Predicting Collapse of Complex Ecological Systems: Quantifying the Stability-Complexity Continuum

Susanne Pettersson[1]*, Van M. Savage[2], Martin Nilsson-Jacobi[1]

Collapse, stability, and dynamical shifts between these states are hallmarks of ecological systems. A major goal in ecosystem research is to identify how limits of these states change with diversity, complexity, interaction-topology, and hierarchies. The primary focus has been on identifying conditions for a system to shift from strict stability to complete collapse. While this boundary is indeed of central importance, it is possible that real ecosystems can show a larger variety of responses to environmental changes.

Here, rather than focusing solely on limits of stability or collapse, we quantify and map the full phase space and the boundaries between regions with different response characteristics. We explore this phase space as biodiversity and complexity are varied for interaction webs in which consumer-resource interactions are chosen randomly and driven by Generalized-Lotka-Volterra dynamics. The ability to pinpoint the location of a system within this phase space and quantify the system's proximity to collapse is made possible via a novel mathematical analysis that we develop. Previous work that focused only on collapse lacks the context within the overall phase space to be able to predict when systems are nearing or are poised to collapse. Moreover, in contrast to previous collapse predictions, we account for the fact that dynamics often lead to single species extinctions. Allowing and accounting for these single-species extinctions reveals more detailed structure of the complexity-stability phase space and introduces an intermediate phase between stability and collapse – Extinction Continuum – that give a more nuanced view of how an ecosystem can respond to internal and external changes. With this extended phase space and our construction of predictive measures based strictly on observable quantities, real systems can be better mapped – than using canonical measures by May or critical slowdown – for proximity to collapse and path through phase-space to collapse.

System stability and collapse are core concepts for ecology and complex systems and have been studied both theoretically and empirically. Over the years many features of ecosystems have been posited as stabilising factors such as species interaction modules[1], long weakly-interacting trophic loops[2], nestedness[3], species body-size ratios[4], correlations in species interaction strengths[5], and trophic coherence[6]. The picture of stability that has emerged is multi-faceted. For example, some features are not purely stabilising or destabilising and can interact in non-trivial ways[7]. In addition, there are multiple perspectives on stability including resilience and resistance[8,9] as well as many ways to represent the web of interactions between species in an ecosystem (trophic, mutualistic/competitive, mixed).

A generic way of investigating ecosystem stability relies on dynamical models. For instance, analysing stabilising effects of Lotka-Volterra-type interaction modules with few species[10,11] or modelling species-rich systems with versions of Generalized-Lotka-Volterra (GLV) models where the species interaction-strengths are typically sampled from statistical distributions[12–14]. Pioneering work on stability was done by Robert May who used Random Matrix Theory to show the importance of system complexity in terms of number of species, number of interactions, and variability of interaction strength[15]. The conclusions of this work ran counter to the previous paradigm that complexity begets stability through functional re-

[1]Department of Space, Earth and Environment, Chalmers University of Technology, Maskingränd 2, 412 58 Gothenburg, Sweden.
[2]Department of Ecology and Evolutionary Biology, Department of Biomathematics, UCLA, Los Angeles, CA, 90095, USA.

dundancies or because of less reliance on a single keystone species[16,17]. This paradigm was based on observations that ecosystems with few species, like agricultural soils, could collapse due to large fluctuations in population abundances of pests. In contrast, species-rich and highly complex ecosystems, like the Amazon rain-forests, had not been observed to exhibit such large fluctuations[18].

A key component was to introduce a theoretic boundary between ecosystem stability and collapse. Closeness to the boundary implies a less robust system with a larger probability to collapse when perturbed–meaning that a seemingly small environmental change and/or a slight increase in interactions among species can instigate a system-wide collapse. Although the model unrealistically assumes random interactions among species, as May acknowledged, it has frequently been used as a null-model. Biologically-inspired structures – such as modularity[7] and predator-prey interactions[19] — can be added and their effects on stability estimated. Some of these features shift the collapse-boundary to higher complexities but are still unable to avoid eventual collapse.

Mapping phase-space and locating boundaries in terms of system properties — such as biodiversity, species interaction strengths, and topology – would provide a theoretical basis for understanding a system's response to changing conditions. Relating this to real systems also then requires locating it in the phase-space based on measurable properties. A substantial hurdle to this are the intricacies of measuring interaction strengths[20], vital components in dynamical models. More theoretical limitations are that most studies either exclude the possibility of species extinctions[7,11,14,19,21,22] or include extinctions but rely on an auxiliary species-pool instead of the actual biodiversity[23].

Excluding the possibility for species extinctions does not take persistence into account — the fraction of species out of an initial species-pool that are viable (i.e., non-extinct) at the locally stable fixed point that corresponds to specific values of diversity, density, interaction strength, and variability. Persistence studies[24,25] clearly show that not all species must be viable at the fixed point that is reached in the population model. This is a crucial result, as also noted by[26], because assuming a feasible fixed point means all species are viable and thus that the derived collapse boundaries only apply to a small subset of possible or real systems. Indeed, the vast majority of systems respond to increasing complexity with a few species extinctions. As we will show in this article, accounting for these species extinctions is important because it substantially shifts the boundary for system wide collapse and highlights a probable system response of continuous species extinction and loss of biodiversity.

Using the species-pool instead of the biodiversity inherently does not assume feasibility and can thus theoretically map phase-space. However, the collapse boundary is defined in terms of the initial species pool, while both the species-pool and persistence are unknowable quantities for a given ecosystem. Consequently, the assessment of a natural system e.g. proximity to collapse and critical response to perturbations, cannot be determined from phase space alone. This makes the theoretical phase space of limited use for understanding real systems.

Here, we develop a framework that overcomes both these limitations. We use large dynamical GLV models and allow for single-species extinctions that typically leads not to a single stability-collapse boundary, but instead to two boundaries. The first boundary is between Strict Biodiversity Stability (SBS), where perturbations do not lead to any extinctions, and an Extinction Continuum (EC) where perturbations lead to single-species extinctions and a new stable state. The second boundary is between the Extinction Continuum and Complete Collapse (CC) with possible chaotic dynamics, limit cycles or new fixed point with substantial loss of species. As such, the Extinction Continuum represents a new intermediate and ecologically important regime between these two boundaries.

Furthermore, based on the measurable quantities of a real system, we use novel techniques to estimate the persistence, the two boundaries, and proximity to collapse in the Extinction Continuum. Thus, we extend the current knowledge and toolkit by providing a framework and a metric that can predict a system's likely responses to change, both in terms of single-species extinction and system-wide collapse.

1 Theory

1.1 Phase-space

To predict responses of real systems, we calculate possible responses to perturbations, quantify the boundaries and regimes in the phase-space of parameters, and provide a formula in terms of empirically measurable quantities to place systems within this space. To capture all possible responses we analyse large GLV models without invoking the as-

sumption of feasibility. The GLV model is

$$\frac{dx_i}{dt} = x_i(r_i - \frac{x_i}{K_i}) + \sigma x_i \sum_{k=1}^{n} A_{ik} x_k. \quad (1)$$

where x_i are relative abundances and r_i and K_i are the intrinsic growth rate and carrying capacity for each species i. The A_{ij} encodes the interaction strengths–fluxes of bio-masses and other materials or processes–between all species. The probability that an off-diagonal entry of A is non-zero is c, and we sample interaction strengths from a normal distribution with mean $\mu = 0$ and a variance of 1. All diagonal entries (self-interactions) of A are zero. The number of species in the initial species-pool is n, and σ is a parameter modulating the interaction strength variability. When $\sigma = 0$ this represents a system with only self-interactions.

The fixed point abundances x_i^* of the system are

$$x_i^* = K_i \left(r_i + \sigma \sum_{k=1}^{n} A_{ik} x_k^* \right) \quad (2)$$

or

$$x_i^* = 0.$$

By not assuming feasibility (dropping the fixed points with zero abundance), we allow for single-species extinctions and can calculate persistence when mapping the region of local-stability for the GLV model. In Fig. 1a we show how the fixed points for a system with specific interactions (A, K, and r) change as the overall variance for interaction strengths increases. Note the single-species extinctions occur as a response to increased complexity (measured here by σ while system size and number of interactions are fixed for each simulation).

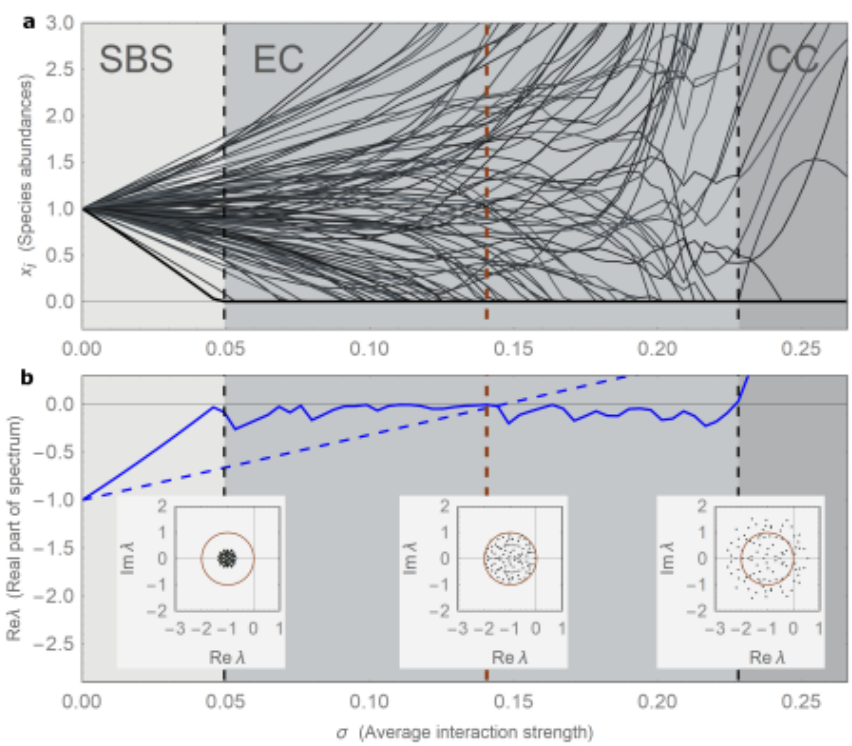

Figure 1 | Effects on stability of increasing interaction strength in a complex system. (**a**,**b**) Example simulation of a system with number of species, $n = 100$, connectance (fraction of realised species interactions), $c = 0.5$, and $\mu = 0$ for the mean of the distribution of inter-specific interaction. Plot shows the species abundances (**a**) and leading eigenvalues (**b**) at stable fixed points for increasing values of interaction strength variability, σ. Three phases of system behaviour, including the phase–Extinction Continuum–are indicated by the different shades of grey. Up to the first extinction, the system is in a stable feasible fixed point (all $n = 100$ species are viable, $x_i > 0$). After the first extinction, the system enters the phase of dynamic self-regulation of complexity via successive extinction events. This phase passes May's stability limit (brown dashed line) and continues until the fixed points lose stability altogether. The leading eigenvalues clearly display the difference in stability behaviour between critical slowdown and the dynamics including extinctions (**b**). The dynamical model's leading eigenvalue (blue line) fluctuates below zero due to single-species extinctions, while the leading eigenvalue models assuming feasibility (blue dashed line) increases linearly and passes through zero at $\sigma \approx \sqrt{cn}$, as predicted. For further clarification, the insets show the eigenvalues for May's model at the transition points of the system, the brown circle indicates that maximum radius of stability.

The local stability (henceforth called stability) of a dynamical system at its fixed points, such as in Fig. 1a, is determined by the eigenvalues of the Jacobian (sometimes referred to as the community matrix). For the GLV, the Jacobian is

$$J_{ij} = \delta_{ij} \left(r_i - \frac{2x_i}{K_i} + \sigma \sum_{k=1}^{n} A_{ik} x_k \right) + \sigma x_i A_{ij}, \quad (3)$$

where δ represents the identity matrix, i.e. δ_{ij} equal to one when $i = j$ and zero otherwise.

For stability the real parts of all eigenvalues must be negative, which implies that after a small disturbance the system will return to the exact same state (*i.e.*, same population abundances for each species). When at least one eigenvalue has a positive real part the system is unstable because a small disturbance will increase exponentially in amplitude.

Changes that cause the real part of at least one eigenvalue to approach zero are therefore usually considered to cause destabilisation, often interpreted as a preamble to collapse. Mathematically, however, collapse is not necessarily implied. Indeed, in the Extinction Continuum the leading eigenvalue fluctuates at or just below zero (Fig. 1b). The key observation is that these single-extinction events actually stabilise the system and prevent collapse.

The extinctions and stability aspects of the GLV model are demonstrated in Fig. 1. Here we see that the system response divides into three distinct phases when we use the interaction variability σ as the phase-space parameter. The first phase Strict Biodiversity Stability is characterised by the existence of a stable feasible fixed point and negative (but increasing with σ) value of the eigenvalue with the largest real part. Systems located in this phase will therefore return to their initial states if perturbed or respond with a slight shift in abundances if the change remains. The boundary of the Strict Biodiversity Stability phase is located where the leading eigenvalue reaches zero and the first single-species extinction occurs.

In the second phase, Extinction Continuum, the system can uphold stability with a leading eigenvalue at or below zero but for a subset of the original species-pool. If we track the system in phase-space when increasing σ, extinctions occur as successive single-extinction events while a reduced viable community stays stable. This implies that perturbations of systems in this region typically causes extinctions but not a system-wide disruption and slowly changing external circumstances can significantly change the species composition of a system. At the end of this phase there can be multiple fixed points[27], meaning that out of all the possible fixed points of subsets of non-extinct species, more than one can be stable. However, even in the Extinction Continuum, the simulated systems almost always converge to a single stable fixed point (which agrees with findings in[27]). We use this single fixed point as the focal point of our analytical calculations.

In the third phase, Complete Collapse, no stable fixed point is found and the system's behaviour is unpredictable. There is the possibility of limit cycles, chaotic behaviour, or a large decrease in the number of viable species

With Fig. 1 in mind it is clear that assuming feasibility amounts to incorrect predictions of collapse and ignoring of single-species extinctions as a stabilising mechanism. Robert May's original result states that a random system is stable when $\sigma < 1/\sqrt{cn}$ which gives that fixed points are unstable for $\sigma > 1/\sqrt{cn}$. May's prediction based on the original number of species actually falls within the Extinction Continuum in Fig. 1 and hence does not predict collapse.

It is interesting to note, however, that the actual system collapse occurs when the reduced system of viable species (n^*) reaches $\sigma > 1/\sqrt{cn^*}$, (with a slight upward bias), better corresponding to an effective version of May's metric. This further implies that the prediction still holds even though the systems that collapse at this boundary no longer have uncorrelated species interactions. Instead, the systems have been dynamically pruned in accordance to dynamics dictated by growth rates r_i and carrying capacities K_i. Out of all possible and reasonable values and combinations of r_i and K_i, only a small number can tolerate completely viable communities for large σ. Moreover, that number decreases exponentially as system size increases[28].

The Extinction Continuum and the accompanying behaviour of the leading eigenvalue is relevant when trying to predict collapse with critical slowdown – the standard method for proximity detection, i.e. that the system starts to recover very slowly after perturbations. In our analysis systems show critical slowdown throughout the Extinction Continuum, but the slowdown only involves one species at the time (the species that will go extinct next). The major transition and predictor of proximity to system collapse however is when the system is close to the collapse-boundary and the critical slowdown involves all species at once (not just a single species). We will show that our new measure can be used to well predict the location of the collapse early in the Extinction Continuum, when the system is still far away from the actual collapsing point in parameter space. This can quantify whether a slow recovery rate is an early warning sign or looming danger.

1.2 First extinction

To construct a measure for predicting collapse, we derive the two phase boundaries introduced that surround the Extinction Continuum. The first is

the boundary that terminates the Strict Biodiversity Stability region in which small perturbations in the GLV parameters will not cause any extinctions and the Extinction Continuum region where the system can react to small parameter perturbations with single-species extinctions. The transition is marked by the first extinction event, corresponding to when the real part of leading eigenvalue hits but does not pass zero.

To locate the first extinction boundary we derive a novel method for analysing the interaction matrix A based on order statistics. This gives a direct estimate of the first extinction event rather than implicit methods[27].

Because we are interested in the boundary σ_f of the Strict Biodiversity Stability phase we can safely look only at feasible fixed points ($x_i^* > 0$) Eq. 2. These non-trivial fixed points can be written as linear equations, and for the case $r_i = K_i = 1$, they are

$$(I - \sigma A)\,\mathbf{x}^* = \mathbf{1} \tag{4}$$

where the 1 is a vector of ones, I is the identity matrix, and $\mathbf{x}^*$ is a vector with fixed point species abundances. The more general case for arbitrary r_i and K_i is given in the supplementary 1.1.

To simplify our analysis, we separate the interaction matrix into $A = A_0 + \mu M$ where the distribution of the entries in A_0 have mean zero and the entries $M_{ij} = 1$ for $A_{ij} \neq 0$. (In the case $c = 1$ when all possible interactions are realised, M is a matrix of ones with zero diagonal). The fixed-point equation can then be written as

$$(I - \sigma A_0 - \sigma\mu M)\,\mathbf{x}^* = \mathbf{1}. \tag{5}$$

We begin with the case $c = 1$, so $M\mathbf{x}^*$ can be approximated with a constant vector $x_{TOT}\mathbf{1}$. This is an approximation since the sum for each row excludes the diagonal element, but it is a good one since the variance of the species abundances x_i^* before first extinction is small. The constant vector $x_{TOT}\mathbf{1}$ can be moved to the other side of Eq. 5 and grouped with 1. We can then solve for the fixed point abundances

$$\mathbf{x}^* = (\sigma\mu x_{TOT} + 1)\,(I - \sigma A_0)^{-1}\,\mathbf{1}. \tag{6}$$

For the inverse a von Neumann series expansion ($(I - B)^{-1} = \sum_{m=0}^{\infty} B^m$ for some matrix B with approximately $|B| < 1$) to first order gives

$$\mathbf{x}^* \approx (\sigma\mu x_{TOT} + 1)\,(I + \sigma A_0)\,\mathbf{1}. \tag{7}$$

Because the entries A_0 are drawn from a random distribution related to interaction terms, the species abundances x_i^* can be treated as stochastic variables. We thus get a set of n stochastic variables

$$x_i^* \approx (\sigma\mu x_{TOT} + 1)\left(1 + \sigma \sum_{j=1}^{n} A_{ij}\right). \tag{8}$$

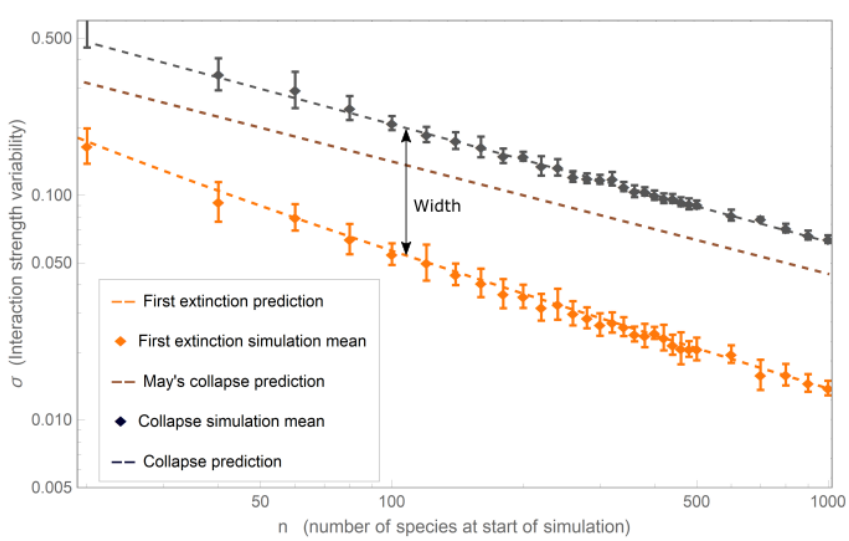

Figure 2 | Stability predictions for complex systems. There are three phases of perturbation behaviour in the biodiversity complexity phase space: Strict Biodiversity Stability, Extinction Continuum, and Complete Collapse. Here we demonstrate that these phases hold across a large range of system sizes n. The plot shows simulation averages of first extinction events (orange dots) with one standard deviation error bars, our theoretical prediction of first extinction (orange dashed line), May's prediction of instability (brown line), simulation averages of final instability (black dots) with one standard deviation error bars, and our theoretical collapse prediction (black dashed line). The width of the Extinction Continuum is indicated by the arrow, note the increase in width for larger systems. All simulations were run with $\mu = 0$ for the distribution of inter-specific interactions and a value of $c = 0.5$ for connectance in the interaction matrix A.

Now that we have estimates of species abundances we can estimate the smallest σ at which there exists an $x_i^* \leq 0$. Because we assumed $c = 1$, a nonzero μ only adds a scaling factor to x_i^*, hence it will not affect our prediction for when the first species abundance x_i^* becomes zero. Cases when $c < 1$ introduce only a slight bias. Although we can no longer treat the sum of abundances $\sum_j x_j^*$ in Eq. 7 as a constant x_{TOT}, we can then treat it as a stochastic variable but with small variance and mean for $\sigma\mu x_{TOT}$ so that this bias is negligible. Thus for small $|\mu|$ the scaling factor x_{TOT} does not impact our estimate of first extinction.

The distributions of the x_i^* are row sums of the interaction matrix A_0. For a normal distribution of the interaction strengths or any distribution so long as the product of biodiversity and connectance is large ($cn >> 1$) enough to invoke the central limit theorem, the sums are normally distributed

$$\begin{aligned} x_i^* &= 1 + \sigma\left((A_0)_{i1} + (A_0)_{i2} + ...(A_0)_{in}\right) \\ &\sim \mathcal{N}(\mu_+ = 1, \sigma\sqrt{cn}\,), \end{aligned} \tag{9}$$

where μ_+ and $\sigma^2 cn$ are the new means and variances for the sums. The extra cn, appearing in variance is the mean of the Binomial distribution of n trials with success probability c.

From order statistics we get the distribution for the minimum x_i^* ($f_{min}(x)$) in the set of abundances from the distribution function of x_i^* ($f(x)$) and its cumulative distribution function, $F(x)$,

$$\begin{aligned} f_{min}(x) &= n(1 - F(x))^{n-1} f(x) \\ &= \frac{ne^{-(x-\mu_+)^2/2\sigma^2 cn}}{\sigma\sqrt{2\pi cn}} \left(\frac{1}{2} - \frac{1}{\sqrt{\pi}} \int_0^{\frac{x-\mu_+}{\sigma\sqrt{2cn}}} e^{-t^2} dt\right)^{n-1}. \end{aligned} \tag{10}$$

The first extinction event is then predicted at the σ_f for which the mean of the above distribution is zero. This prediction is in well agreement with simulations as shown in Fig.2.

1.3 Collapse boundary and persistence

To locate the collapse boundary we use a novel mathematical method connected to persistence. We estimate persistence by the non-trivial fixed point solutions of $\mathbf{x}^* = (KI - \sigma A)^{-1}\mathbf{r}$, corresponding to specific values of r_i, K_i, and other GLV parameters. For systems of size n with values of $\sigma > \sigma_f(n)$ – larger interaction variability than that for the predicted first extinction boundary – the non-trivial solution will include negative x_i. Although the positive x_i are not the fixed point abundances of a reduced viable system of size n^*, the number of non-negative abundances can be used for an approximation of the persistence. In Fig. 3 we show the statistics of prediction of the persistence obtained from the non-negative abundances compared to the actual persistence for systems of varying sizes ($n = 20 - 1000$).

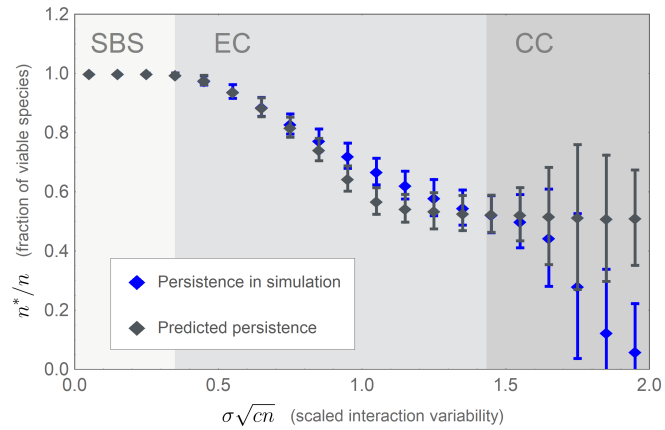

Figure 3 | Predicted persistence. The plot shows the average persistence from simulations of systems ranging from $n = 20$ to $n = 1000$ (blue dots) with one standard deviation errorbars and predicted persistence for randomly generated A in the same system size range (grey dots) with one standard deviation errorbars. The background shading indicates the three phases of phase-space with the boundaries at the theoretical predictions. The difference between predicted persistence and simulated persistence in the CC phase is because we set a systems size to zero after collapse in simulations causing the statistics of the fraction n^*/n to drop. Before this artifical drop in the CC phase we can predict the persistence well in the Extinction Continuum.

The reduced system with non-negative abundances and persistence opens up two new ways of predicting collapse. First, we use the reduced interaction matrix of non-negative species $\sigma A(n^*) - I(n^*)$, and calculate collapse as the σ where the reduced system loses stability. Second, we find the σ at which the persisting species n^* reach May's limit, which is possible since we know that May's stability limit is a good predictor of collapse when assuming feasibility. For both these choices, statistics were collected and their average used as a predictor of collapse.

The collapse prediction and persistence together with first extinction means that, when the size of the species-pool is known, we can predict the two boundaries in phase-space in a direct way for any system size, as seen in Fig. 2, and its approximate rate of extinctions in the Extinction Continuum. Fig. 2 displays simulation averages and theoretical estimates of first extinctions along with May's stability limit and the boundary to collapse for systems as a function of species-pool size. Our analytic

estimates are in good agreement with the results of numerical simulations.

These predictions hold not only for systems with a standard setup of random interactions between species (no imposed structure), interaction strengths sampled from a normal distribution with a mean of zero and variance of one, and of specific sizes as shown in Fig. 2. In Figs. 1-5 in supplementary 1.2 we show that there are no qualitative differences in results even if any of the assumptions on distribution, structure, mean, r_i, and K_i respectively are modified. The robustness of the Extinction Continuum follows from the fact that there are exponentially many fixed points that the system can switch between to uphold stability. Indeed, for this reason we expect the Extinction Continuum to exist in more general population models than the GLV.

1.4 Instability metric

As stated previously it is not enough to map phase space by using the species-pool since its size is not necessarily equal to the actual biodiversity of a system (it can reside in the Extinction Continuum). The lack of species-pool information when working with real systems, prompts us to find a way to connect our theoretical results to a system's measurable quantities. This we can do by estimating the species-pool size n from a system of size n^* with our use of persistence. With an estimated species-pool n_{pred} we can properly place the system in phase-space and calculate its proximity to collapse.

A first thing to note is that a system with $\sigma < \sigma_f(n^*)$ can directly and immediately be located in the Strict Biodiversity Stability phase space. If $\sigma \approx \sigma_f(n^*)$ or $\sigma > \sigma_f(n^*)$ the system likely resides in the Extinction Continuum and we need to approximate the size of the species-pool. Again, we use the amount of non-negative abundances in the non-trivial solution to Eq. 2 this time to obtain the rate of extinctions for increasing σ. We use this rate to extrapolate backwards to σ_f for the approximation n_{pred} (more details in supplementary 1.3).

With the approximate species-pool size we retrieve the two boundaries in relation to the system and construct the measure γ to quantify the proximity to collapse

$$\gamma(n^*) = \frac{\sigma - \sigma_f(n_{pred}(n^*))}{w(n_{pred}(n^*))}. \tag{11}$$

Where $w(n_{pred}(n^*))$ is the width of the Extinction Continuum $\sigma_c(n_{pred}(n^*)) - \sigma_f(n_{pred}(n^*))$. The measure γ is a real number $\in [0, 1]$, although the boundaries of this interval are not exact since we are inferring from n^*, that quantifies the robustness of the system. We posit that the higher the value of γ the more likely a perturbation or external pressure will lead to collapse.

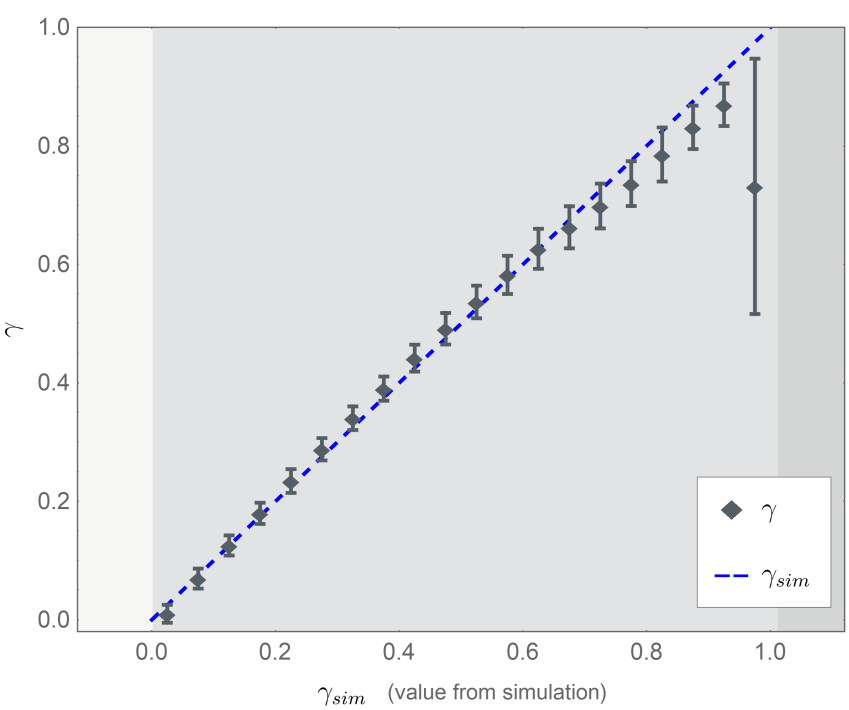

Figure 4 | Maesure of stability of a complex system. The plot shows the γ (dark dots) with one standard deviation errorbars compared to γ_{sim} (blue dashed line). The γs were calculated from simulations ranging between $n = 20$ and $n = 1000$ and 10 repeats on each. The background shading indicates the three phases of behaviour of the phase-space. Here we see that our measure of proximity to collapse is within one standard deviation of the correct value right up until collapse, which means it is a good predictor of position in phase-space.

2 Results

2.1 Proximity to collapse

To evaluate the γ metric we simulated systems with size ranging from $n = 20$ to $n = 1000$ for values of σ placing the system in the Extinction Continuum. We calculated γ using available information of the viable species: the reduced interaction matrix A^*, effective community size n^*, and $x_i^* > 0$. The analytic γ measure from these simulations are shown in Fig. 4 in comparison to the actual measure from simulations γ_{Sim}.

It can be seen in Fig. 4 that γ follows the actual γ_{sim} closely. The variance of the prediction increases during the continuum which is to be expected, but the actual value is within one standard deviation right up until collapse.

All simulations for Fig. 4 were performed for a standard setup with $r_i = K_i = 1$, no symmetries in A and the interactions strengths drawn from a Normal distribution with zero mean. For small deviations from zero of the mean ($|\mu| < 0.05$) the γ measure remains a reliable approximation. For larger deviations the prediction of collapse becomes more complicated due to the influence of x_i^* on the actual Jacobian of the system ($X(\sigma A - I)$ instead of $(\sigma A - I)$), a problem shared with May's prediction and discussed further in supplementary 1.4.

As an additional verification that γ quantifies proximity to collapse we gathered statistics for stability of viable systems before and after perturbations to the interaction matrix for varying γ. The perturbation was an addition of Gaussian noise ($\mathcal{N}(0, \sigma_{noise})$ with $\sigma_{noise} = 0.25, 0.5, 0.75$) to all the realised interactions in the interaction matrix, meaning the density or topology of interactions was not altered nor the stabilising self-interactions, but the variability in the interaction strength distribution and relative interaction strengths. The results from these simulations can be seen in Fig. 5. The top panel shows the percentage of systems that find a stable configuration after the perturbation and the bottom panel shows the relative size of the stable communities after the perturbation. The first thing to note is that the stable communities after perturbation are consistently of smaller size, even for small perturbations Fig. 5b. It is also evident that systems with larger γ closer to collapse loose more species before finding stability. This is consistent with our previous observation that correlations are introduced in the dynamic pruning process that are needed for stability. It is also clear from Fig. 5a that systems with a large γ are more likely to collapse after a perturbation. Both these findings verifying γ as a metric for collapse.

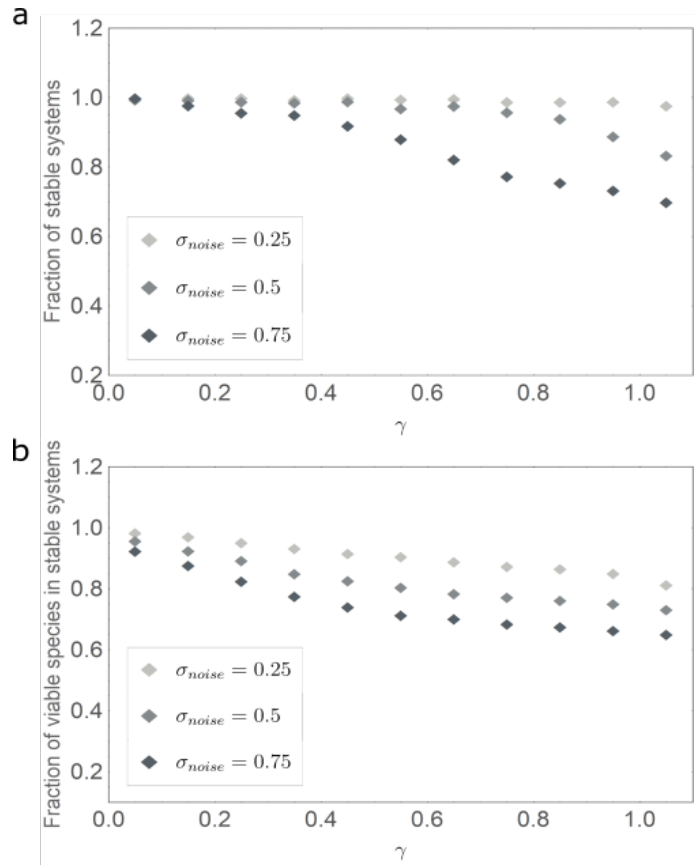

Figure 5 | Predicting collapse. The plot shows in the top panel the fraction of stable systems after perturbations for certain γ values generated from random systems with $n = 70$ and connectivity $c = 0.5$. The bottom panel shows the fraction of viable species in the systems that found a stable configuration after perturbation. Here we see that even for small perturbations in the Extinction Continuum some species are lost when finding a new stable configuration. This effect increases for larger γ indicating more specificity of interactions in more dynamically pruned systems. Together the plots demonstrate that a larger γ indicates collapse both in terms of a substantial loss of species and a higher probability of loss of stability altogether.

2.2 Extinctions for decreasing complexity

As we have emphasised previously the inclusion of extinctions reveals a wider range of system behaviour in response to perturbation or pressure. One interesting result is that in the Extinction Continuum extinctions occur also for decreasing interaction variability σ –corresponding to decreasing complexity – in contrast to almost all other complexity-stability measures. Both systems with random interaction matrices but with intrinsic growth rates r_i specifically chosen to make

them stable and feasible at May's stability limit $\sigma \approx 1/\sqrt{cn}$ and our reduced viable systems can exhibit species extinctions when σ decreases (supplementary 1.5 Fig. 7).

One interpretation of this response is that for systems to be stable at complexities high enough to reside in the Extinction Continuum correlations or interaction patterns have to be in place to accommodate stability. These correlations in turn seem to be stabilising at specific complexities rather than stabilising up to a certain complexity limit. This makes them sensitive not only to an increase in complexity but also decreases.

3 Discussion

When interacting with and studying ecosystems it is helpful to have a catalogue of possible responses a system can exhibit when disturbed, and a way of assessing which of those responses are likely for the system at hand. In this study our aim is to provide such a metric. The γ measure is a quantification of the proximity to collapse in a phase-space corresponding to likely behavioural responses, which can be calculated based on data/knowledge of real systems.

We have extended the response repertoire of a perturbed ecosystem from stability and collapse to also include single extinctions as a stabilising mechanism. A smaller stable community is sustained through extinctions when the interaction variability is increased since the extinction are not random and thus introduce correlations between the r_i, K_i, A_{ij} and x_i^* in the system. In[29] they find correlations between A_{ij} and A_{ji} in dynamically pruned matrices and also correlations between the interaction strengths and abundances x_i^* which they find to have the largest stabilising effect. Although these correlations are not necessarily solely stabilising since they might increase the sensibility to other perturbations, as in the case of extinctions for decreasing σ. What still holds true though from the former single-boundary phase-space to our elaborated phase-space is that random systems cannot be stable above a certain complexity level. Although the limit is no longer the collapse limit by May but the first extinction boundary between the Strict Biodiversity Stable phase and the Extinction Continuum, all more complex systems need some structure or correlations to uphold stability.

We have shown that our γ measure constructed from observable quantities of an ecosystem captures proximity to collapse in the σ phase-space. Since changes in the GLV parameters (n, c, r_i, K_i, distribution of interaction strengths, and symmetries in interactions) induce "mere" shifts in boundaries we argue that γ not only quantifies likelihood of collapse in terms of σ perturbations but quantifies structural stability in general. This makes γ a much wider measure than its definition (Eq. 11) might indicate. Although we cannot rule out possible synergistic phenomena of combinations of GLV parameter compared to the standard set-up that would change the phase-space structure, we have not seen any indications of such synergies in simulations.

Even though we posit a wide scope for γ there are of course routes we have not explored and limitations of applicability. One such route is investigating if the same pattern of extinctions as a stabilising mechanism remains and the impact on γ in case of more biologically informed topologies in the interaction matrix. Besides from symmetric (mutualistic/antagonistic) and anti-symmetric (predator/pray) structures we have not investigated topology. Thus introducing more realistic structures such as trophic hierarchies into the interaction matrix is an important direction for future work.

Another limitation to our approach, the GLV approach in general, and most studies on these topics, is that the interaction matrices are static "averages" of interactions among species. While in reality interactions might fluctuate with external factors such as seasons. Since such seasonal changes in general do not lead to extinctions, it is assumed that they do no not introduce fluctuations large enough to breach the limits of structural stability for systems in the Extinction Continuum. Still, we do not know if cycles of interaction strengths play a role in the stability of communities. A further limitation regarding fluctuations is that we focus on fixed points, corresponding to communities with relatively stable populations over time. The γ metric does not capture communities for which the abundances have significant abundance cycles.

In conclusion we have shown that with the use of theoretical analysis it is possible to quantify a system's proximity to collapse and position in phase-space based solely on measurable information such as biodiversity, density, and strengths of species interactions. We provide a consistent framework to quantify multiple behaviours of a complex ecosystem and derive a new measure of proximity to collapse that is both detectable at larger parameter-"distances" from collapse than previous methods and also easier to evaluate as compared, for example, to critical slowdown. Together the framework and collapse-metric expand our ability to estimate a system's likely response to perturbations and assess systems that are at risk of collapse.

References

1. Stouffer, D. B. & Bascompte, J. Understanding food-web persistence from local to global scales. *Ecology letters* **13,** 154–161. ISSN: 1461-0248 (2010).
2. Neutel, A.-M., Heesterbeek, J. A. & de Ruiter, P. C. Stability in real food webs: weak links in long loops. *Science* **296,** 1120–1123. ISSN: 0036-8075 (2002).
3. Bastolla, U. *et al.* The architecture of mutualistic networks minimizes competition and increases biodiversity. *Nature* **458,** 1018. ISSN: 1476-4687 (2009).
4. Brose, U., Williams, R. J. & Martinez, N. D. Allometric scaling enhances stability in complex food webs. *Ecology letters* **9,** 1228–1236. ISSN: 1461-0248 (2006).
5. Tang, S., Pawar, S. & Allesina, S. Correlation between interaction strengths drives stability in large ecological networks. *Ecology letters* **17,** 1094–1100. ISSN: 1461-0248 (2014).
6. Johnson, S., Domínguez-García, V., Donetti, L. & Muñoz, M. A. Trophic coherence determines food-web stability. *Proceedings of the National Academy of Sciences* **111,** 17923–17928. ISSN: 0027-8424 (2014).
7. Grilli, J., Rogers, T. & Allesina, S. Modularity and stability in ecological communities. *Nature communications* **7,** 12031. ISSN: 2041-1723 (2016).
8. Grimm, V. & Wissel, C. Babel, or the ecological stability discussions: an inventory and analysis of terminology and a guide for avoiding confusion. *Oecologia* **109,** 323–334. ISSN: 0029-8549 (1997).
9. Ives, A. R. & Carpenter, S. R. Stability and diversity of ecosystems. *science* **317,** 58–62. ISSN: 0036-8075 (2007).
10. McCann, K., Hastings, A. & Huxel, G. R. Weak trophic interactions and the balance of nature. *Nature* **395,** 794. ISSN: 1476-4687 (1998).
11. Pimm, S. L. & Lawton, J. Number of trophic levels in ecological communities. *Nature* **268,** 329–331. ISSN: 0028-0836 (1977).
12. Bairey, E., Kelsic, E. D. & Kishony, R. High-order species interactions shape ecosystem diversity. *Nature communications* **7,** 12285. ISSN: 2041-1723 (2016).
13. Allesina, S. & Tang, S. Stability criteria for complex ecosystems. *Nature* **483,** 205–208. ISSN: 0028-0836 (2012).
14. Mougi, A. & Kondoh, M. Stability of competition–antagonism–mutualism hybrid community and the role of community network structure. *Journal of theoretical biology* **360,** 54–58. ISSN: 0022-5193 (2014).
15. May, R. M. Will a large complex system be stable? *Nature* **238,** 413–414. ISSN: 0028-0836 (1972).
16. MacArthur, R. Fluctuations of animal populations and a measure of community stability. *ecology* **36,** 533–536. ISSN: 1939-9170 (1955).
17. Paine, R. T. A note on trophic complexity and community stability. *The American Naturalist* **103,** 91–93. ISSN: 0003-0147 (1969).
18. Elton, C. S. *The ecology of invasions by animals and plants* ISBN: 0226206386 (University of Chicago Press, 1958).
19. Allesina, S. & Pascual, M. Network structure, predator–prey modules, and stability in large food webs. *Theoretical Ecology* **1,** 55–64. ISSN: 1874-1738 (2008).
20. Paine, R. T. Road maps of interactions or grist for theoretical development? *Ecology* **69,** 1648–1654. ISSN: 1939-9170 (1988).
21. Allesina, S. & Tang, S. The stability–complexity relationship at age 40: a random matrix perspective. *Population Ecology* **57,** 63–75. ISSN: 1438-3896 (2015).
22. Gibbs, T., Grilli, J., Rogers, T. & Allesina, S. The effect of population abundances on the stability of large random ecosystems. *arXiv preprint arXiv:1708.08837* (2017).
23. Bunin, G. Ecological communities with Lotka-Volterra dynamics. *Physical Review E* **95,** 042414 (2017).
24. Stouffer, D. B. & Bascompte, J. Compartmentalization increases food-web persistence. *Proceedings of the National Academy of Sciences* **108,** 3648–3652. ISSN: 0027-8424 (2011).
25. Saavedra, S., Stouffer, D. B., Uzzi, B. & Bascompte, J. Strong contributors to network persistence are the most vulnerable to extinction. *Nature* **478,** 233. ISSN: 1476-4687 (2011).

26. Dougoud, M., Vinckenbosch, L., Rohr, R. P., Bersier, L.-F. & Mazza, C. *The feasibility of equilibria in large ecosystems* Report (2018).

27. Bunin, G. Interaction patterns and diversity in assembled ecological communities. *arXiv preprint arXiv:1607.04734* (2016).

28. Grilli, J. *et al.* Feasibility and coexistence of large ecological communities. *Nature communications* **8,** 0. ISSN: 2041-1723 (2017).

29. Biroli, G., Bunin, G. & Cammarota, C. Marginally stable equilibria in critical ecosystems. *New Journal of Physics.* ISSN: 1367-2630 (2018).

Author Contributions Conceived the project: M.N.J.; Performed simulations and wrote paper: S.P.; Analytic results: M.N.J., S.P.; Interpretation of results and comments on manuscript: M.N.J., S.P., V.M.S.

Author Information The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to S.P. (susannep@chalmers.se).