, 836–846.

Sommers, H. J., Crisanti, A., Sompolinsky, H., Stein, Y., 1998. Spectrum of large random asymmetric matrices. Physical Review Letters 60, 1895–1898.

Tang, S., Allesina, S., 2014. Reactivity and stability of large ecosystems. Frontiers in Ecology and Evolution.
URL `http://dx.doi.org/10.3389/fevo.2014.00021`

Tang, S., Pawar, S., Allesina, S., 2014. Correlation between interaction strengths drives stability in large ecological networks. Ecology Letters 17, 1094–1100.

Tao, T., Vu, V., Krishnapur, M., 2010. Random matrices: Universality of esds and the circular law. Annals of Probability 38 (5), 2023–2065.

Trefethen, L. N., Embree, M., 2005. Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press, New Jersey, USA.

Williams, R. J., Martinez, N. D., 2000. Simple rules yield complex food webs. Nature 404, 180–183.

Woodward, G., Speirs, D. C., Hildrew, A. G., 2005. Quantification and resolution of a complex, size-structured food web. Advances in Ecological Research 36, 85–135.

Yodzis, P., 1988. The indeterminacy of ecological interactions as perceived through perturbation experiments. Ecology 69, 508–515.

Yodzis, P., 2000. Diffuse effects in food webs. Ecology 81, 261–266.

Predicting global community properties from uncertain estimates of interaction strengths

Supporting Information

György Barabás & Stefano Allesina

Abstract

This Supporting Information presents a step-by-step description of our methodology for generating random interaction matrices and Allometric Trophic Networks (Section 1), shows further results obtained from our simulations (Section 2), and gives extra details on the theoretical background of matrix binning (Section 3). We also describe a simple algorithm for obtaining the ε-pseudospectral contour lines of any given matrix (Section 4). Finally, we derive a new metric for assessing the degree of nonnormality of a matrix, which is useful in assessing how sensitive its spectrum is expected to be to perturbing its entries (Section 5).

Contents

1 More details on the simulation methods

1.1 Random interaction matrices

Here is a step-by-step breakdown of how we created our random interaction matrices.

1. We first pick a species richness S and a connectance C (i.e., the fraction of nonzero interactions in the matrix). In our simulations S was either 50, 100, 250, or 500, and C was chosen from 0.1, 0.25, 0.5, or 1.

2. Consider the following procedure for generating a probability distribution. First pick a distribution shape from two options: either lognormal or Gamma. Then determine the given distribution's mean by sampling it uniformly from $[0.1,\ 10]$. Finally, the standard deviation is also sampled uniformly, from $[1,\ 10]$.

3. Use this procedure to generate four separate probability distributions. Call the first one "PredPrey" for predator-prey, the second one "Mut" for mutualism, the third "Comp" for competition, and the last one "Diag" for diagonal.

4. How do we decide whether to pick a lognormal or a Gamma distribution? We actually repeated every simulation with all 16 possible combinations of the distributions, with "PredPrey", "Mut", "Comp", and "Diag" all taking on both possible values, so nothing was left out.

5. Generate two random numbers, both of them sampled from $[0,\ 1]$. Call them C_t and C_m (for "trophic" and "mutualistic").

6. Now generate an $S \times S$ matrix A of all zeros.

7. Make a fraction $C_t C$ of the offdiagonal entries predator-prey, a fraction $C_m C$ mutualistic, and a fraction $C(1 - C_t - C_m)$ competitive (leave the rest of the $1 - C$ entries as zeros). The predator-prey entries are drawn from "PredPrey", the mutualistic ones from "Mut", and the competitive ones from "Comp". For predator-prey interactions, make sure that the (i, j)th and (j, i)th entries of A have the opposite sign; for mutualism, both should be positive, and for competition, both negative.

8. Finally, the diagonal of the matrix is generated: fill out the diagonal entries by sampling from the distribution "Diag". Multiply the diagonal by -1 to make its entries negative.

9. Note: one could introduce a "diagonal connectance" C_d, i.e., the fraction of nonzero diagonal entries. However, in our simulations we always set this to 1, so no diagonal entries were left equal to zero.

Here is a function, written in R (R Development Core Team 2008), that we used to implement this procedure:

```
# This function generates a random interaction matrix with a prescribed
# proportion of trophic, mutualistic, and competitive interactions.
#
# Input
# -----
# S: number of species.
# C: connectance (fraction of nonzero interactions); between 0 and 1.
# ppred: probability that a nonzero link is predator-prey (between 0 and 1).
# pcomp: probability that a nonzero non-predator-prey link is competitive
#        (between 0 and 1); the rest are mutualistic.
# conv: conversion efficiency of predators; realistically between 0 and 1.
# pdist, mdist, cdist: distribution from which trophic, mutualistic, and
#        competitive entries are drawn. Values: 1 = lognormal; 2 = Gamma.
# mp, sp, mm, sm, mc, sc: sample means and standard deviations for each
#        distribution. The (mp, sp) are mean and std dev for predator-prey,
#        (mm, sm) are for mutualism, and (mc, sc) for competition.
# Cd: fraction of diagonal entries that are nonzero.
# ddist: distribution from which nonzero diagonal entries are drawn.
#        Values: 1 = lognormal; 2 = Gamma. Note: the diagonal entries are
#        never positive (i.e., the values drawn from the chosen distribution
#        get multiplied by -1).
# md, sd: sample mean and standard deviation of nonzero diagonal entries.
#
# Output
# ------
# An S-by-S matrix
#
GenerateMatrix <- function(S, C, ppred, pcomp, conv, pdist, mdist, cdist,
                           mp, sp, mm, sm, mc, sc, Cd, ddist, md, sd) {
    # Generate adjacency matrix. Only upper triangle is generated; the
    # diagonal and lower triangular part are discarded. Entry is 0 for
    # absence and 1 for presence of interaction.
    A <- matrix(sample(0:1, S*S, replace=TRUE, prob=c(1-C, C)), nrow=S, ncol=S)
    A[lower.tri(A)] <- 0
    diag(A) <- 0
    # Now set trophic (= 1) vs nontrophic (= 2) interactions:
    A <- A * matrix(sample(1:2, S*S, replace=TRUE, prob=c(ppred, 1-ppred)),
                    nrow=S, ncol=S)

    # Create predator-prey adjacency matrix AP
    AP <- A
    AP[AP!=1] <- 0 # discard nontrophic (= 2) interactions
    AP <- AP * matrix(sign(rnorm(S*S)), nrow=S, ncol=S) # either -1 or 1
    AP <- AP - t(AP)
```

```
# Create mutualistic and competitive adjacency matrices AM and AC
AM <- A
AM[AM!=2] <- 0
AM[AM==2] <- 1
AM <- AM * matrix(sample(c(-1, 1), S*S, replace=TRUE,
                         prob=c(pcomp, 1-pcomp)), nrow=S, ncol=S)
AM <- AM + t(AM)
AC <- AM
AM[AM==(-1)] <- 0 # mutualism
AC[AC==1] <- 0 # competition

# Predator-prey values
AP[AP==1] <- conv * AP[AP==1] # conversion efficiencies
if (pdist==1) {
    p1 <- log(mp) - log(1 + sp^2/mp^2)/2
    p2 <- sqrt(log(1 + sp^2/mp^2))
    AP <- AP * matrix(rlnorm(S*S, p1, p2), nrow=S, ncol=S)
}
if (pdist==2) {
    p1 <- mp^2 / sp^2
    p2 <- sp^2 / mp
    AP <- AP * matrix(rgamma(S*S, shape=p1, scale=p2), nrow=S, ncol=S)
}

# Mutualism values
if (mdist==1) {
    p1 <- log(mm) - log(1 + sm^2/mm^2)/2
    p2 <- sqrt(log(1 + sm^2/mm^2))
    AM <- AM * matrix(rlnorm(S*S, p1, p2), nrow=S, ncol=S)
}
if (mdist==2) {
    p1 <- mm^2 / sm^2
    p2 <- sm^2 / mm
    AM <- AM * matrix(rgamma(S*S, shape=p1, scale=p2), nrow=S, ncol=S)
}

# Competition values
if (cdist==1) {
    p1 <- log(mc) - log(1 + sc^2/mc^2)/2
    p2 <- sqrt(log(1 + sc^2/mc^2))
    AC <- AC * matrix(rlnorm(S*S, p1, p2), nrow=S, ncol=S)
}
if (cdist==2) {
    p1 <- mc^2 / sc^2
    p2 <- sc^2 / mc
    AC <- AC * matrix(rgamma(S*S, shape=p1, scale=p2), nrow=S, ncol=S)
```

```
    }

    # Finally, create the diagonal of the matrix
    diagvec <- sample(0:1, S, replace=TRUE, prob=c(1-Cd, Cd))
    if (ddist==1) {
        p1 <- log(md) - log(1 + sd^2/md^2)/2
        p2 <- sqrt(log(1 + sd^2/md^2))
        diagvec <- diagvec * rlnorm(S, p1, p2)
    }
    if (ddist==2) {
        p1 <- md^2 / sd^2
        p2 <- sd^2 / md
        diagvec <- diagvec * rgamma(S, shape=p1, scale=p2)
    }

    # The sum of all components is the full interaction matrix
    return(AP + AM + AC - diag(diagvec))
}
```

1.2 The Allometric Trophic Network model

To generate community matrices using the Allometric Trophic Network model, we followed the equations and parameterization described by Berlow et al. (2009). First, a food web's adjacency matrix w_{ij} is created using the niche model (Williams and Martinez 2000; see below), where w_{ij} is equal to 1 if species i eats species j and to 0 otherwise. The dynamical equations of the model read

$$\frac{\mathrm{d}B_i}{\mathrm{d}t} = r_i G_i(N_1,N_2) B_i - x_i B_i + \sum_{j=1}^{S} w_{ij} x_i y B_i F_{ij} - \sum_{j=1}^{S} w_{ji} x_j y B_j F_{ji}/e_{ji}, \tag{1}$$

where B_i is the biomass of species i, r_i is a nutrient uptake-dependent maximum growth rate, $G_i(N_1,N_2)$ is the growth achieved on the two nutrients N_1 and N_2, x_i is a metabolic rate, y is the maximum consumption rate of consumers relative to their metabolic rate, e_{ij} is species i's assimilation efficiency when eating species j, and F_{ij} is a generalized functional response given by

$$F_{ij} = \frac{\omega_{ij} B_j^h}{B_0^h + cB_i B_0^h + \sum_{k=1}^{S} w_{ik}\omega_{ik} B_k^h}. \tag{2}$$

In this functional response, ω_{ij} is the proportion of i's maximum consumption rate targeted at consuming j, B_0 is a half-saturation density, h is an exponent determining the Holling type of the functional response, and c is a predator interference parameter.

We assume there are two limiting nutrients N_1 and N_2 which the primary producers can consume. Their dynamics is given by

$$\frac{\mathrm{d}N_i}{\mathrm{d}t} = D(s_i - N_i) - \sum_{j=1}^{S} c_{ij} r_i G_i(N_1, N_2) B_i, \tag{3}$$

where N_i is the concentration of nutrient i, D is a turnover rate, c_{ij} is the content of nutrient i in the biomass of species j, s_i is nutrient i's supply rate, and $G_i(N_1, N_2)$ is the same function that appeared in Eq. (1). It is defined by

$$G_i(N_1, N_2) = \min\left(\frac{N_1}{K_{1i} + N_1},\ \frac{N_2}{K_{2i} + N_2}\right), \tag{4}$$

where "min" picks the smaller of the two arguments, and K_{ij} is species j's half saturation density for nutrient i.

The parameters were assigned as follows.

1. $S = 50$ for the initial number of species.
2. Out of these 50, the number of basal species (primary producers) was sampled uniformly as an integer from $[2,\ 10]$. In our ordering of species, they are the last ones, so if there are two producers, they will be species 49 and 50.
3. The food web adjacency matrix w_{ij} (equal to 1 if species i eats species j and to 0 otherwise) is generated by the niche model (Williams and Martinez 2000), in the following way.
   - Primary producers only consume nutrients, therefore $w_{ij} = 0$ for all i that are primary producers.
   - Each consumer species i consumes a range of other species $j, j+1, \ldots, j+k$, where the starting index j is uniformly sampled between $i+1$ and S, and the length k is an integer sampled from $[-(S-i+1)/3,\ (S-i+1)/3]$.
   - Note: if $j+k$ turns out to be larger than S, it is set equal to S, and if it turns out less than $i+1$, it is set to $i+1$.
4. $\omega_{ij} = 1/\sum_{j=1}^{S} w_{ij}$ if the sum in the denominator is nonzero; otherwise, $\omega_{ij} = 0$.
5. $r_i = 1$ for species i that are primary producers and 0 otherwise.
6. Determine the trophic level T of each species. Trophic level can be calculated by defining the matrix $A_{ij} = w_{ij}/\sum_{k=1}^{S} w_{ik}$ and the vector $u_i = 1$ (for all i); we then have $T_i = \sum_{j=1}^{S} (I-A)^{-1}_{ij} u_j$, where I is the $S \times S$ identity matrix and $(I-A)^{-1}_{ij}$ is the (i,j)th entry of the inverse of the matrix $I - A$.

7. Define a vector Z of predator-prey body mass ratios; its S entries are sampled from a lognormal distribution with mean 10 and variance 100.

8. Define the vector of body masses. For predators, the body masses are given by the formula $M_i = Z_i^{T_i - 1}$. For species that are not predators, the body masses are equal to 1.

9. Using this, the metabolic rates x_i are calculated as $x_i = (a_x/a_r)\, M_i^{-0.25}$, where M_i is species i's body mass, and (a_x/a_r) is equal to 0.138 for primary producers (i.e., species with trophic level exactly equal to one), and to 0.314 otherwise.

10. Assimilation efficiencies: $e_{ij} = 0.45$ if the trophic level of species i is exactly 2; otherwise, $e_{ij} = 0.85$.

11. The K_{ij} are uniformly and independently sampled from $[0.1,\ 0.2]$.

12. $c_{1i} = 1$ and $c_{2i} = 0.5$ for all i, $y = 8$, $B_0 = 0.5$, $D = 0.25$, $h = 2$, and $s_i = 1$ for both nutrients.

13. Initial conditions: all the B_i and N_i at $t = 0$ are uniformly sampled from $[0.05,\ 0.2]$.

With this parameterization, the model equations Eq. (1) can be integrated with any reputable algorithm for solving systems of ordinary differential equations. We used the `NDSolve` routine implemented in Mathematica (Wolfram Research Inc. 2014).

For each of our $10,000$ random parameterizations, the equations were solved until they reached a fixed point. At that point, extinct species were removed, and the Jacobian of the nonextinct part of the system was evaluated, yielding the community matrix.

2 Further results on our simulated matrices

The first thing to note is the overall insensitivity of the binned community matrices to the misclassification rate (i.e., the probability per entry of the community matrix that it gets classified into an incorrect bin). In the main text we always show results with 10% misclassification. Below, we show the same results with 0% and 20% misclassification rates. For the leading eigenvalues (related to stability), we have Figs. S1 and S2 for randomly generated matrices, and Figs. S6 and S7 for Allometric Trophic Networks (Berlow et al. 2009). For the leading eigenvalues of the Hermitian parts (related to reactivity), we have Figs. S3-S5 (randomly generated matrices) and Figs. S8-S10 (Allometric Trophic Networks); here the results with 10% misclassification rate are also included, as these were referred to but not shown in the main text. As can be seen, the sensitivity to an increased misclassification rate is small, though in general it does lead to lower overall prediction accuracy.

We also looked at the effects of connectance on prediction accuracy, both for stability and reactivity (Figs. S11-S18); the effects, however, do not appear systematic. For instance, for binning constant $b = 4$ and number of bins $k = 7$, increasing connectance improves accuracy when the number of species is $S = 50$, but leads to worse accuracy for $S = 500$. This trend is apparent both in the case of stability (Figs. S11, S14) and of reactivity (Figs. S15, S18).

For matrices whose leading eigenvalues are close to the imaginary axis relative to the total spread of the eigenvalues, stability may be incorrectly predicted even if the leading eigenvalue of the binned matrix is not very different from that of the original one. To see how often this would happen, we first take all those matrices for which $|r_A|/\sigma_A < 0.05$, i.e., the magnitude of the leading eigenvalue is less than one twentieth of the total spread of the real parts of all the eigenvalues. The result depends heavily on the binning resolution b, and especially on the number of bins k (Fig. S19). For $k = 3$ (3 bins), stability is more likely to be misclassified than correctly predicted. However, for $b \geq 4$ and $k \geq 5$, stability is correctly predicted in the majority of cases (for instance, exactly 90% of the time when $b = 4$, $k = 7$). If we are more inclusive and consider all those matrices for which $|r_A|/\sigma_A < 0.1$, this number goes up to 97% (Fig. S20).

The same can be done with reactivity (Figs. S21, S22), where we look at the eigenvalues of the Hermitian parts $H(A)$, $H(B)$ of the original and binned matrices A and B. For $k = 7$ bins and binning resolution $b = 4$, reactivity is correctly predicted 92% of the time for matrices with $|r_{H(A)}|/\sigma_{H(A)} < 0.05$, and 99.4% of the time for matrices with $|r_{H(A)}|/\sigma_{H(A)} < 0.1$.

Finally, we show that our results do not depend strongly on the particular form of the probability distribution from which the entries of the random matrices are drawn. We get the same qualitative results if we draw the entries from a lognormal distribution as if we draw them from a Gamma distribution, both for stability (Figs. S23, S24) and for reactivity (Figs. S25, S26).

3 Random matrices

Here we describe in general how random matrices are binned when their entries are drawn independently from the same underlying probability distribution $p_A(x)$. Imagine we are given a binning scheme with k bins whose values are $(x_1, x_2, \ldots, x_k)$. The binned distribution $p_B(x)$ is then given by

$$p_B(x) = \sum_{i=1}^{k} w_i \delta(x - x_i), \tag{5}$$

where $\delta(x - x_i)$ is the Dirac delta function (which can be though of as a normal distribution with mean x_i and zero variance), and w_i is the probability that a given matrix entry will get classified into the ith bin; $\sum_{i=1}^{k} w_i = 1$. This distribution is well-normalized, and its mth moment μ_m has a very simple form:

$$\mu_m = \int_{-\infty}^{+\infty} x^m p_B(x)\,\mathrm{d}x = \sum_{i=1}^{k} w_i \int_{-\infty}^{+\infty} x^m \delta(x - x_i)\,\mathrm{d}x = \sum_{i=1}^{k} w_i x_i^m, \tag{6}$$

where we used the property of the Dirac delta function that $\int_{-\infty}^{+\infty} f(x)\delta(x - x_0)\,\mathrm{d}x = f(x_0)$ for any function $f(x)$. The w_i are calculated from $p_A(x)$ using the criterion that each entry should be lumped into the bin with the closest value. Let us define corresponding integration limits $\Omega_0, \Omega_1, \ldots, \Omega_k$, with $\Omega_0 = -\infty$, $\Omega_k = +\infty$, and $\Omega_i = (x_i + x_{i+1})/2$ for $1 \leq i < k$. We then have

$$w_i = \int_{\Omega_{i-1}}^{\Omega_i} p_A(x)\,\mathrm{d}x. \tag{7}$$

After the binning has been obtained, the eigenvalue distribution can be calculated based on the circular law and its extensions (Ginibre 1965, Girko 1984, Tao et al. 2010, Allesina and Tang 2012). If one is interested in the leading eigenvalue, it can be obtained (Tang et al. 2014) from the formula

$$r = \max\left\{\sqrt{SV} - E + d,\ (S-1)E + d\right\}, \tag{8}$$

where S is the number of rows/columns of the matrix, E is the expected value of $p_A(x)$, V is its variance, and d is a possible constant that has been added to the diagonal of the matrix.

The example given in the main text involves the uniform distribution $p_A(x) = \mathcal{U}[-1,1]$, having mean $E_A = 0$, variance $V_A = 1/3$, and $d = 0$. The leading eigenvalue r_A is therefore simply given by $\sqrt{SV_A}$, or $\sqrt{S/3}$. We bin this matrix with three bins $(-1,\ 0,\ 1)$. The probability of a given $x \in \mathcal{U}[-1,1]$ falling into the -1 or 1 bins is $1/4$, while the probability of falling into the 0 bin is $1/2$. We therefore have $w_1 = 1/4$, $w_2 = 1/2$, and $w_3 = 1/4$. The binned distribution $p_B(x)$ therefore reads

$$p_B(x) = \sum_{i=1}^{3} w_i \delta(x - x_i) = \frac{\delta(x+1) + 2\delta(x) + \delta(x-1)}{4}. \tag{9}$$

The mean of $p_B(x)$ is

$$E_B = \int_{-\infty}^{+\infty} x\, p_B(x)\, \mathrm{d}x = \frac{1}{4}\int_{-\infty}^{+\infty} x\left[\delta(x+1) + 2\delta(x) + \delta(x-1)\right] \mathrm{d}x = \frac{1}{4}\left(1 + 2\cdot 0 + (-1)\right) = 0, \tag{10}$$

while its variance is

$$\begin{aligned} V_B &= \int_{-\infty}^{+\infty} (x - E_B)^2\, p_B(x)\, \mathrm{d}x = \int_{-\infty}^{+\infty} x^2 p_B(x)\, \mathrm{d}x \\ &= \frac{1}{4}\int_{-\infty}^{+\infty} x^2\left[\delta(x+1) + 2\delta(x) + \delta(x-1)\right] \mathrm{d}x = \frac{1}{4}\left(1^2 + 2\cdot 0^2 + (-1)^2\right) = \frac{1}{2}. \end{aligned} \tag{11}$$

The leading eigenvalue r_B of the binned matrix is then given by Eq. (8), which in this case simplifies to $r_B = \sqrt{SV_B} = \sqrt{S/2}$. The ratio of r_B and r_A is then

$$\frac{r_B}{r_A} = \frac{\sqrt{S/2}}{\sqrt{S/3}} = \sqrt{\frac{3}{2}} \approx 1.22. \tag{12}$$

The same calculation and comparison can be made with a different binning scheme involving five bins $(-1,\ -1/2,\ 0,\ 1/2,\ 1)$. Using Eq. (7), we obtain $w_i = (1,\ 2,\ 2,\ 2,\ 1)/8$, therefore $p_B(x)$ is given by

$$p_B(x) = \sum_{i=1}^{5} w_i \delta(x - x_i) = \frac{\delta(x+1) + 2\delta(x+1/2) + 2\delta(x) + 2\delta(x-1/2) + \delta(x-1)}{8}. \tag{13}$$

The mean is again

$$E_B = \int_{-\infty}^{+\infty} x\, p_B(x)\, \mathrm{d}x = 0, \tag{14}$$

but the variance now reads

$$\begin{aligned} V_B &= \int_{-\infty}^{+\infty} (x-E_B)^2\, p_B(x)\,\mathrm{d}x = \int_{-\infty}^{+\infty} x^2 p_B(x)\,\mathrm{d}x \\ &= \frac{1}{8}\int_{-\infty}^{+\infty} x^2\left[\delta(x+1)+2\delta(x+1/2)+2\delta(x)+2\delta(x-1/2)+\delta(x-1)\right]\mathrm{d}x = \frac{3}{8}. \end{aligned} \tag{15}$$

The new ratio of the leading eigenvalues, based on Eq. (8), is

$$\frac{r_B}{r_A} = \frac{\sqrt{3S/8}}{\sqrt{S/3}} = \sqrt{\frac{9}{8}} \approx 1.06. \tag{16}$$

All the above formulas assumed that the matrix entries were all independently drawn from the same distribution. The same approach may in principle be extended to different types of random matrices. For instance, the elliptic law (Sommers et al. 1998, Allesina and Tang 2012, Tang et al. 2014) concerns the eigenvalues of matrices where all symmetric pairs of entries (A_{ij},A_{ji}) are sampled from some bivariate probability distribution with identically distributed marginals and correlation ρ. The eigenvalues are then uniformly distributed within an ellipse of horizontal and vertical semi-axes $\sqrt{SV}(1+\rho)$ and $\sqrt{SV}(1-\rho)$, respectively. The leading eigenvalue r is then equal to

$$r = \max\left\{(1+\rho)\sqrt{SV}-E+d,\ (S-1)E+d\right\}, \tag{17}$$

where d is the mean of the diagonal entries, $E=\overline{A_{ij}}$ is the expected value of the offdiagonal entries, $V=\mathrm{Var}(A_{ij})$ is their variance, and

$$\rho = \frac{\overline{A_{ij}A_{ji}}-E^2}{V} \tag{18}$$

is the expected correlation between symmetric pairs of (offdiagonal) entries. To determine the ratio of the leading eigenvalues of the original matrix A and the binned matrix B, one simply calculates E_A, V_A, ρ_A and E_B, V_B, ρ_B, applies Eq. (17) to obtain both r_A and r_B, and calculates their ratio. There is one difference in how $p_B(x,y)$ is calculated, however: the integrals in Eq. (7) now become double integrals, going over both directions of the bivariate distribution $p_A(x,y)$.

4 Calculating pseudospectral contour lines

There are several methods for computing the pseudospectra of matrices (Trefethen and Embree 2005, chapters 39-44), each adapted to different situations. Here we only describe the simplest algorithm (Trefethen and Embree 2005, p. 371), which we used to generate the pseudospectral contour plots of the main text.

The ε-pseudospectrum of a matrix A is defined as the set of complex numbers that are eigenvalues of all possible perturbed matrices $A+P$ with $\|P\|<\varepsilon$, where the matrix norm $\|\cdot\|$ is defined as $\|P\|=\sqrt{\lambda_{\max}(P^*P)}$, with P^* being the conjugate transpose of P, and $\lambda_{\max}(P^*P)$ the largest eigenvalue of P^*P. An equivalent definition uses the concept of singular values (the singular values

of a matrix A are the square roots of the eigenvalues of A^*A): the ε-pseudospectrum of A is the set of complex numbers z such that the smallest singular value of the matrix $zI - A$ is smaller than ε, where I is the identity matrix (Trefethen and Embree 2005, pp. 16-17). The simplest way of obtaining pseudospectra is to compute the smallest singular values of $zI - A$ on a regular grid of points in the complex plane and then visualize this information via a contour plot.

If the matrix to be analyzed is given by `A`, the grid consists of `m` points in both the real and imaginary directions, and the starting and endpoints along the real and imaginary directions are given, respectively, by `reMin`, `reMax`, `imMin`, and `imMax`, the following R code (R Development Core Team 2008) computes the minimum singular values on the grid and stores them in the matrix `sigmin`:

```
x <- seq(from=reMin, to=reMax, len=m)
y <- seq(from=imMin, to=imMax, len=m)
sigmin <- matrix(0, m, m)
for (k in 1:m) {
  for (j in 1:m) {
    sigmin[j,k] <- min(svd((x[k]+y[j]*1i)*diag(ncol(A))-A, nu=0, nv=0)$d)
  }
}
```

This information may then be visualized by, say, invoking the `contour` function. Note that there are more efficient algorithms available; see the book by Trefethen and Embree (2005, p. 375) for one that is much faster but still reasonably simple.

5 Departure from normality

Despite the avaliability of algorithms for obtaining pseudospectral contour lines, the computation itself can be quite time-intensive, especially when S is large and the resolution of the grid on the complex plane is high. It would be good to have a way of assessing, before performing the computations, whether one expects large or small pseudospectral regions. Since normal matrices possess the smallest possible pseudospectra (Trefethen and Embree 2005, Theorem 2.2), a natural way to see if the spectrum of a matrix would be sensitive to perturbations is to measure the matrix's departure from normality. Any such metric is bound to be approximate, because nonnormality is a complex property that cannot be encompassed into a single number (Trefethen and Embree 2005, p. 446). Nevertheless, such measures do provide a quick and useful way of checking whether we expect a matrix to be oversensitive to perturbations.

One commonly used measure of the departure from normality (Henrici 1962, Lee 1996) is

$$\mathrm{dep}(A) = \sqrt{\sum_{i=1}^{S} \sigma_i^2 - \sum_{i=1}^{S} |\lambda_i|^2}, \tag{19}$$

where λ_i and σ_i are the ith eigenvalue and singular value of A, respectively (singular values are the square roots of the eigenvalues of A^*A). This metric is equal to zero for normal matrices, and can get arbitrarily large for highly nonnormal ones. As there is no upper limit to $\mathrm{dep}(A)$, it is not immediately clear whether any given nonzero value of this metric should be considered large or small.

One way of finding a point of comparison is to calculate $\mathrm{dep}(M)$ for a simple $S \times S$ random matrix M, where each entry is drawn independently from the same distribution with mean zero and variance V. The reason is that such matrices are known to possess only a very mild degree of nonnormality, being fairly robust to perturbations of their entries (Edelman 1988). Any matrix whose departure from normality is comparable to or smaller than that of the aforementioned random matrix is therefore likely to possess a spectrum that does not change much in response to binning. Note that the eigenvalue and singular value distributions of these random matrices do not depend on the shape of the probability distribution from which their entries are drawn, merely its variance.

The departure from normality of the random matrix described above can be derived analytically. The probability distribution $p_\sigma(x)$ of their singular values follows

$$p_\sigma(x) = \frac{2}{\pi\sqrt{SV}}\sqrt{1-\frac{x^2}{4SV}} \tag{20}$$

for $0 \leq x \leq 2\sqrt{SV}$ and 0 otherwise (Marčenko and Pastur 1967). In turn, the probability distribution of the absolute values of their eigenvalues $p_\lambda(x)$, as a function of the distance x from the origin of the complex plane, follows

$$p_\lambda(x) = \frac{2x}{SV} \tag{21}$$

for $0 \leq x \leq \sqrt{SV}$ and 0 otherwise (Tao et al. 2010, Allesina and Tang 2012).

Eq. (19) instructs us to sum the squared singular values and eigenvalues. We therefore transform $p_\sigma(x)$ and $p_\lambda(x)$ by introducing the new variable $y = x^2$, from which $x = \sqrt{y}$ and $|\mathrm{d}x/\mathrm{d}y| = \left(2\sqrt{y}\right)^{-1}$. The transformed probability distribution $g_\sigma(y)$ of singular values then reads

$$g_\sigma(y) = p_\sigma\left(x(y)\right)\left|\frac{\mathrm{d}x}{\mathrm{d}y}\right| = \frac{1}{\pi\sqrt{SVy}}\sqrt{1-\frac{y}{4SV}} \tag{22}$$

for $0 \leq y \leq 4SV$ and 0 otherwise. In turn, the transformed probability distribution $g_\lambda(y)$ of the absolute eigenvalues reads

$$g_\lambda(y) = p_\lambda\left(x(y)\right)\left|\frac{\mathrm{d}x}{\mathrm{d}y}\right| = \frac{1}{SV} \tag{23}$$

for $0 \leq y \leq SV$ and 0 otherwise.

We can now evaluate Eq. (19) by approximating the sums with integrals (this assumes S is large). The sum of all the squared singluar values is equal to S times their mean:

$$\sum_{i=1}^{S}\sigma_i^2 = S\int_0^{4SV} y\,g_\sigma(y)\,\mathrm{d}y = S\int_0^{4SV}\frac{y}{\pi\sqrt{SVy}}\sqrt{1-\frac{y}{4SV}} = S^2V. \tag{24}$$

The sum of the absolute squares of all eigenvalues is again S times their mean:

$$\sum_{i=1}^{S} |\lambda_i|^2 = S \int_0^{SV} y\, g_\lambda(y)\, \mathrm{d}y = S \int_0^{SV} \frac{y}{SV}\, \mathrm{d}y = \frac{S^2 V}{2}. \tag{25}$$

For a random matrix M, Eq. (19) then reads

$$\mathrm{dep}(M) = \sqrt{\sum_{i=1}^{S} \sigma_i^2 - \sum_{i=1}^{S} |\lambda_i|^2} = \sqrt{S^2 V - \frac{S^2 V}{2}} = S\sqrt{\frac{V}{2}}. \tag{26}$$

We can now introduce a measure of the *scaled departure from normality* $\mathrm{depn}(A)$, comparing the degree of nonnormality with that of a random matrix which is of the same size and its entries are of the same variance as the original matrix. Take an $S \times S$ matrix A and determine the sample variance V of all its entries. The metric then reads

$$\mathrm{depn}(A) = \frac{\mathrm{dep}(A)}{S\sqrt{V/2}}. \tag{27}$$

Whenever this metric is equal to one, the degree of nonnormality of A is that of a random matrix.

The following R function (R Development Core Team 2008) calculates this metric, given that the sample variance of the entries of A is nonzero:

```
depn <- function(A) {
  eA <- eigen(A, only.values=TRUE)$values
  sA <- svd(A, nu=0, nv=0)$d
  depA <- sqrt(sum(sA^2) - sum(abs(eA)^2))
  S <- nrow(A)
  V <- var(as.vector(A))
  depRand <- S*sqrt(V/2)
  return(depA/depRand)
}
```

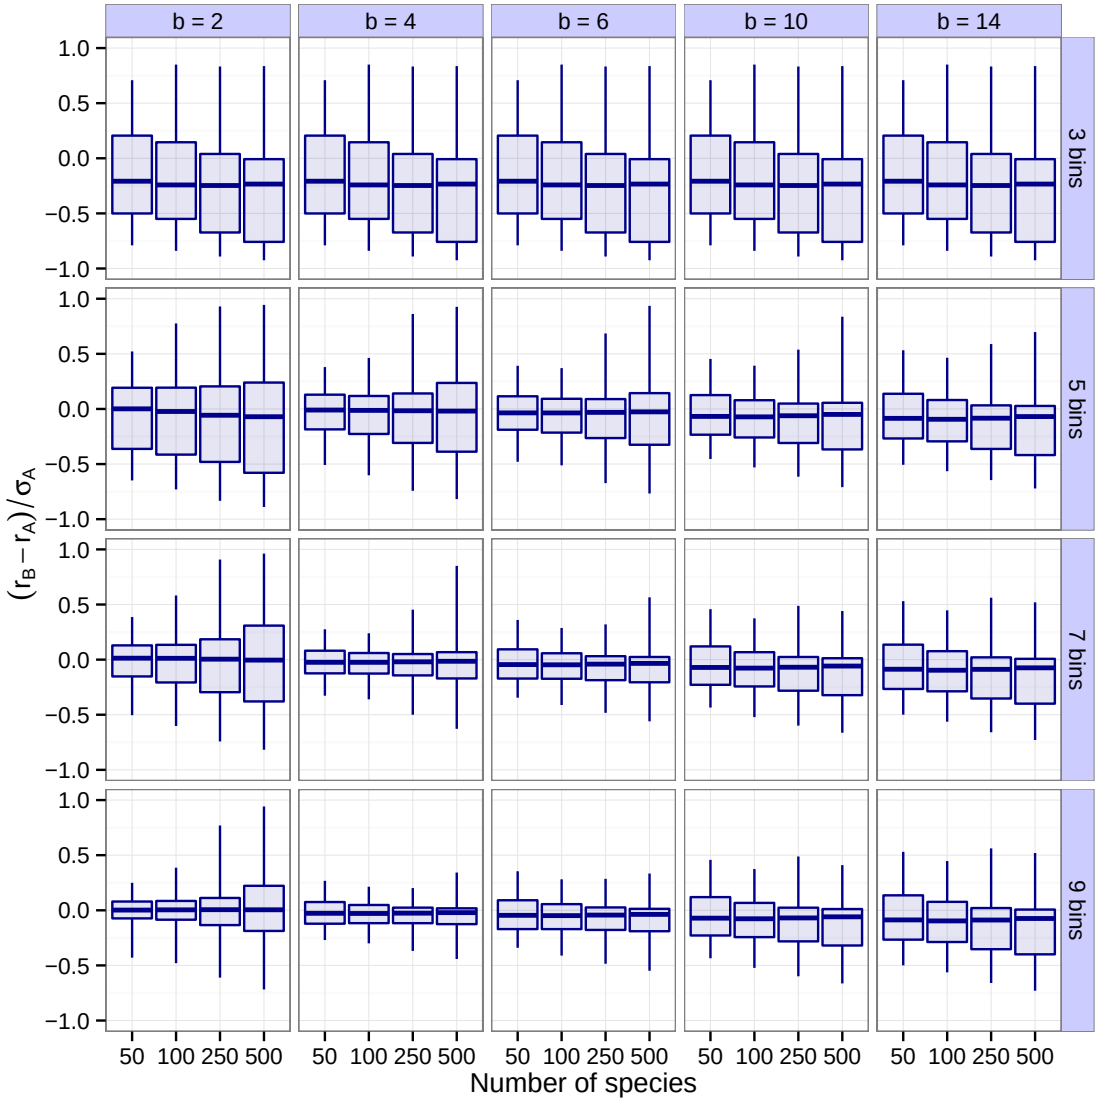

Figure S1: Box plots of how the leading eigenvalues of randomly generated community matrices A are captured by those of their binned counterparts B. Each matrix is binned with no misclassification. Rows correspond to different values of the binning resolution b; columns to different numbers of bins k. The data in each panel are separated based on the number of species. The ordinate of the panels shows the difference between the leading eigenvalue r_A of the original and r_B of the binned matrices relative to σ_A, the total range of the real parts of the original matrix's eigenvalues. Interpretation of the box plots: median (lines), 5% to 95% quantiles (boxes; note that they encompass 90% instead of the usual 50% of the data), and ranges (whiskers).

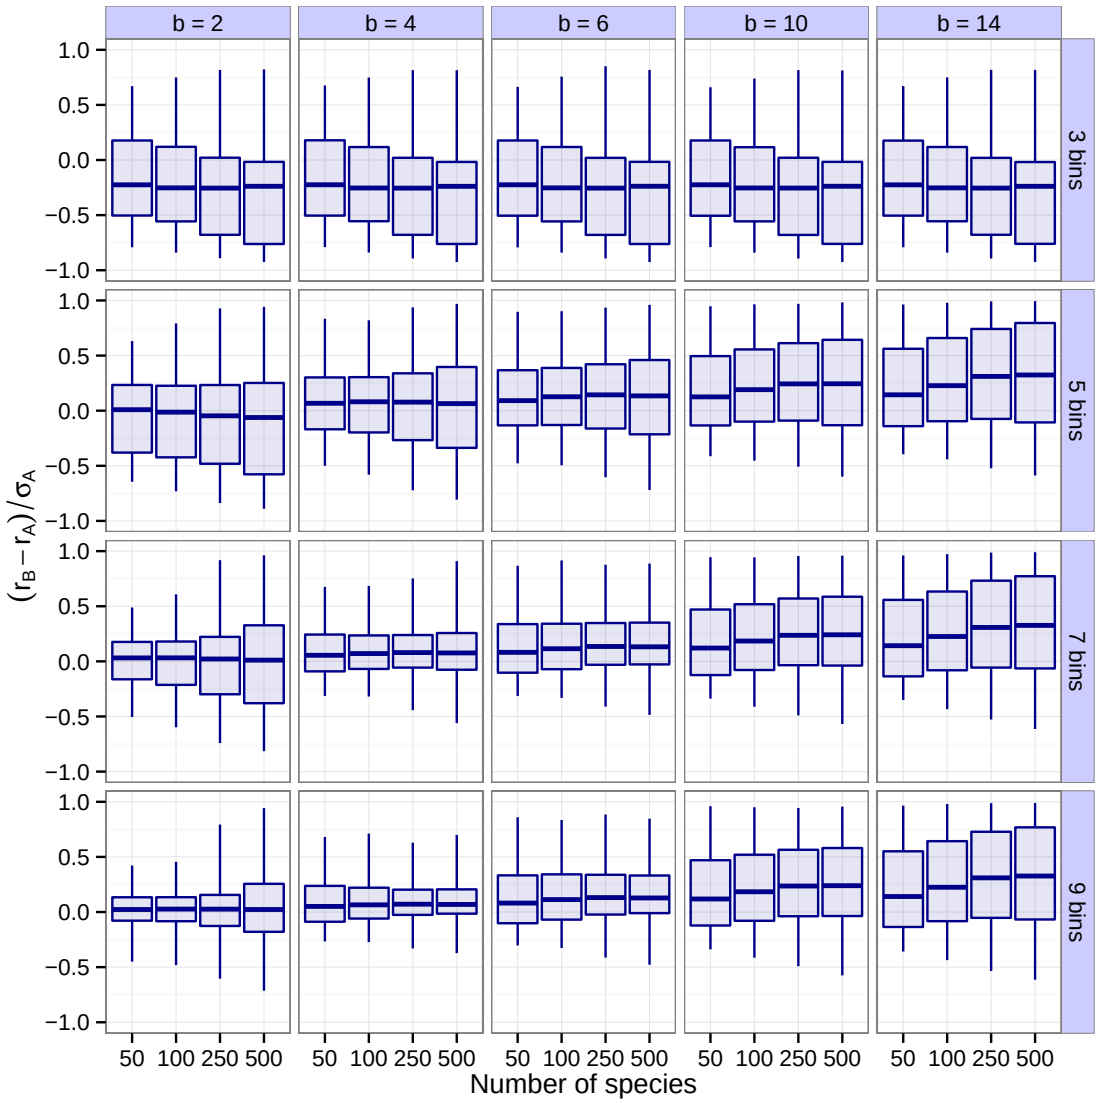

Figure S2: Box plots of how the leading eigenvalues of randomly generated community matrices A are captured by those of their binned counterparts B. Each matrix is binned with a 20% misclassification rate. Otherwise, the plot is organized just like Fig. S1.

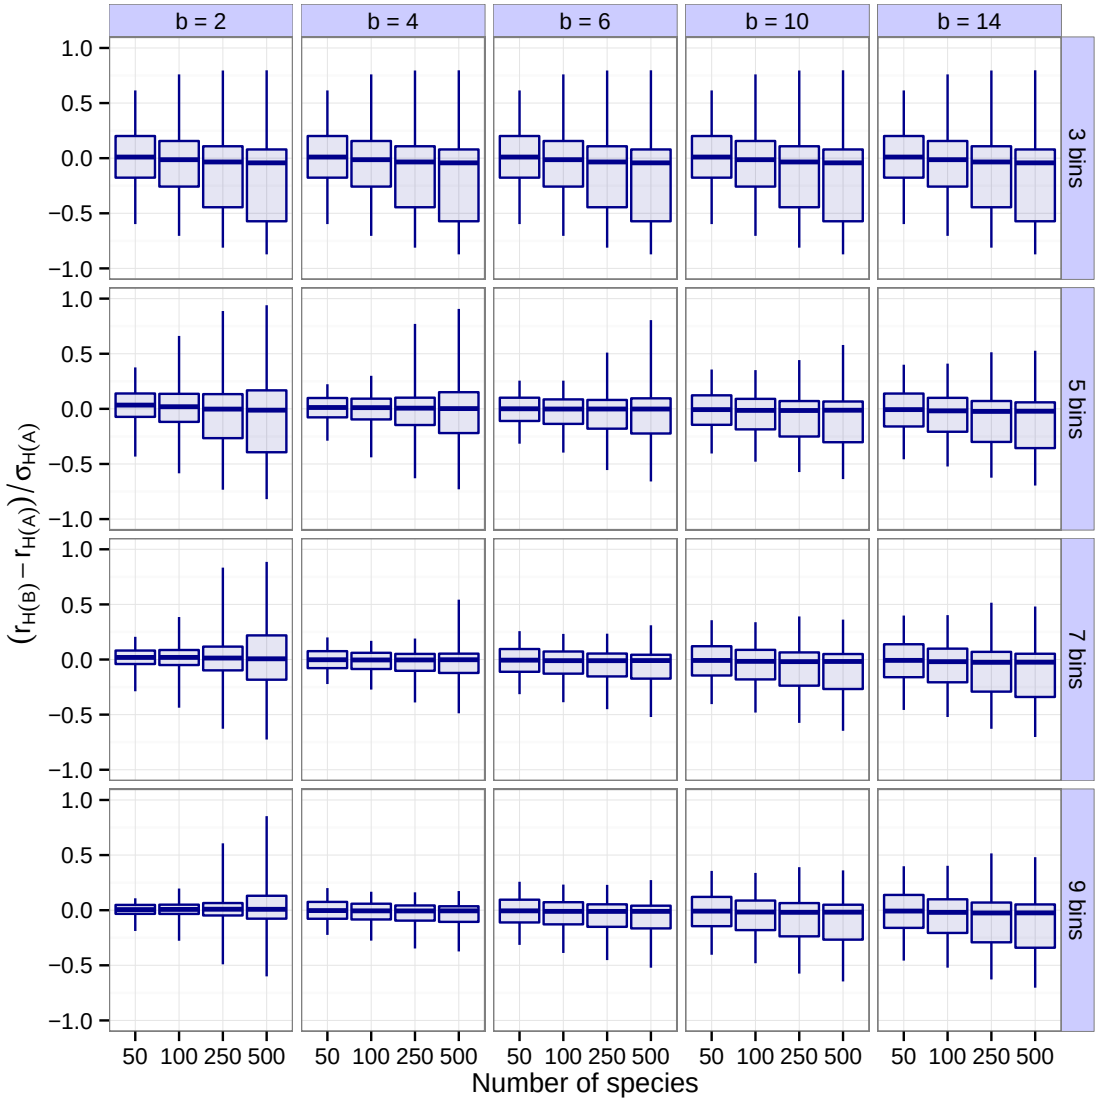

Figure S3: Box plots of how well the leading eigenvalue of A's Hermitian part $H(A)$ is approximated by the leading eigenvalue of B's Hermitian part $H(B)$, where B is the binned counterpart of A. The ordinate of each panel shows the difference between the leading eigenvalues $r_{H(A)}$ of the Hermitian parts of the original and $r_{H(B)}$ of the Hermitian parts of the binned matrices relative to $\sigma_{H(A)}$, the total range of the eigenvalues of the Hermitian parts of the original matrices. The rate of misclassification is 0%. Rows, columns, and panels are organized just as in Fig. S1.

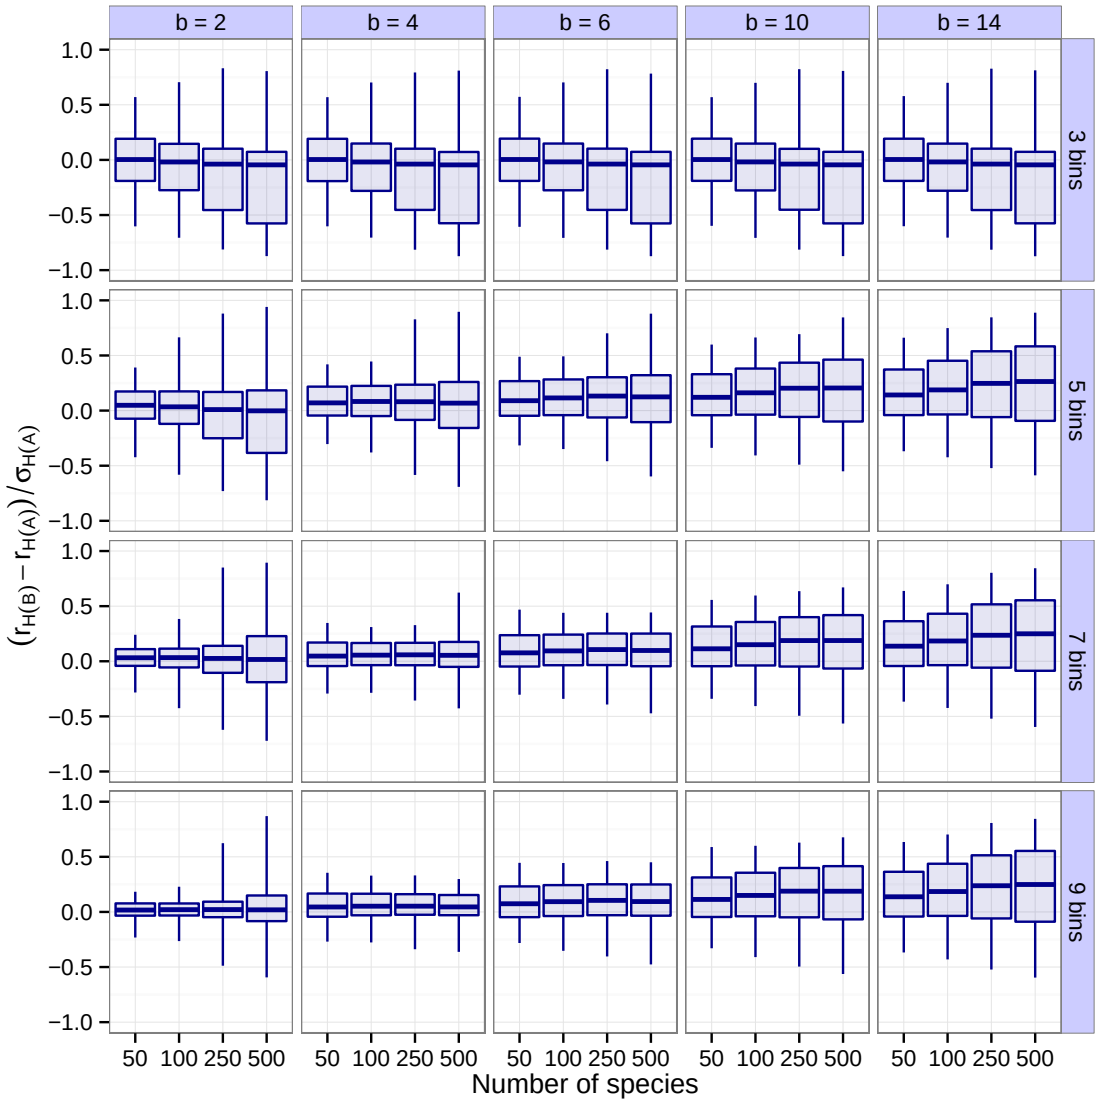

Figure S4: As Fig. S3, except each matrix is binned with a 10% misclassification rate.

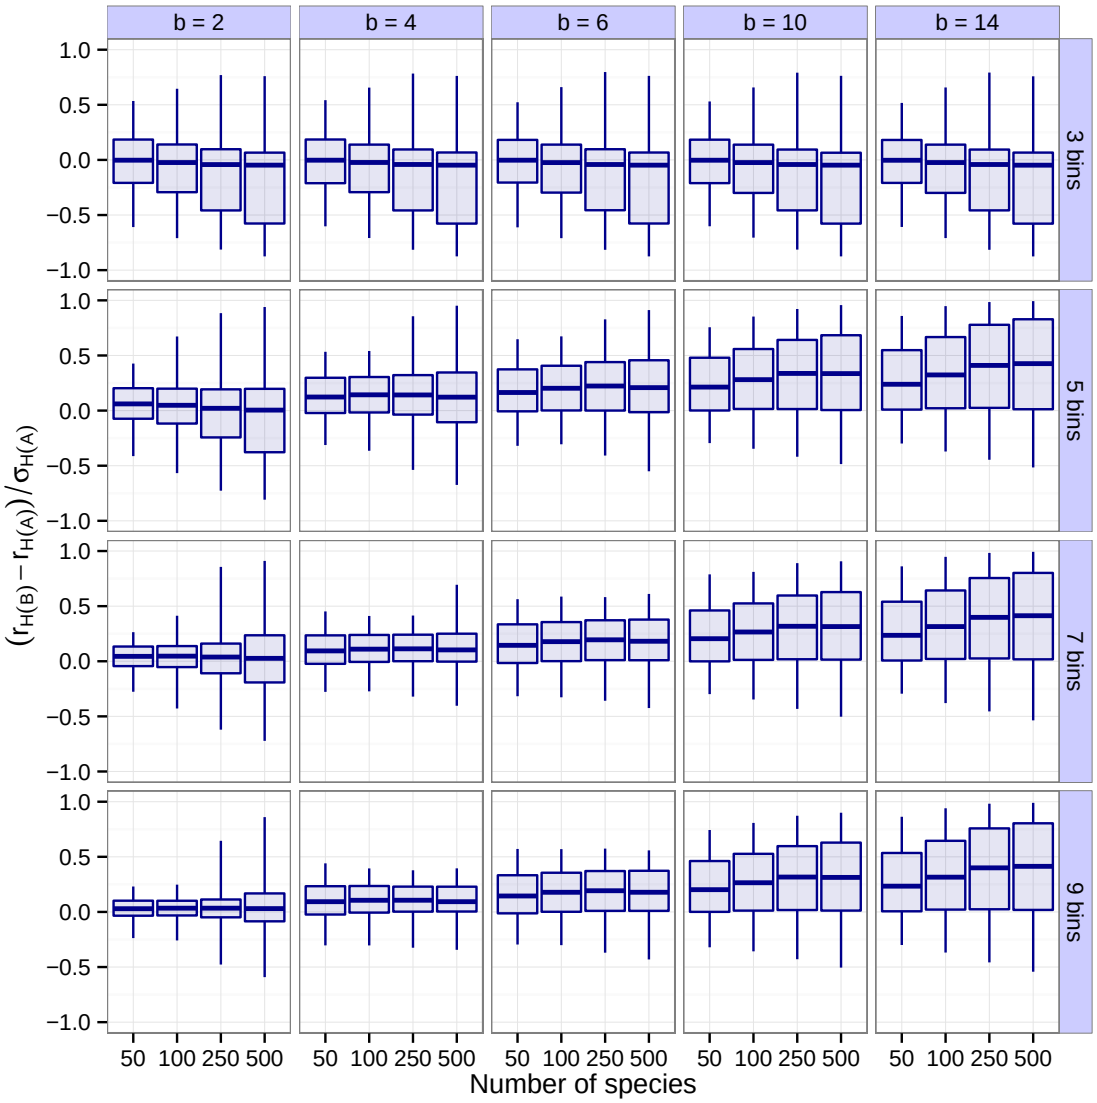

Figure S5: As Fig. S3, except each matrix is binned with a 20% misclassification rate.

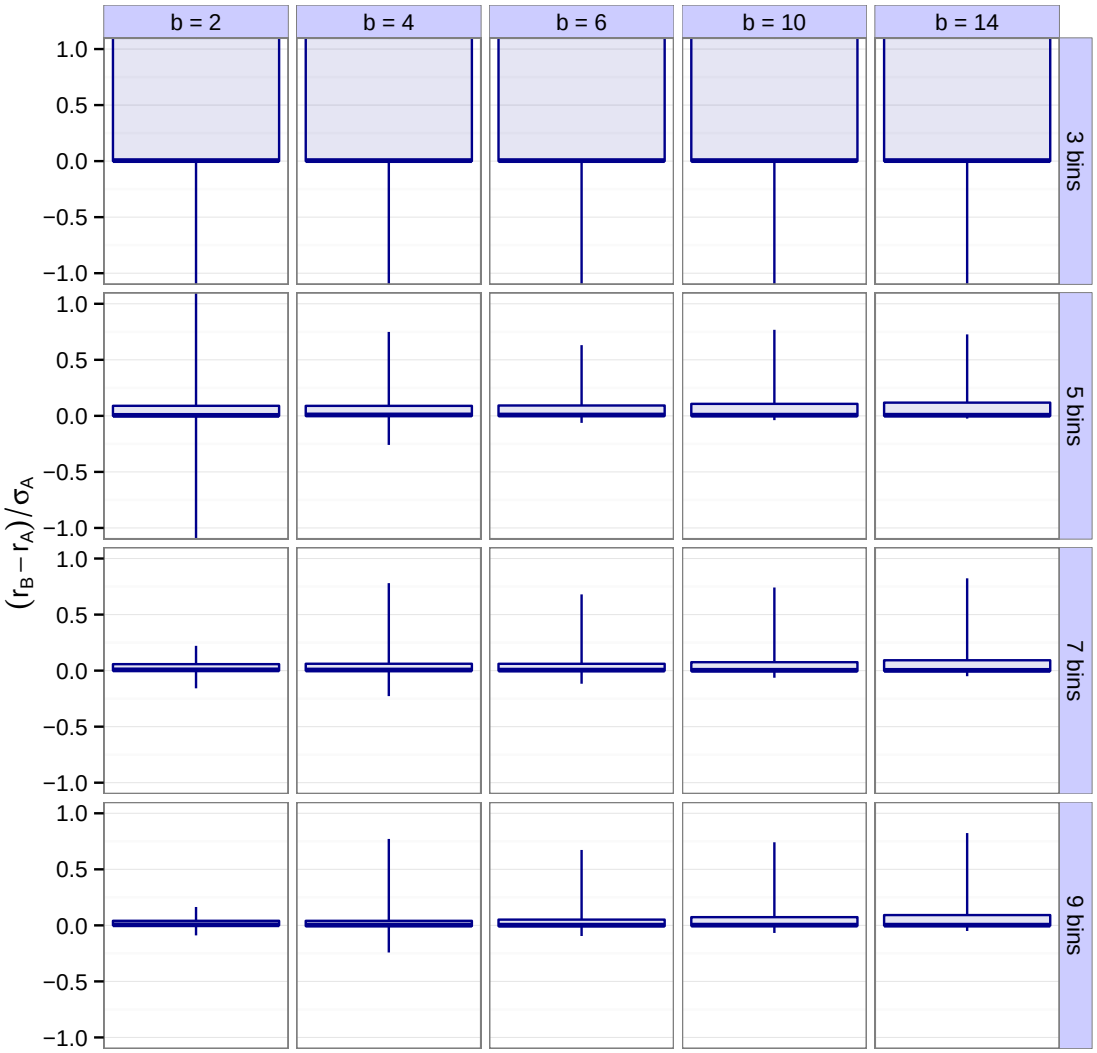

Figure S6: Box plots of how the leading eigenvalues of community matrices A are captured by those of their binned counterparts B, where the matrices A are generated by the Allometric Trophic Network model (Berlow et al. 2009). Each matrix is binned with no misclassification. The figure is organized just like Fig. S1, except the data in the panels are not separated based on the number of species.

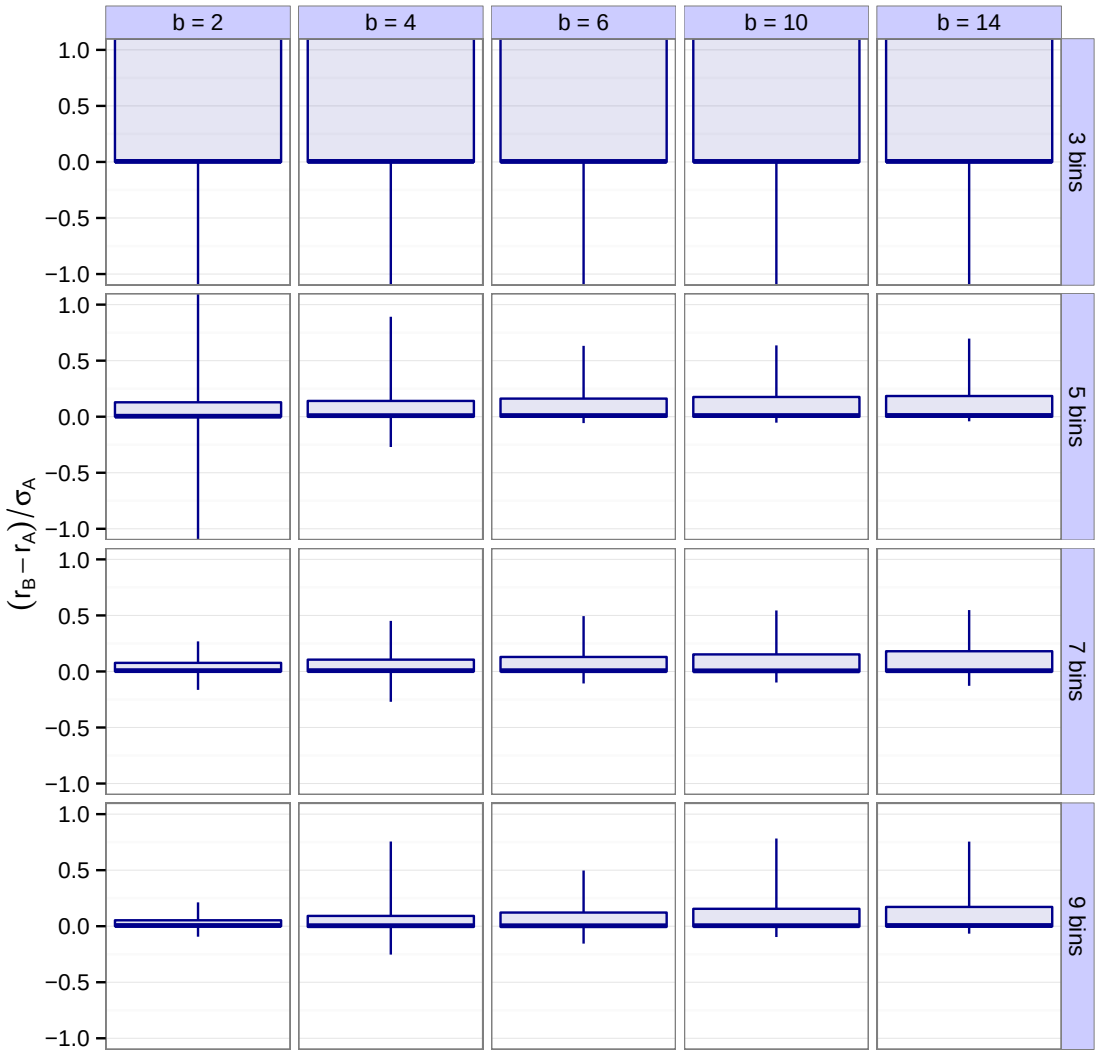

Figure S7: As Fig. S6, but with a 20% misclassification rate.

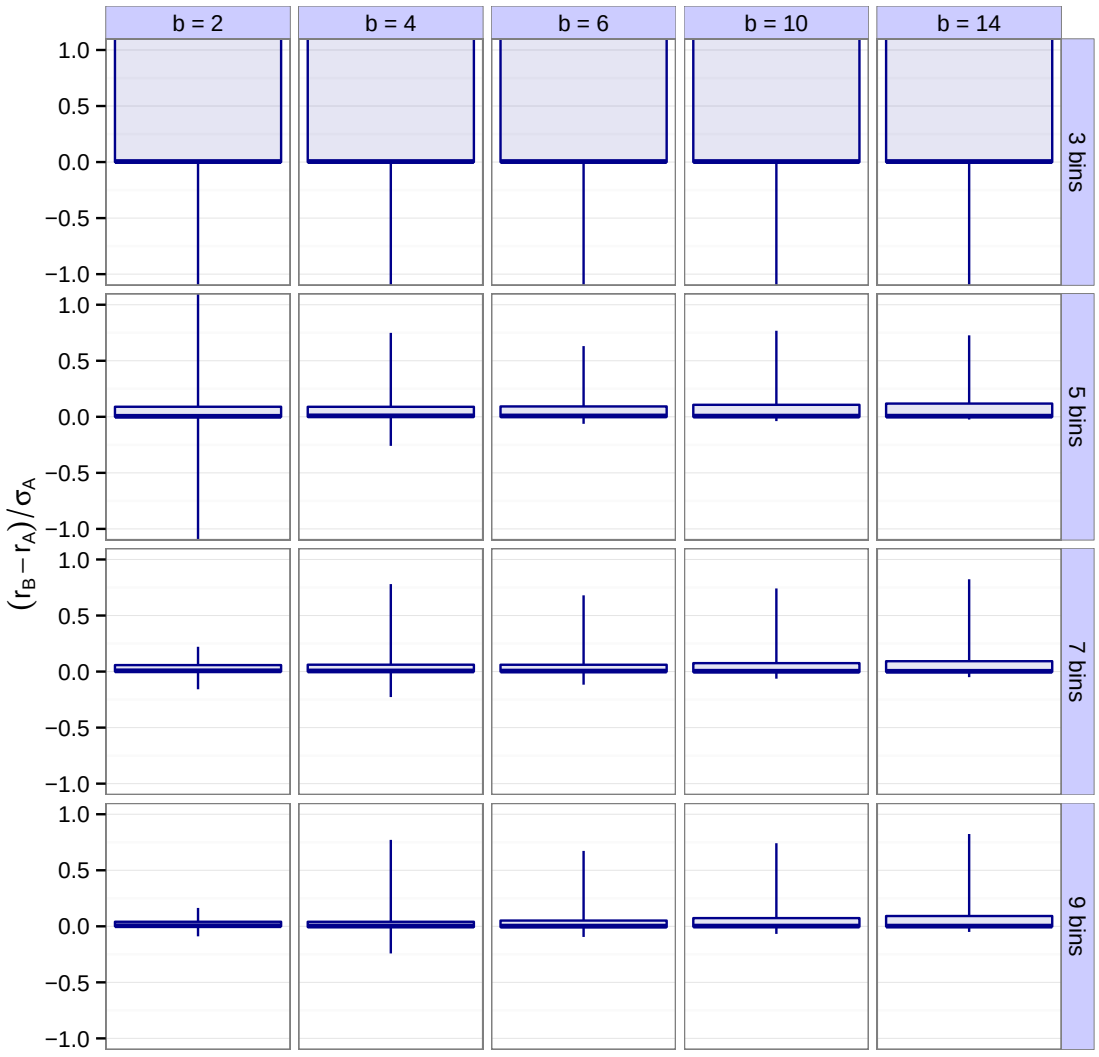

Figure S8: As Fig. S6, but for reactivity instead of stability; i.e., how well the leading eigenvalue of A's Hermitian part $H(A)$ is approximated by the leading eigenvalue of B's Hermitian part $H(B)$ in the Allometric Trophic Network model. Each matrix is binned with zero rate of misclassification.

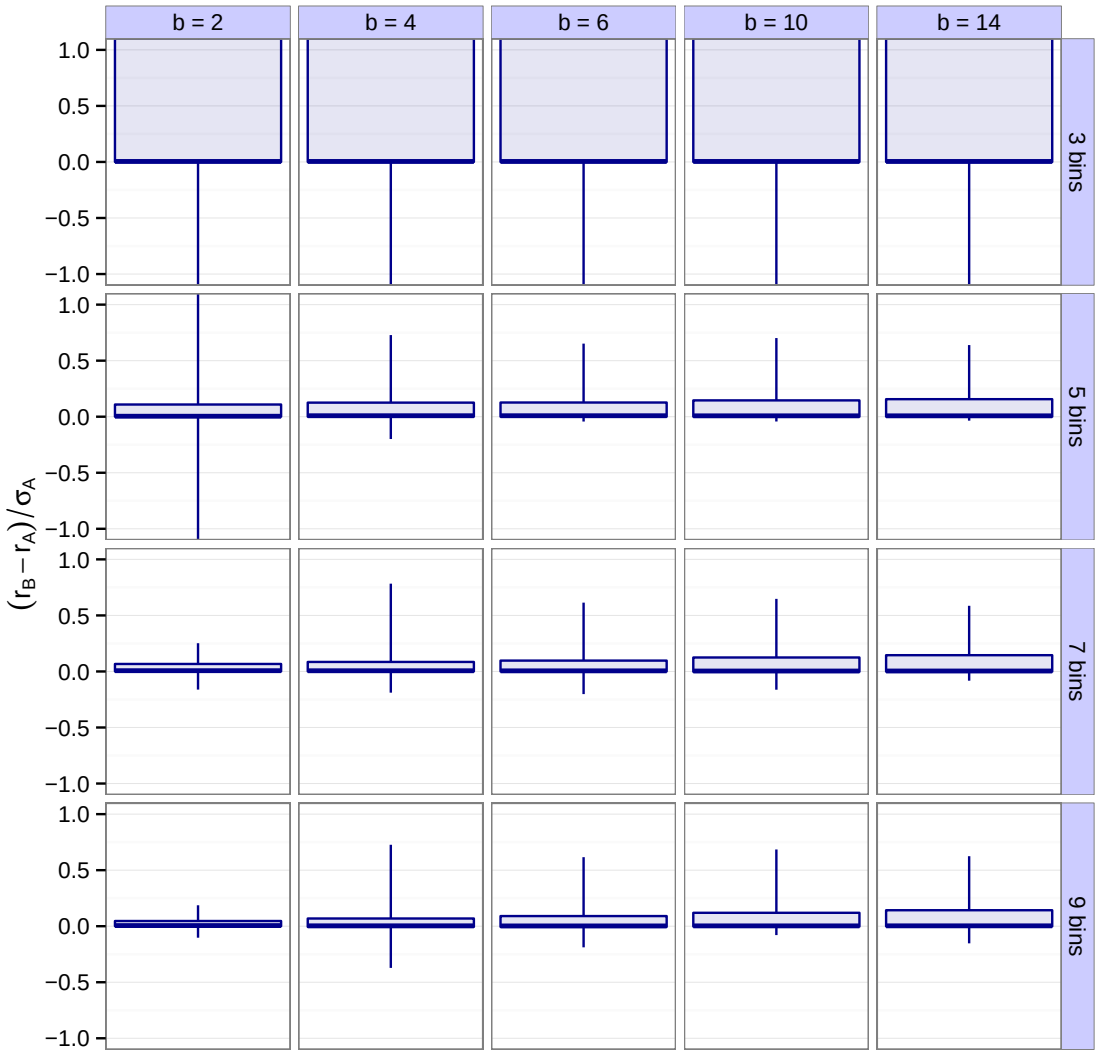

Figure S9: As Fig. S8, but with a 10% misclassification rate.

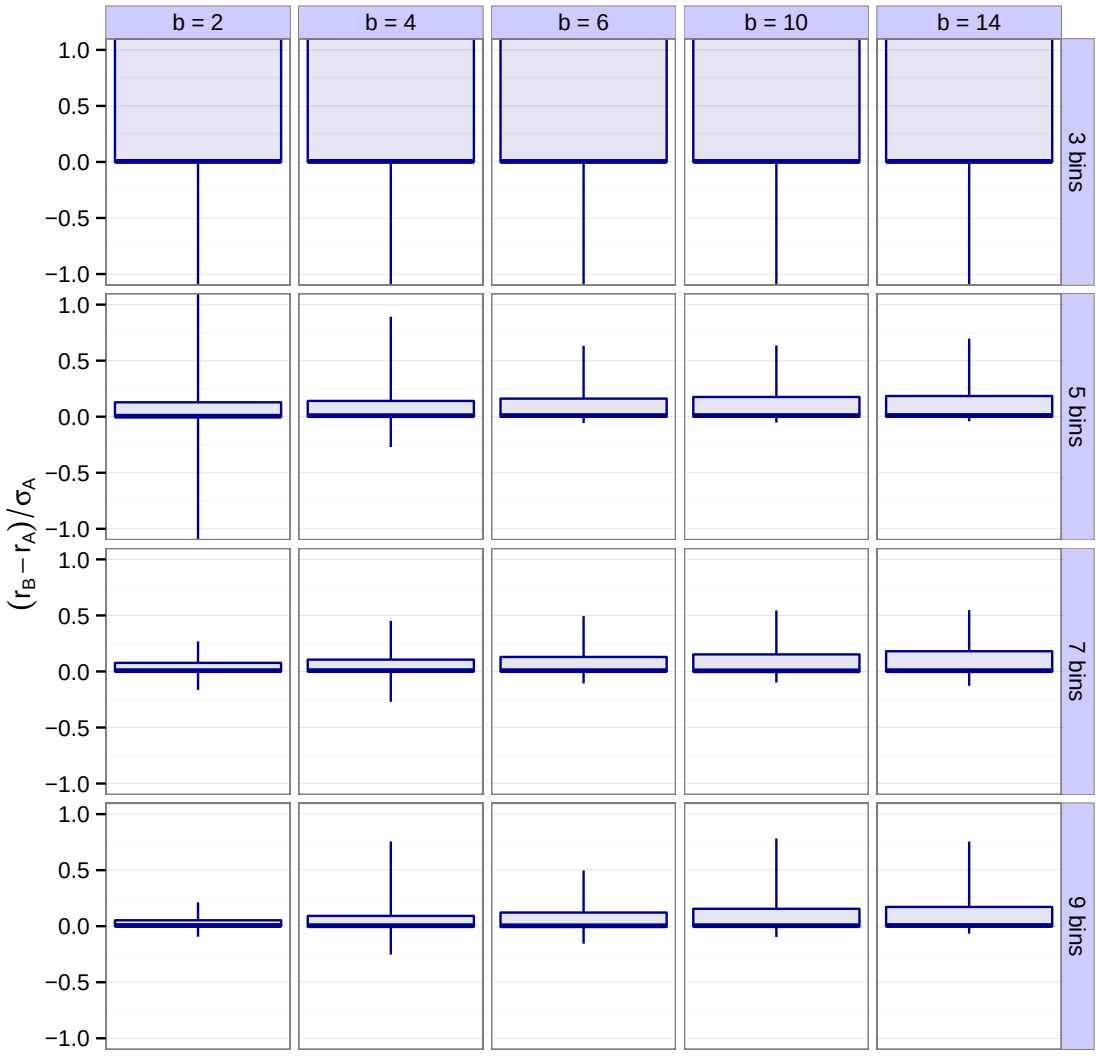

Figure S10: As Fig. S8, but with a 20% misclassification rate.

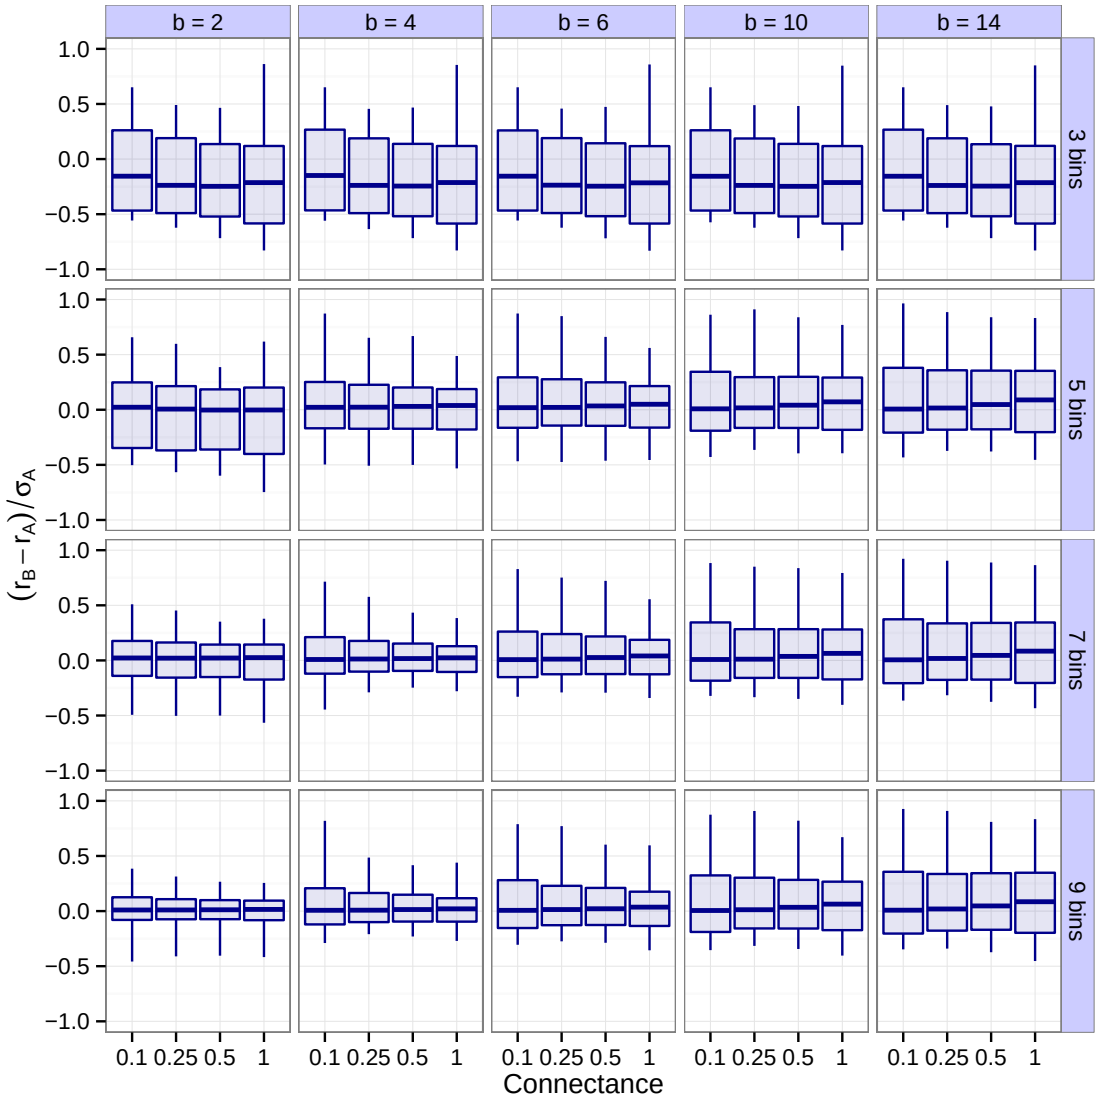

Figure S11: The effect of varying network connectance on the prediction accuracy of the leading eigenvalue. Organized as Fig. S1, except with 10% misclassification rate, the abscissa is organized by different values of the connectance instead of different values of the number of species, and the number of species is fixed at $S = 50$.

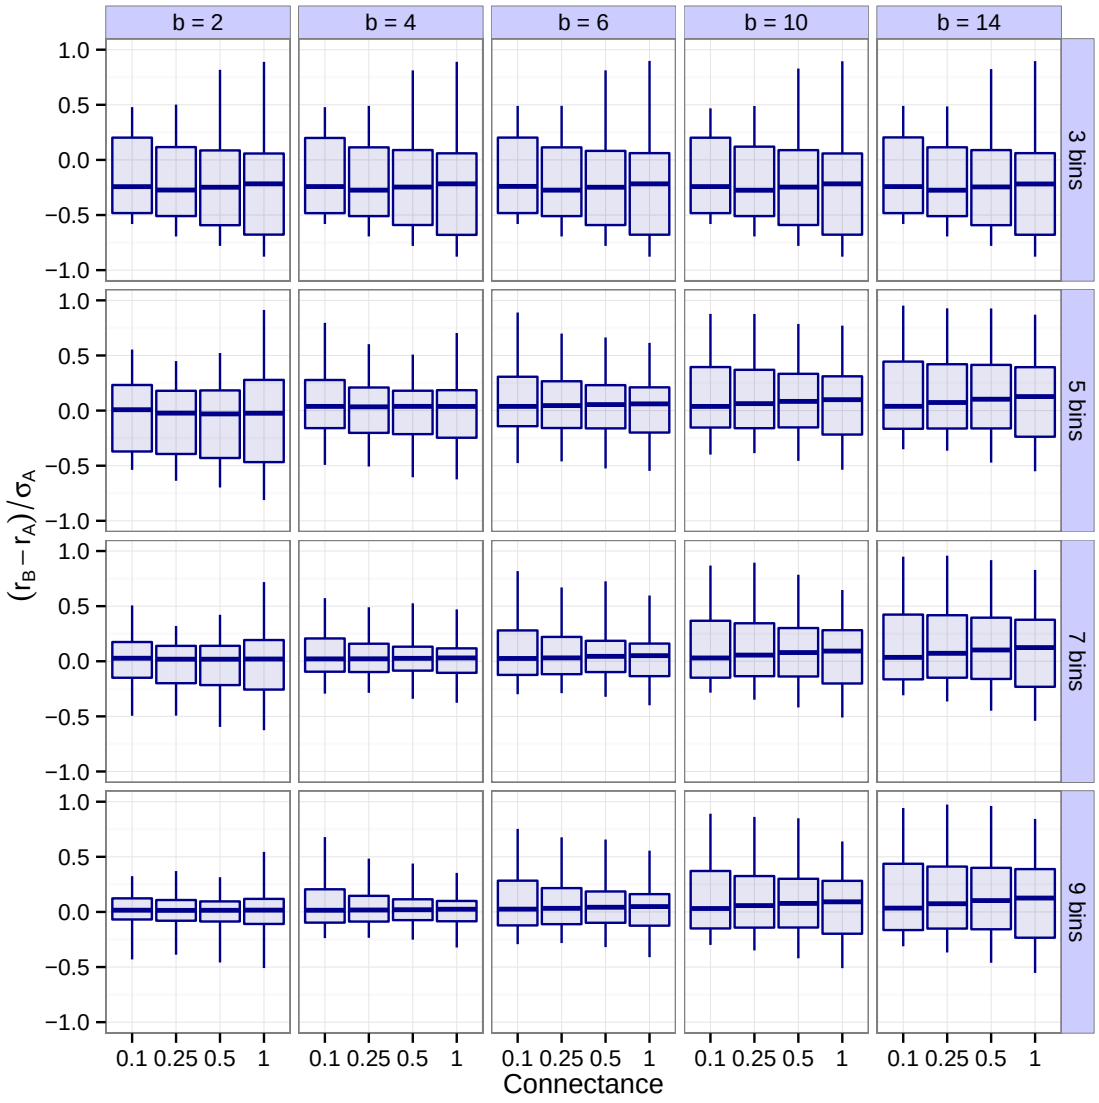

Figure S12: As Fig. S11, but with the number of species fixed at $S = 100$.

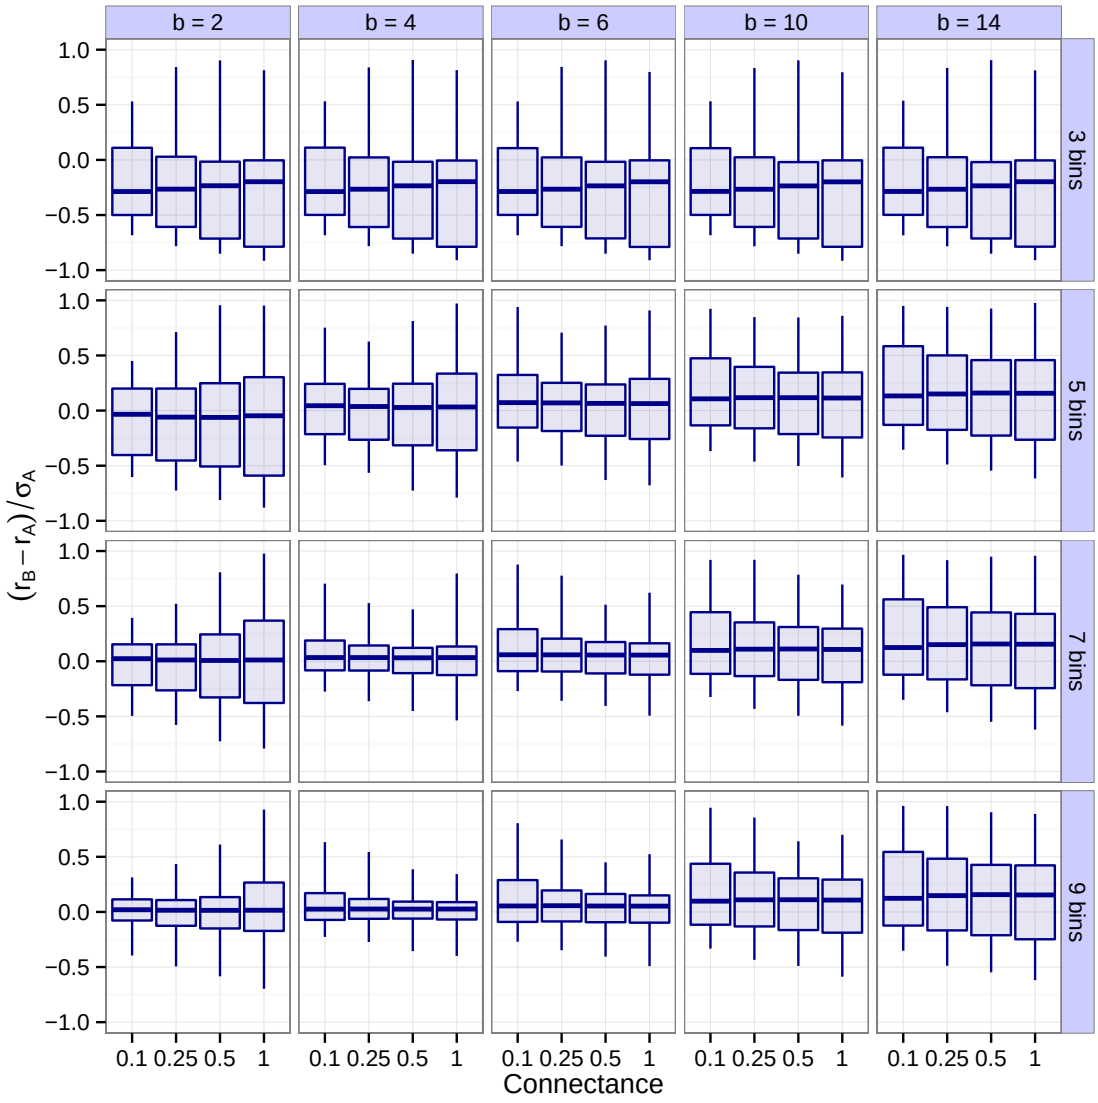

Figure S13: As Fig. S11, but with the number of species fixed at $S = 250$.

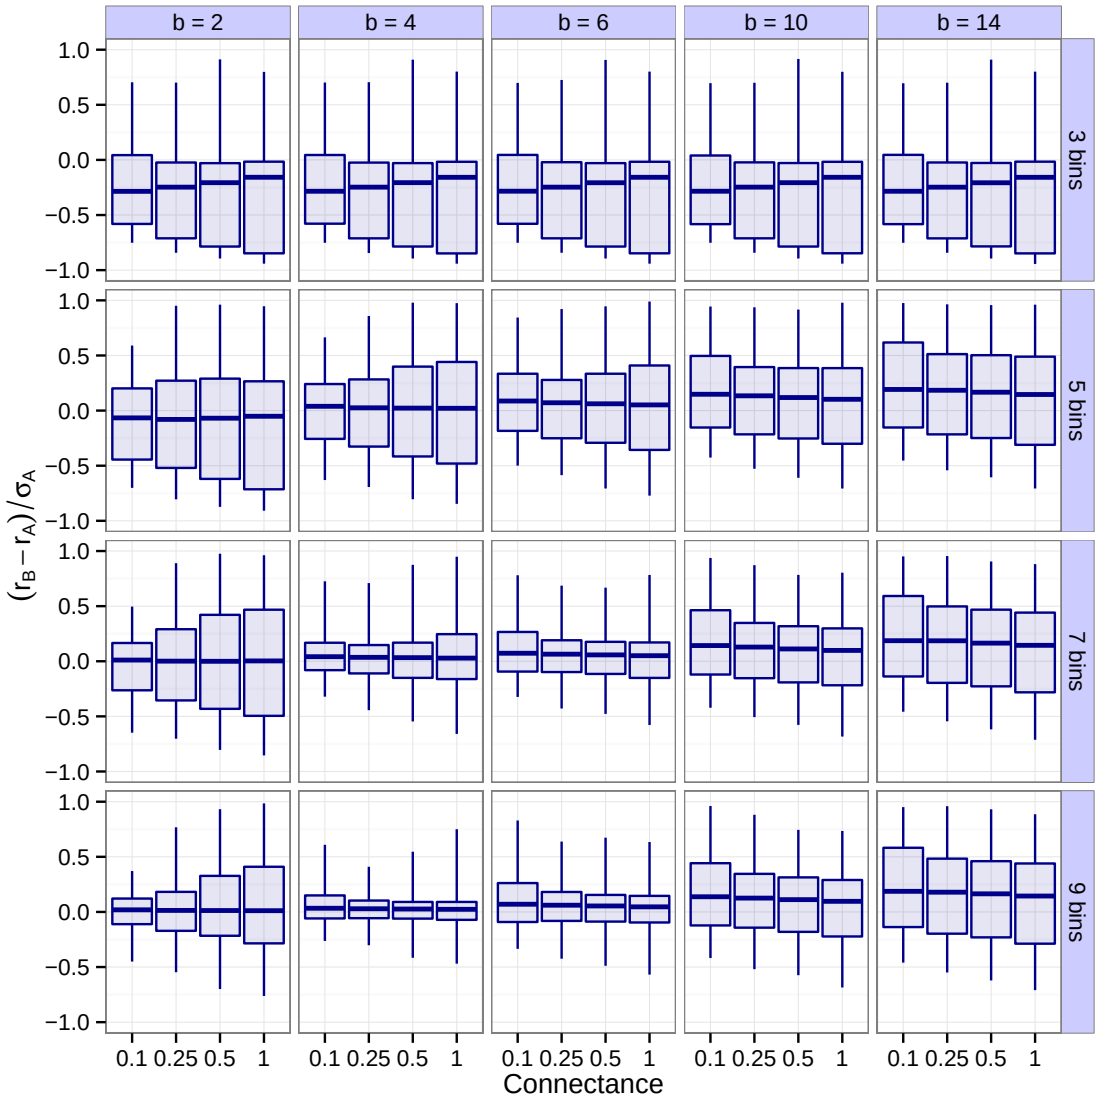

Figure S14: As Fig. S11, but with the number of species fixed at $S = 500$.

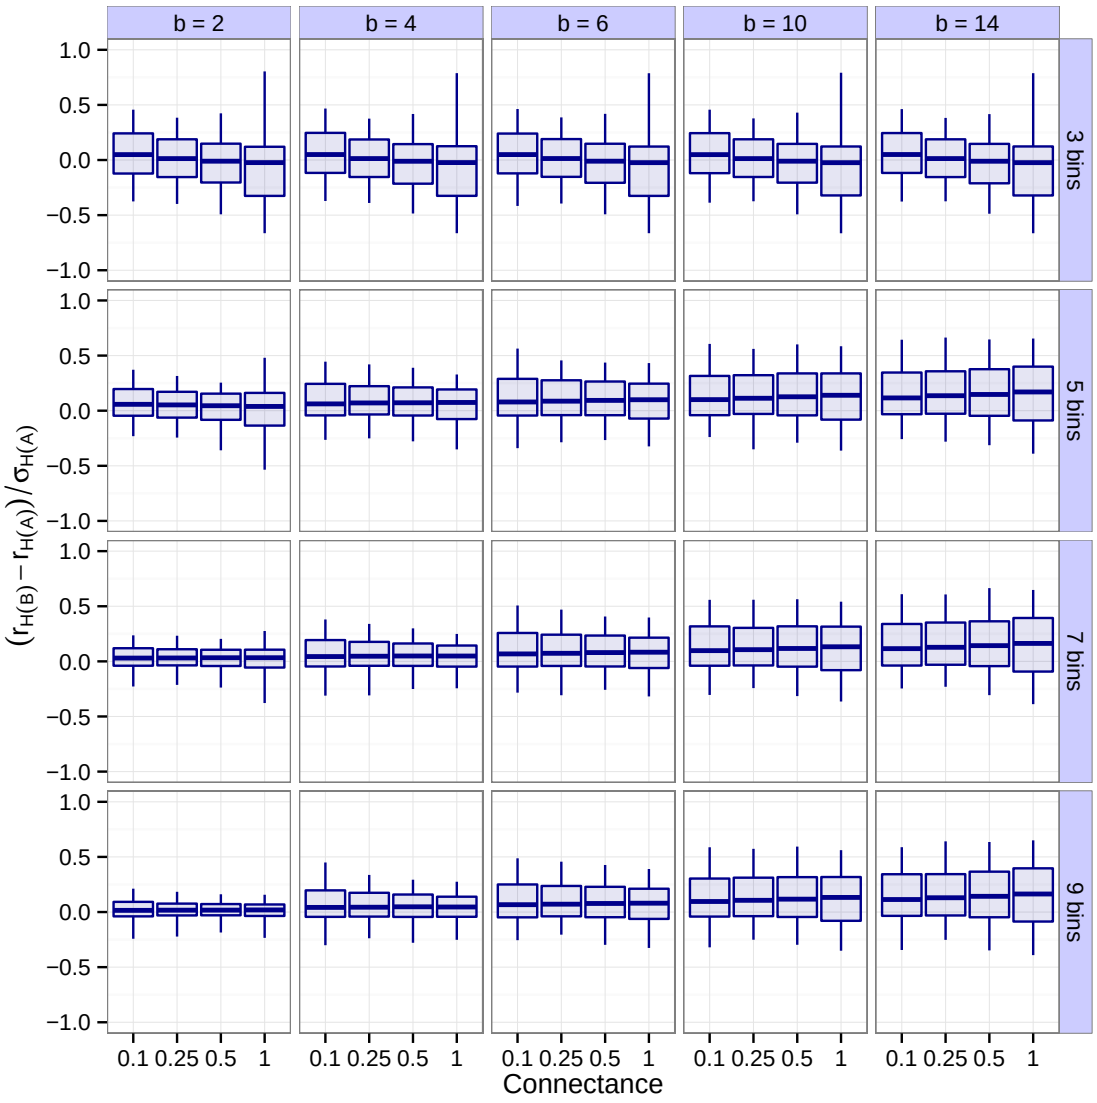

Figure S15: The effect of varying network connectance on the prediction accuracy of the leading eigenvalue of the Hermitian part (related to reactivity). As Fig. S11, but the ordinate represents predictions of reactivity instead of stability.

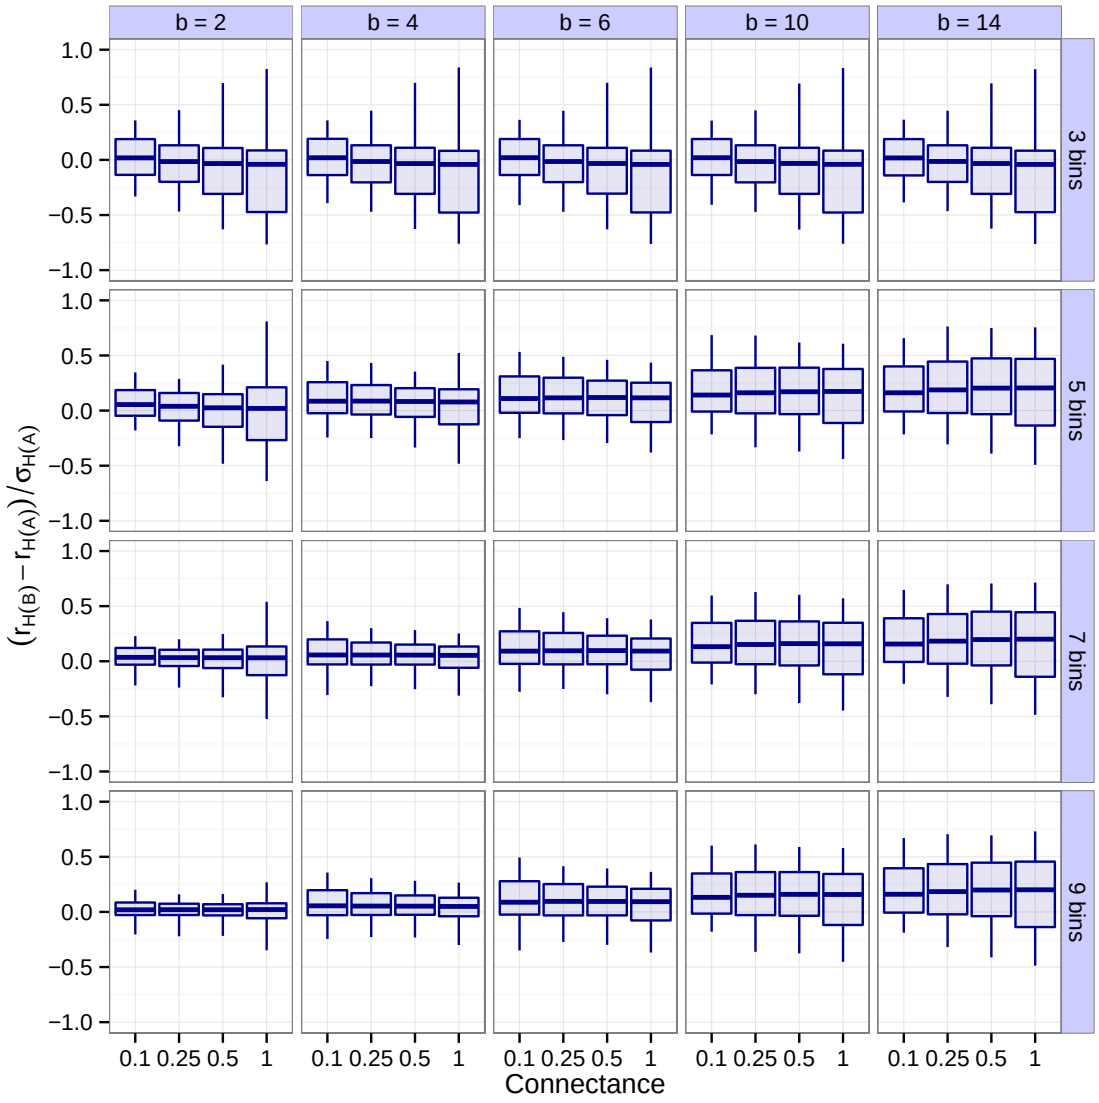

Figure S16: As Fig. S11, but with the number of species fixed at $S = 100$.

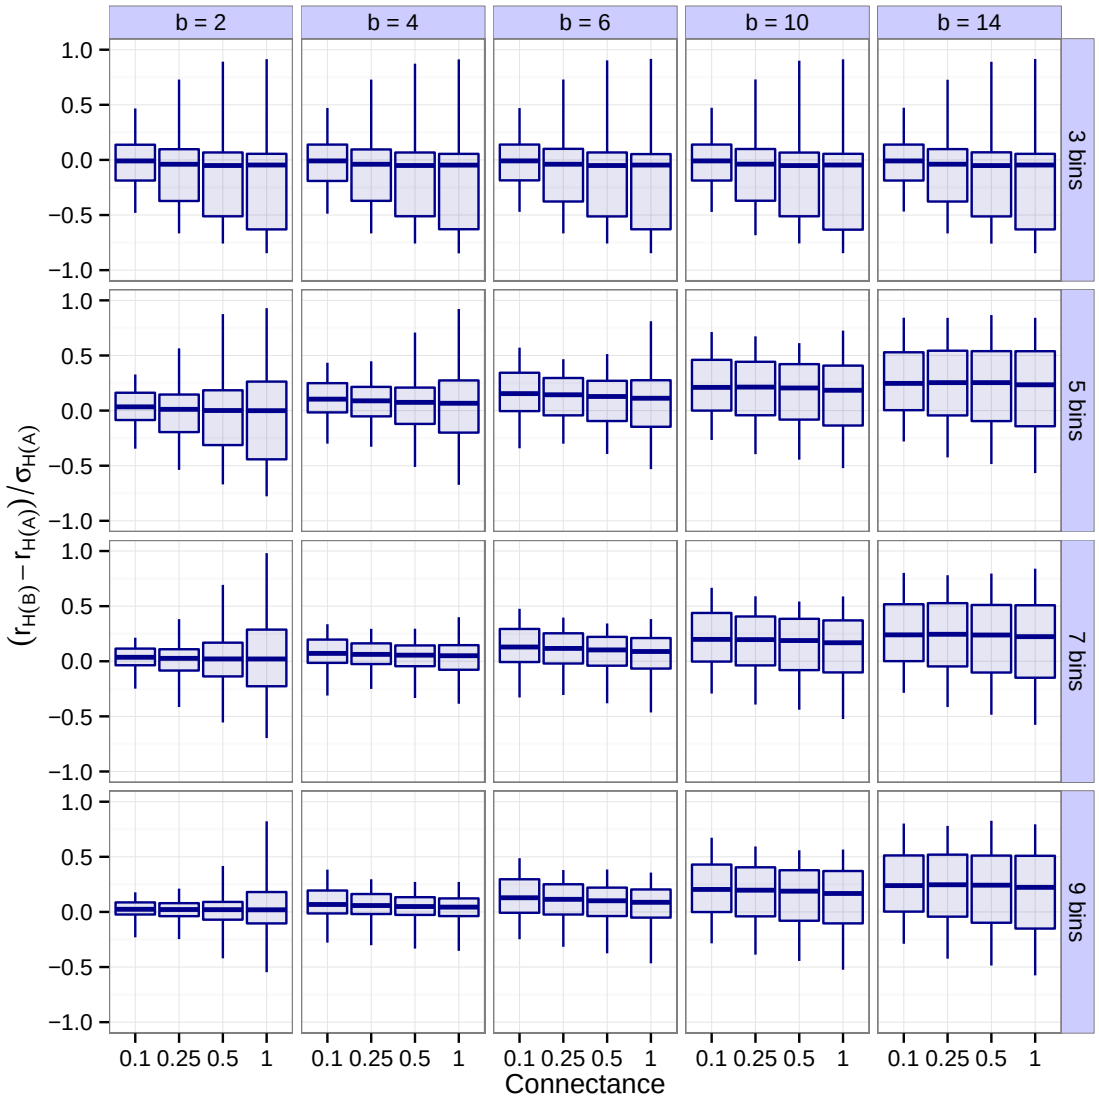

Figure S17: As Fig. S11, but with the number of species fixed at $S = 250$.

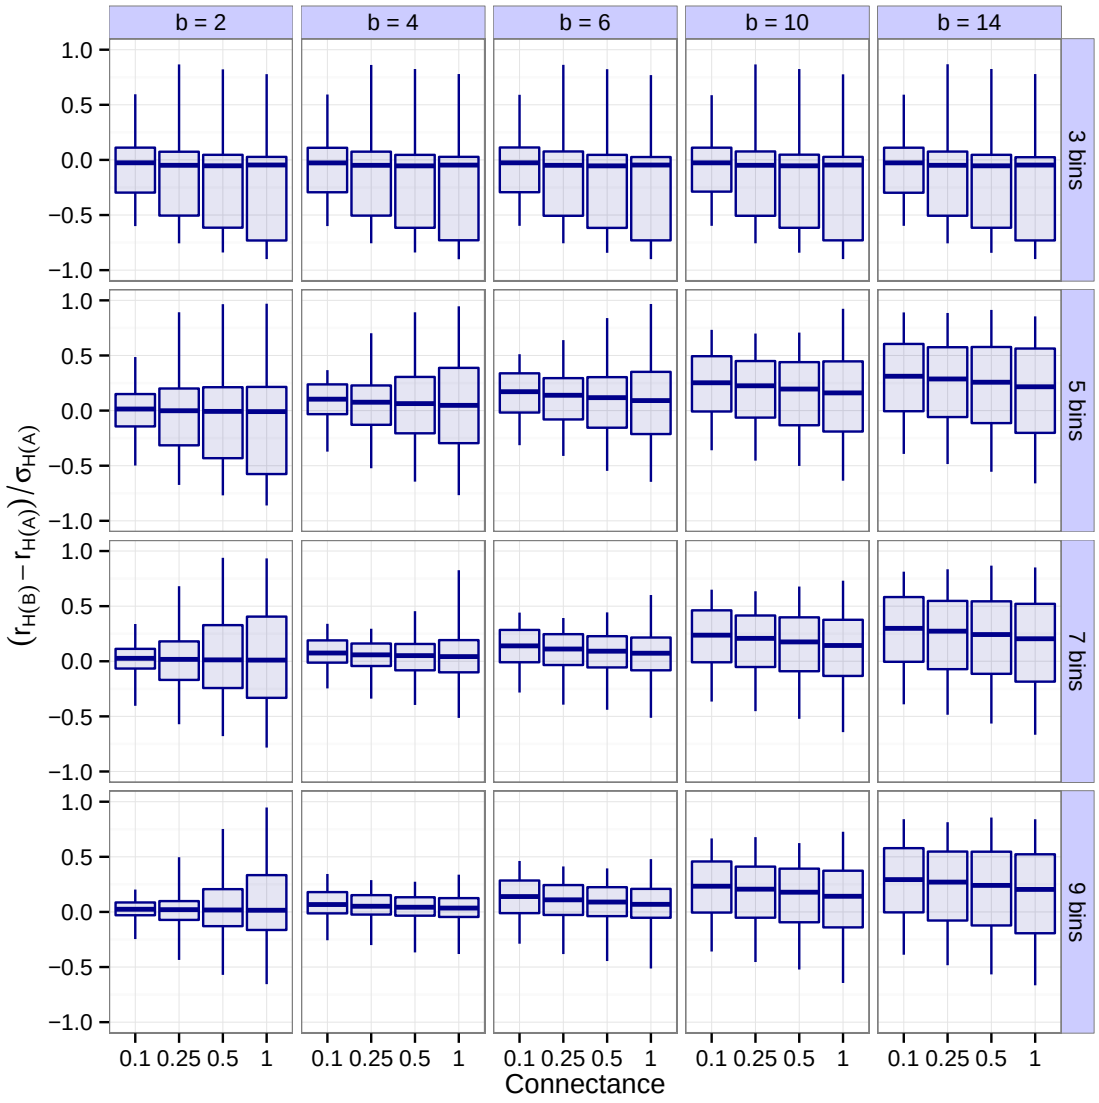

Figure S18: As Fig. S11, but with the number of species fixed at $S = 500$.

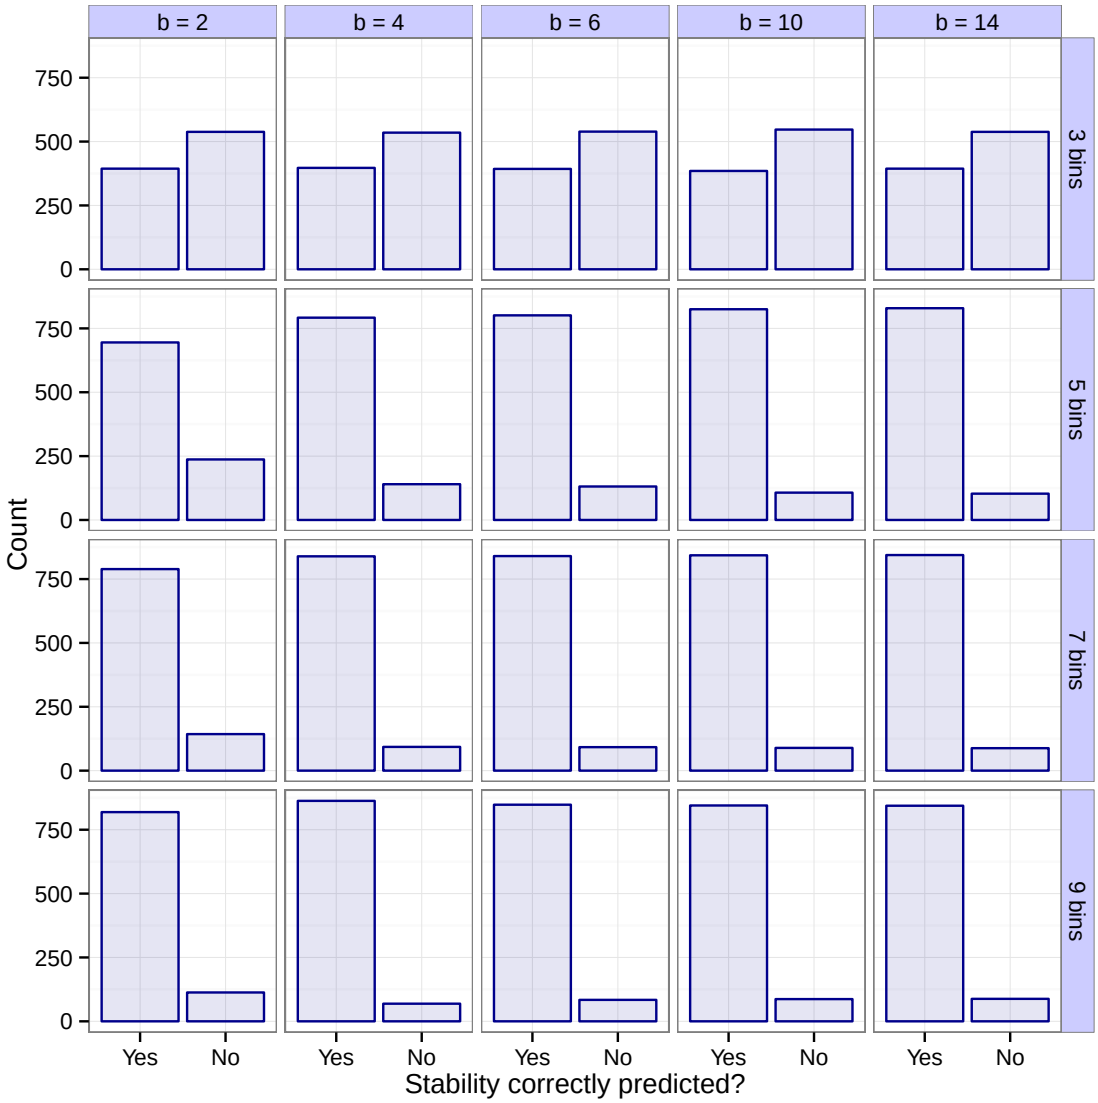

Figure S19: The number of cases in which stability is correctly predicted, out of those matrices whose leading eigenvalues have a magnitude less than one twentieth of the total spread of the real parts of all eigenvalues: $|r_A|/\sigma_A < 0.05$. Rows correspond to different values of the binning resolution b; columns to different numbers of bins k. The rate of misclassification is 10%.

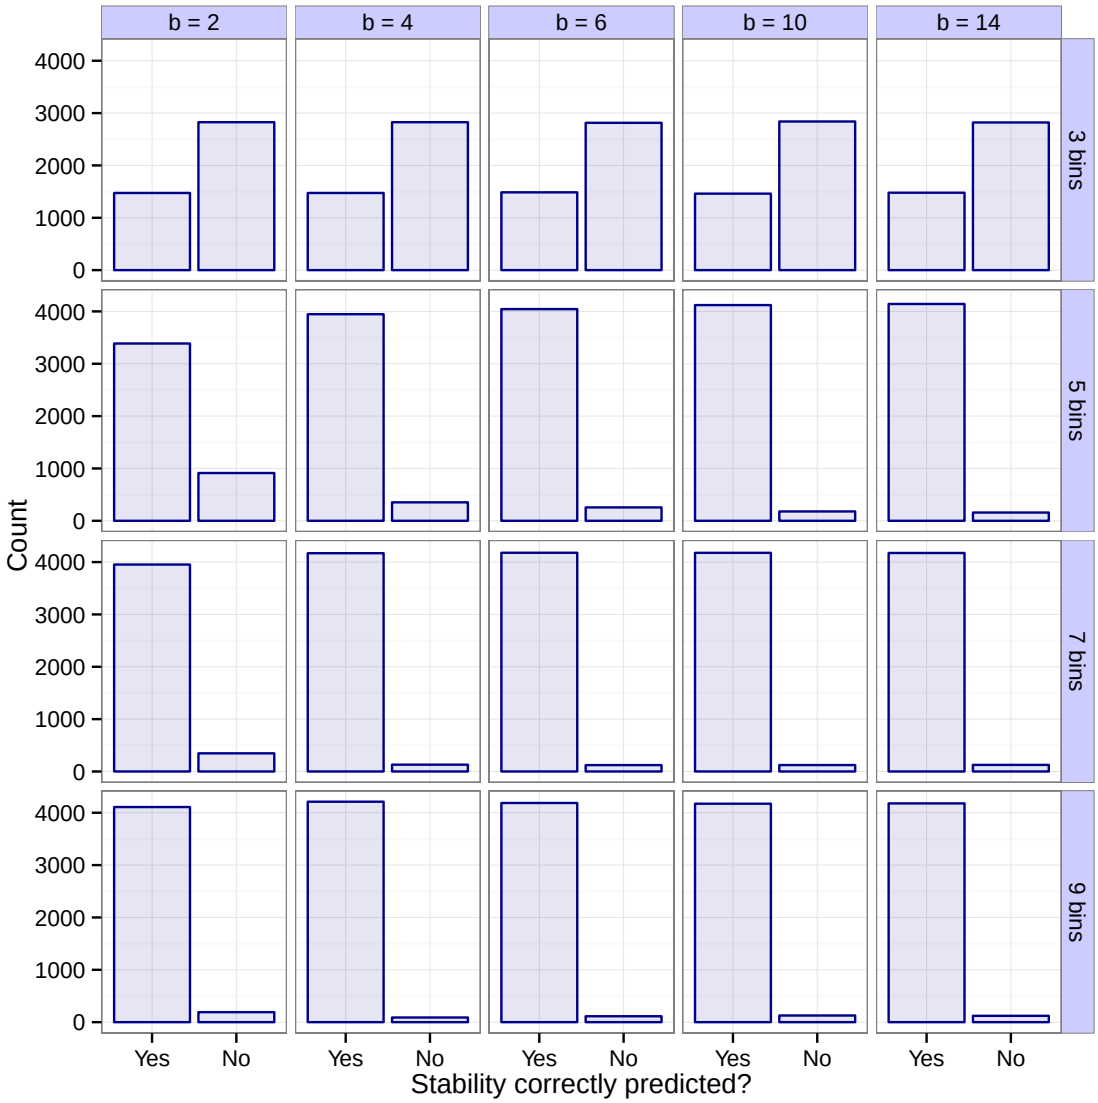

Figure S20: As Fig. S19, but for all those matrices for which $|r_A|/\sigma_A < 0.1$.

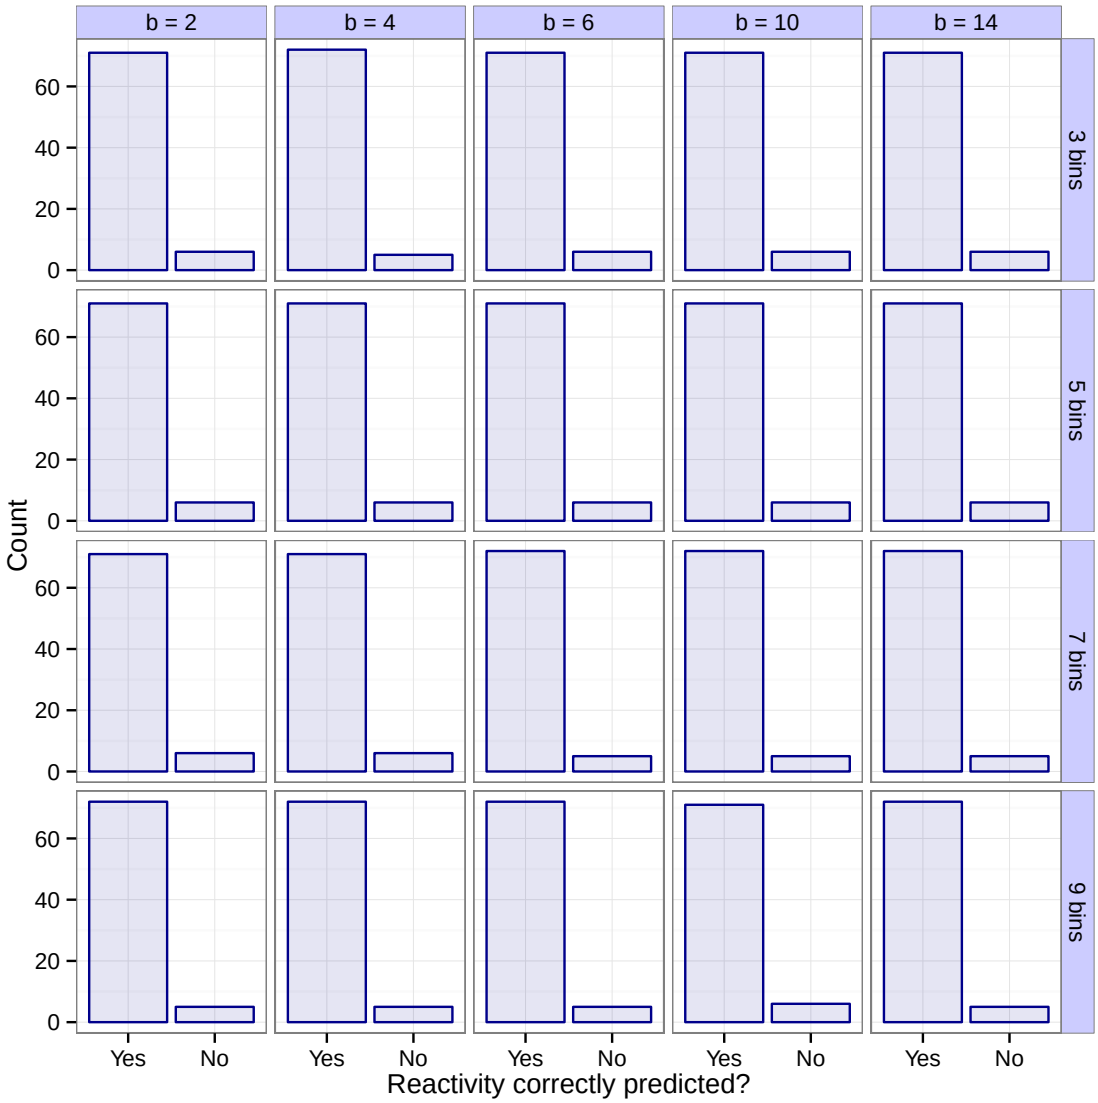

Figure S21: As Fig. S19, but for reactivity instead of stability.

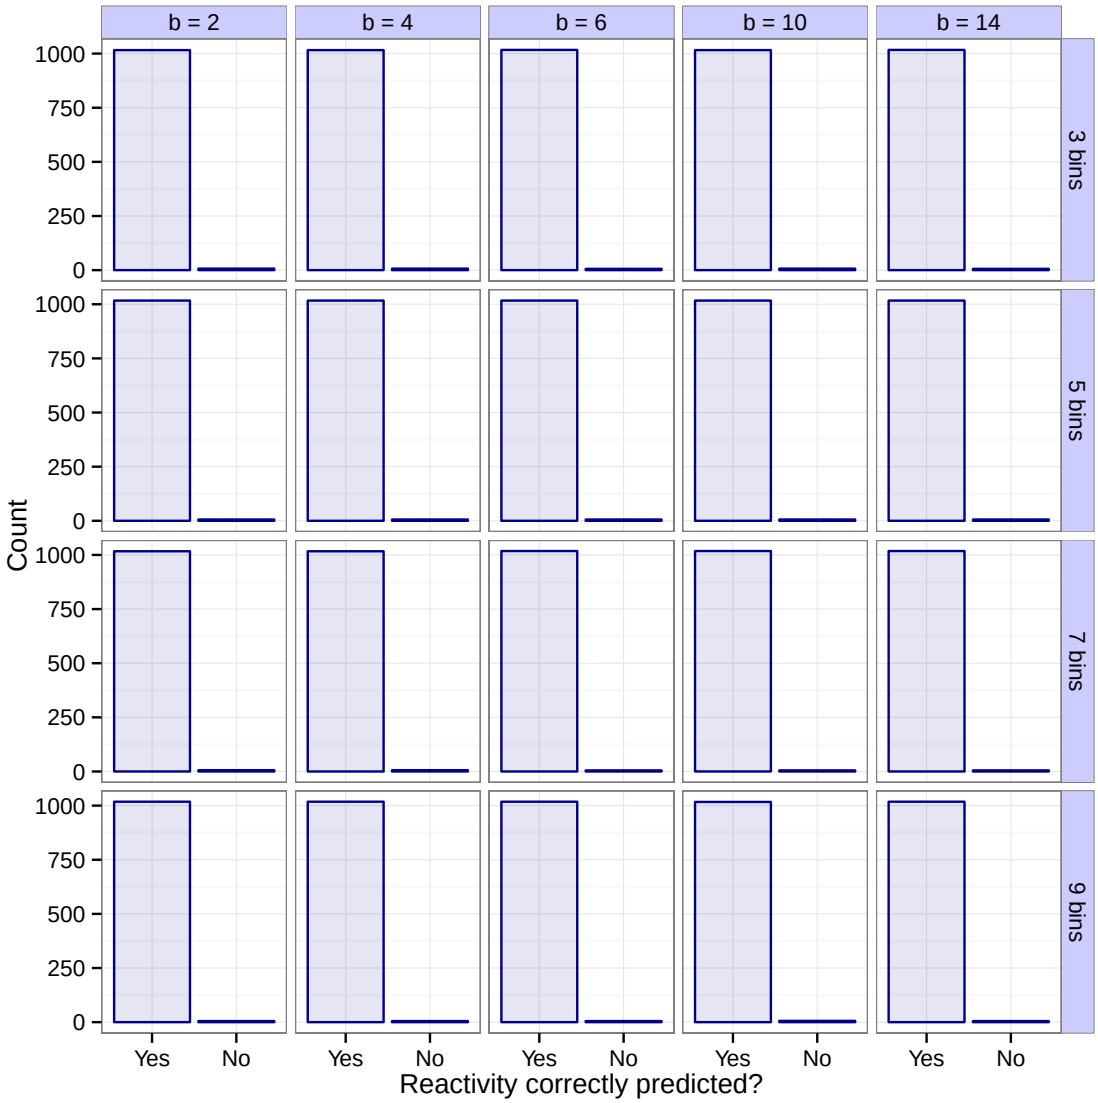

Figure S22: As Fig. S20, but for reactivity instead of stability.

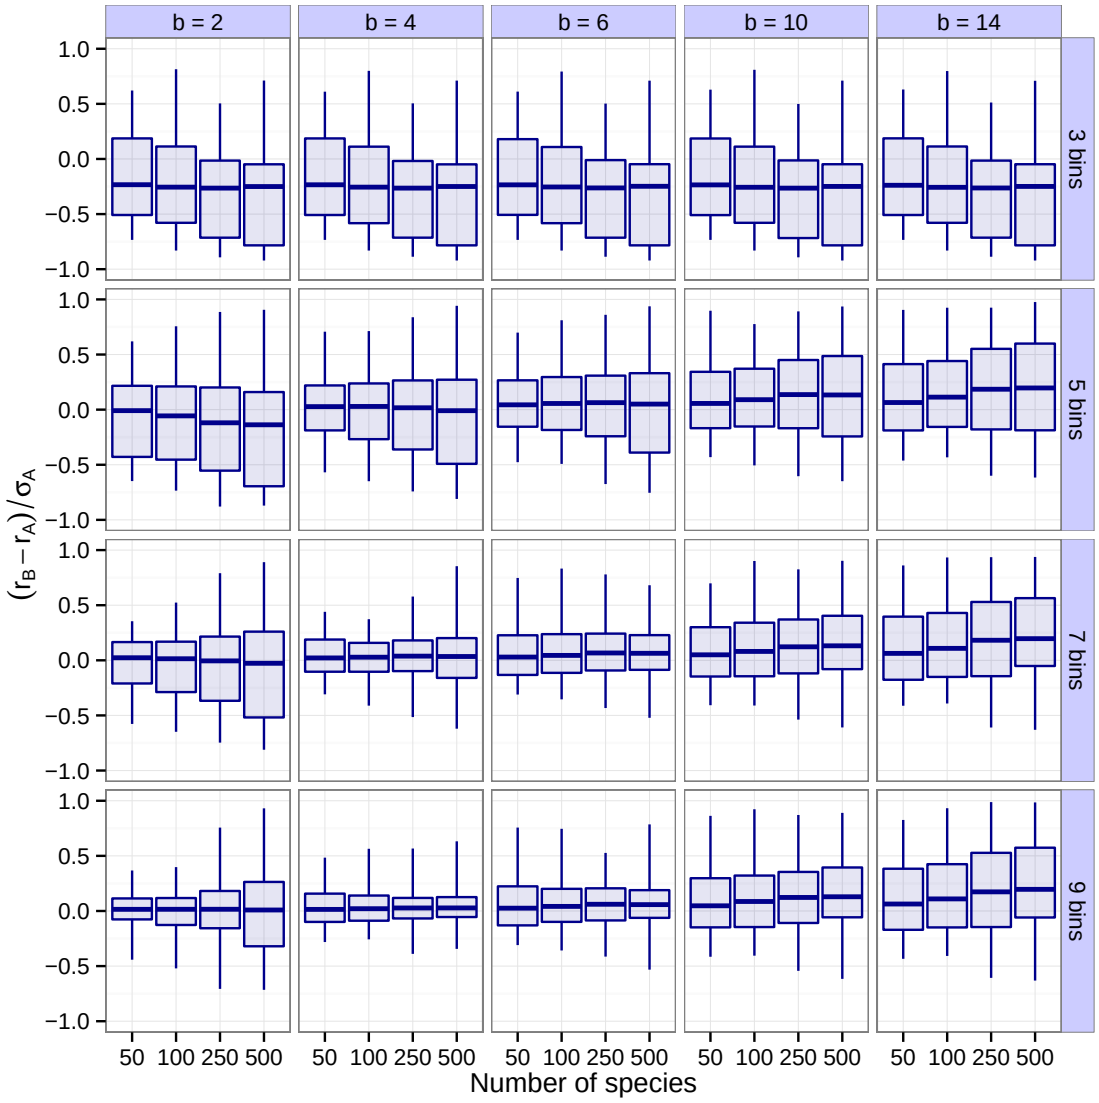

Figure S23: As Fig. S1, but with a misclassification rate of 10%, and all interaction strengths drawn from lognormal distributions of varying means and variances (see main text).

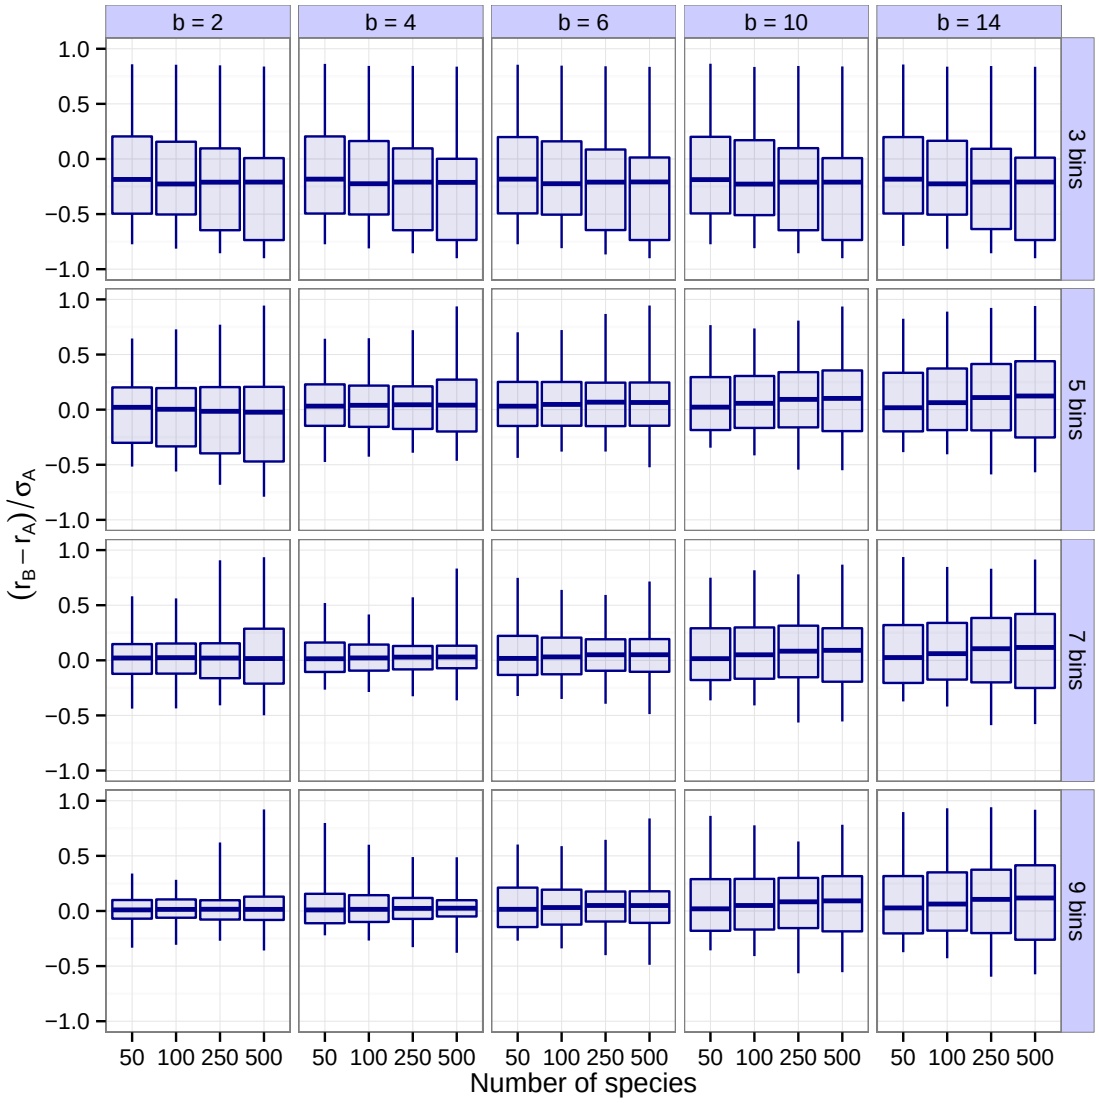

Figure S24: As Fig. S23, but with all interaction strengths drawn from variously parameterized Gamma distributions.

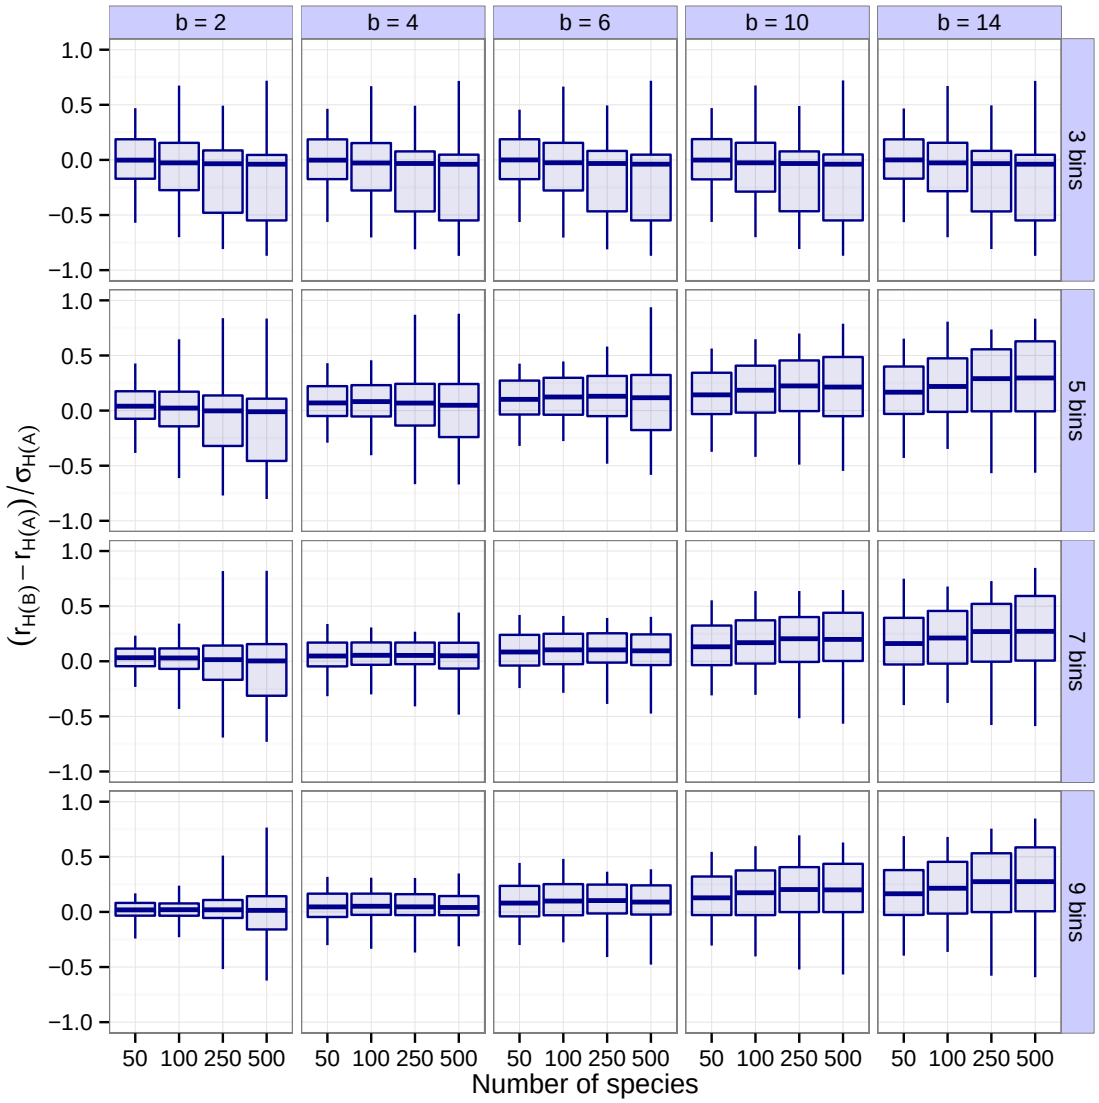

Figure S25: As Fig. S23, but for reactivity instead of stability.

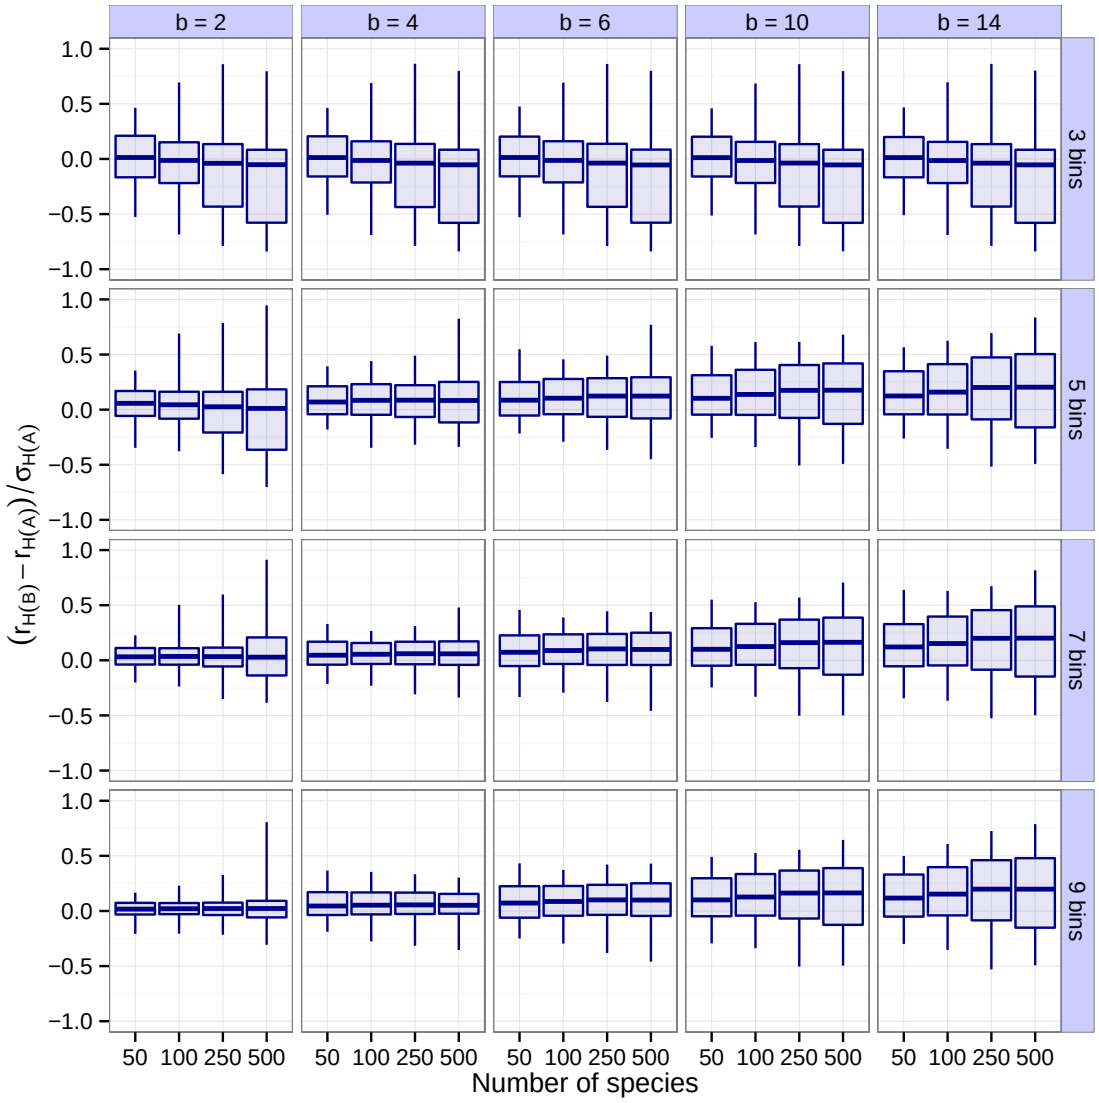

Figure S26: As Fig. S25, but with all interaction strengths drawn from variously parameterized Gamma distributions.

References

Allesina, S., Tang, S., 2012. Stability criteria for complex ecosystems. Nature 483, 205–208.

Berlow, E. L., Dunne, J. A., Martinez, N. D., Stark, P. B., Williams, R. J., Brose, U., 2009. Simple prediction of interaction strengths in complex food webs. Proceedings of the National Academy of Sciences USA 106, 187–191.

Edelman, A., 1988. Eigenvalues and condition numbers of random matrices. SIAM Journal on Matrix Analysis and Applications 9, 543–560.

Ginibre, J., 1965. Statistical ensembles of complex, quaternion, and real matrices. Journal of Mathematical Physics 6, 440–449.

Girko, V. L., 1984. The circle law. Theory of Probability and its Applications 29, 694–706.

Henrici, P., 1962. Bounds for iterates, inverses, spectral variation and fields of values of non-normal matrices. Numerische Mathematik 4, 24–40.

Lee, S. L., 1996. Best available bounds for departure from normality. SIAM Journal on Matrix Analysis and Applications 17, 984–991.

Marčenko, V. A., Pastur, L. A., 1967. Distribution of eigenvalues for some sets of random matrices. Math USSR Sbornik 1, 457–483.

R Development Core Team, 2008. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.

Sommers, H. J., Crisanti, A., Sompolinsky, H., Stein, Y., 1998. Spectrum of large random asymmetric matrices. Physical Review Letters 60, 1895–1898.

Tang, S., Pawar, S., Allesina, S., 2014. Correlation between interaction strengths drives stability in large ecological networks. Ecology Letters 17, 1094–1100.

Tao, T., Vu, V., Krishnapur, M., 2010. Random matrices: Universality of esds and the circular law. Annals of Probability 38 (5), 2023–2065.

Trefethen, L. N., Embree, M., 2005. Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press, New Jersey, USA.

Williams, R. J., Martinez, N. D., 2000. Simple rules yield complex food webs. Nature 404, 180–183.

Wolfram Research Inc., 2014. Mathematica, version 10.0. Wolfram, Champaign.