The cavity method for community ecology

Matthieu Barbier¹ and Jean-François Arnoldi¹

June 23, 2017

Abstract

This article is addressed to researchers and students in theoretical ecology, as an introduction to “disordered systems” approaches from statistical physics, and how they can help understand large ecological communities. We discuss the relevance of these approaches, and how they fit within the broader landscape of models in community ecology. We focus on a remarkably simple technique, the cavity method, which allows to derive the equilibrium properties of Lotka-Volterra systems. We present its predictions, the new intuitions it suggests, and its technical underpinnings. We also discuss a number of new results concerning possible extensions, including different functional responses and community structures.

Contents

1	Context	2
1.1	Disordered systems	2
1.2	Models of ecological dynamics	4
1.2.1	The disordered systems framework	4
1.2.2	Example: History dependence	6
2	General framework	7
2.1	The assembly process	7
2.2	The assembled state	7
2.2.1	Equilibrium relations	7
2.2.2	Control parameters	8
2.2.3	The special case of competitive communities	9
3	What can we predict?	10
3.1	Functioning	11
3.2	Diversity and coexistence	11
3.3	Stability	13
3.4	Selected features	14
3.5	Discussion	16

¹Centre for Biodiversity Theory and Modelling, Theoretical and Experimental Ecology Station, CNRS and Paul Sabatier University, 09200 Moulis, France.

4	Mathematical setting	16
4.1	Scaling of interactions: when are large disordered communities not neutral?	19
4.1.1	Inter-species versus intra-species effects	19
4.1.2	Comparing communities of different sizes and connectivities	19
4.1.3	Tradeoffs and global stability in large systems	21
4.2	Couplings rather than interactions	22
4.2.1	Allowing heterogeneous self-interactions	22
4.2.2	Reinterpreting γ	23
4.3	Computing the assembled state	24
5	Detailed calculations	26
5.1	The cavity trick	27
5.2	Universality of cavity results	29
5.3	Abundance distribution	31
5.3.1	Correlations between K_i and A_{ij}	32
5.3.2	Relationship to empirics	33
5.4	Selection by assembly: the case of carrying capacities	34
5.5	Productivity	35
5.6	Response to perturbations	36
5.6.1	Multistability order parameter	36
5.6.2	Secondary extinctions	37
5.6.3	Variability	37
6	Sensitivity testing	38
6.1	Distributions	38
6.2	Interaction strength	39
7	Extending the model	39
7.1	Functional response	40
7.2	Heterogeneous means	43
7.3	Species groups	44
7.4	Ordered species	45
7.5	One-sided order	46
7.5.1	Exponential order	46
7.5.2	Linear order	48
8	Appendices	52

1 Context

1.1 Disordered systems

Many foundational insights of ecological theory were derived from the careful study of simplified patterns of interaction between a few distinct entities (individuals or species). Examples include trophic chains [20], competition and

exclusion over a few resources [21, 39], mutualism [13] and parasite-host interactions [25]. These archetypal structures can be used as modules [15, 5] to understand more complex systems, when this complexity does not interfere with their order and structure. For instance, our intuitions on multitrophic communities combine simple trophic mechanisms to understand inter-guild relations, and simple competitive mechanisms for intra-guild dynamics. Many theories also distinguish between ecosystems shaped by top-down control from predators, and those shaped by bottom-up control from resources. In all cases, these intuitions can only work if these interaction structures remain distinct even in a complex community.

In parallel, neutral models [16] have shown that many quantitative predictions can be made with minimal assumptions of structure, or even of distinctiveness of the entities. In most such works, the only accounting for an ecosystem’s diversity is its number of otherwise identical species. These theories belong to the *mean-field* class, meaning that intrinsic differences between entities are erased statistically, owing to large numbers and some form of stochasticity. More precisely, such differences lose their importance in the collective dynamics – and therefore also in individual species’ dynamics, provided that those are mostly shaped by collective factors (e.g. aggregate competition from many other species).

Disordered systems stand between these extremes, in that they exhibit heterogeneity without order: they represent situations where one cannot neglect the existence of differences between species, but no clear-cut structural patterns can be identified. As in neutral theories, it is either impossible or unnecessary to assign clear “roles” to individual species. But instead of being modelled as identical, all species appear to be completely heterogeneous in their biological features and interactions, without regular motifs. What we can know is the range and frequency of these traits in the community as a whole, and we investigate how much can be predicted about an ecosystem from that limited knowledge. While the notion of disordered system originates in physics¹, perhaps the simplest analogy is found in models of neural networks [27]. Each neuron has a distinct set of interactions, but it is difficult to identify the role of any given neuron, or group them into functional classes; mainly, we can access summary statistics of “how different” neurons and their interactions are in a given network.

This approach shares a conceptual starting point with Random Matrix Theory, which has long become a staple of theoretical ecology [24, 1]. In both cases, one considers a “simple sort of complexity”: in large disordered communities, most details cease to matter and a few general mechanisms emerge, allowing us to unify many superficially different systems and models. However, Random Matrix Theory has developed with a specific focus on stability, whereas disordered systems techniques give access to properties – functioning, coexistence, abundance distributions, trait selection – which are usually predicted using neu-

¹It became prominent with the study of disordered materials such as glasses [29], before being applied to algorithmics and information theory [28], chemistry [30], economics [33] and systems biology [35].

tral or structured models. In fact, many ecological models already fall into this category, especially numerical models that contain many species with uncertain attributes (lacking empirical or mechanistic constraints) such as large random Lotka-Volterra communities. What has perhaps not been sufficiently realized is the availability of mathematical techniques for making broader and more robust statements about such systems, thus helping guide, validate or even bypass heavy simulation work. It is only recently that these techniques and ideas have been brought to bear on ecological problems [12, 18, 6, 37, 38].

Of particular interest is the application of this framework to *ecosystem assembly*, i.e. how dynamics and interactions shape the realized community from a pool of potential invaders. Indeed, it is easier to imagine that the pool itself is disordered: it may consist of species immigrating from a number of different environments, with independent evolutionary histories. By contrast, the assembled state that emerges from that pool will be less random. By following the process of assembly from that pool, it is possible to see which properties are selected by the ecological dynamics, i.e. how traits in the assembled community differ from those in the pool. Multiple works have thus explored what shapes species coexistence and invadability, in the context of assembly from a fixed regional pool [6], from immigration and stochastic extinction in an island biogeography perspective [18], or from two consumer communities “coalescing” – starting to compete on the same resources [37]. These methods have proven apt at generalizing to many species what was previously understood for modules of a few heterogeneous species, e.g. the conditions of competitive coexistence [38].

1.2 Models of ecological dynamics

1.2.1 The disordered systems framework

Practically speaking, the notion of disordered systems, and the fact that it allows for a simple and generic analytical approach, comes with significant methodological potential: it connects a class of commonly used models in ecology, synthesizing their insights and explaining why and when they may give similar or different predictions. We illustrate this in another work [3] with a “model meta-analysis” demonstrating that a variety of community structures, including mutualistic networks and resource competition, can still fall within the realm of disordered systems and create assembly patterns that we can predict with this framework.

This also means that if there were reasons to outright reject our results as a valid representation of ecological reality, then an entire class of models would become disqualified at once, as well as many qualitative arguments based on certain notions of complexity. That would be a useful result for ecological theory, limiting the field of investigation in models and simulations.

More concretely, to help situate this framework among preexisting work in theoretical ecology, we show in Fig. 1 a bird’s eye view of the main types of dynamical models we identified in the field. They roughly fall between three extremes. First, fully independent populations, where we use physiological and be-

Mathematical structure of ecological models

Figure 1: An abstract representation of the main classes of dynamical models in theoretical ecology, and the regime of collective organization where the notions and methods developed here are meant to apply.

havioral traits to predict existence and functioning within a given environment. Second, fully dependent (i.e. very strongly interacting) populations where, for instance, the coexistence of competitors requires precise tradeoffs between their traits. Third, fully stochastic models, where the characteristics of species and their interactions are all blurred and dominated by noise coming from internal (demographic) or external (environmental) sources. On this graph, our position is represented by the orange dot, a regime of collective organization where interactions are individually weak but collectively strong.

The two main axes of opposition on this graph are the following. First, deterministic versus stochastic: in deterministic approaches such as ours, we generally expect that real ecosystem properties can be understood from the equilibrium or near-equilibrium properties of a given dynamical model, controlled by some combination of intrinsic species fitness and by their interactions. The opposite pole includes theories that are dominated by stochastic and transient dynamics, where “noise” – in the demographics, the environment, or the immigration process – prevents the system from ever being close to some deterministic assembled equilibrium². Those include neutral theories [16] and island biogeography and

²The term of “equilibrium” is ambiguous as it has been used for the Theory Island Biogeography among others. A fully stochastic system can indeed be characterized by probability distributions that attain a stationary state – in physics, this is called a Non-Equilibrium Stationary State, or driven stationary state, to distinguish it from the concept of equilibrium in

its variants [22, 11, 18]. These two viewpoints are complementary; only careful confrontation to empirical data can decide whether an ecosystem’s dynamics are governed by transient, far-from equilibrium dynamics such as neutral turnover, or by the sort of near-equilibrium features described here.

The second axis of opposition confronts models whose dynamics are dominated by species characteristics, and those that are determined by interactions. Independent populations illustrate the former, competitive exclusion the latter. In these two extremes, results are often very sensitive to model assumptions: for instance, the colonization-competition tradeoff model ensures global coexistence by requiring a very precise ordering of competitive and dispersive ability. It may be surprising to see that the middle-ground – collective organization, where interactions are individually weak but collectively strong (because distributed among many species) – gives rise to very generic and predictable patterns.

This is the class of systems that we can describe with this framework. Note however that it could, fairly easily, be extended to include noise; by contrast, systems where species are strongly influenced by other individual species tend to have very dramatic dynamics (e.g. predator-prey cycles) which typically preclude the existence of any non-trivial equilibrium. Such dynamics are not entirely beyond the reach of disordered system approaches, but will almost certainly require more involved techniques than the simple one presented here.

1.2.2 Example: History dependence

Importantly, each class of models makes different predictions for the mechanisms of various ecological phenomena. We can illustrate it in the case of history-dependence (e.g. founder effects). In a deterministic equilibrium viewpoint, the long-term state of a system is dependent on its history only if its dynamics allow for multiple attractors. Of course, even with a single global attractor (e.g. a single final, stable and uninvadable community), if the time to reach it is very long, the system may retain some memory of its previous states at visible timescales. But this memory should become fainter as time passes, and any shock should only reset the dynamics to some earlier stage along a similar trajectory. This perspective is compatible with the intuitive notion of maturity³ as found e.g. in Margalef [23].

By contrast, in a stochastic perspective, any composition is to some degree an accident, though it may hold for a long time if there are positive feedbacks (e.g. self-reinforcing density dependence) that transform an early advantage into a long-lasting one. Nevertheless, the composition may well undergo a dramatic shift, due to a shock or even out of steady noise, without it signalling any change in intrinsic species fitness. This is closer to the usual intuition behind the founder effect.

deterministic dynamics and from thermodynamic equilibrium.

³Even with many local attractors, this picture may remain valid if the dynamics exhibit *aging*, i.e. the system is prone to falling into deeper and deeper attractors as time passes.

2 General framework

2.1 The assembly process

Here is the central conceit of this theory of assembly: we define a pool of species that can invade the community, and whose dynamical traits, such as growth rates and interactions, are known only statistically. We then allow sequences of invasions until the community reaches an uninhabitable state, and investigate the properties of this “assembled” state – among others, species diversity, abundance, and how traits selected by the dynamics differ from the pool baseline.

This can represent a variety of physical settings: the pool could be a entire region and the community a local patch [31], they could be a mainland and an island connected by migration fluxes [22, 11, 18], or even oceanic water and the microbiome of an individual sponge [36]. More generally, this suggests a stable reservoir of biodiversity with some known aggregate properties, and an environment that is being populated from that reservoir, even if the reservoir is not larger (e.g. a forest serving as a biodiversity shelter for surrounding plains), provided that it can serve as a stable supply of potential invaders for the timescales considered.

But this can also represent, more abstractly, a “separation of concerns” between, for example, physiology and ecology. Properties of the pool reflect the traits of species “in potential”; by contrast, the assembled community retains only the subset of traits that are dynamically selected to survive together. Thus, in an abstract model, the pool could represent the range of species that are allowed to exist by physiological constraints, while assembly would single out combinations of such species that can coexist in the same ecosystem. Even in the spatial interpretation suggested before, there is such a separation of concerns: species-species interactions are properties of the pool, but species-environment interactions (e.g. carrying capacities) could be specific to a location, leading to different assembled states in different locales.

This means that the size of the pool, S , admits very different interpretations depending on the ecological scenario represented: S could be regional diversity or mainland diversity, if one is interested in invasion through immigration, but it could also represent the mutant strains that try to invade a microbial community over a given timescale (the pool being disordered would then reflect an assumption of random mutations). Carefully teasing apart these various interpretations and the hypotheses they hide is one of our purposes in this article.

2.2 The assembled state

2.2.1 Equilibrium relations

Within this general setting, we decide to look only at the final assembled state, i.e. a feasible, stable and uninhabitable community. As discussed in Sec. 1.2.1, this stands in stark contrast with theories that are dominated by stochastic and transient dynamics.

The mathematical techniques of disordered systems can be applied to many different dynamics. For convenience, we investigate the very commonly used Lotka-Volterra dynamics⁴:

$$\frac{d}{dt}N_i = N_i \left(R_i - D_i N_i - \sum_{j \neq i}^S A_{ij} N_j \right) \quad (1)$$

Here A_{ij} is the interaction between species i and j (i.e. the rate at which a unit of species i affects a unit of species j dynamically). R_i is the growth rate of species i , that combines growth on external resources and mortality (and could thus be negative). D_i is intra-species competition, hereafter *self-interaction*. We can define

$$D_i = R_i/K_i \quad (2)$$

where K_i is properly the carrying capacity of species i in the case of consumers with $R_i > 0$. If $R_i < 0$, then K_i is negative and cannot be interpreted as a carrying capacity, but remains an important mathematical quantity (a measure of density-dependent mortality).

We will concern ourselves only with the equilibrium solutions, meaning $dN_i/dt = 0$ and

$$0 = R_i - D_i N_i - \sum_{j \neq i}^S A_{ij} N_j. \quad (3)$$

Only the species for which a positive solution $N_i > 0$ exists will then survive together in an assembled state. This state is obtained after long immigration sequence, during which all species that could invade did. As we explain below in Sec. 3.3, in a certain region of parameters, the precise sequence of invasions is irrelevant to the final state, which is a global attractor of the dynamics.

2.2.2 Control parameters

Important: To simplify the exposition, we assume starting now that all the self-interaction terms are equal $D_i = \langle D \rangle = 1$, so that $K_i = R_i$. This assumption is relaxed and the consequences of heterogeneous D_i are discussed starting in Sec. 4.

We will proceed to show that, to predict many properties of the assembled state, it is sufficient to know four simple statistical features of the pool of potential invaders. First, we need a measure of intrinsic fitness differences between species; for this we take ζ the standard deviation of the growth rates R_i . Second, we need to characterize interactions at a collective level. We thus ask how much a species is affected by the *sum* of all its interactions: how adverse this total interaction is on average, represented by a parameter μ ($\mu > 0$ means that

⁴Note however that in Sec. 7.1, we extend these results to a saturating functional response.

biomass is on average lost through interactions, $\mu < 0$ means it is gained), and how much that varies from species to species, represented by a parameter σ . Hence, for any species, the total effect of interactions can be understood to fall roughly in the range $\mu \pm \sigma$. Finally, to understand feedbacks between species, we need to know how similar the effect of species i and j is to the reciprocal effect of j on i , and we encapsulate this in a parameter γ .

Formally, these parameters are defined as follows:

$$\begin{aligned}\zeta^2 &= \langle R_i^2 \rangle - \langle R_i \rangle^2, & \mu &= s \langle A_{ij} \rangle, \\ \sigma^2 &= s(\langle A_{ij}^2 \rangle - \langle A_{ij} \rangle^2), & \gamma &= \frac{s}{\sigma^2} (\langle A_{ij} A_{ji} \rangle - \langle A_{ij} \rangle^2).\end{aligned}\quad (4)$$

where $s = Sc$ is the average number of interaction partners for each species (c being the connectivity of the interaction network). This factor s means that μ and σ^2 indeed represent the *total* interaction strength and variance faced by one species from all its partners. Given our definition of γ , it takes values in the range $[-1, 1]$: $\gamma = 1$ means that two partners affect each other identically, while $\gamma = -1$ means that reciprocal effects are maximally different, see Sec. 4.2.2. Examples from competitive and trophic communities are given in the next section.

These parameters suggest how disordered systems differ from both mean-field and structured systems. The various species traits retain a non-negligible heterogeneity, encapsulated by ζ and σ . Yet, little more information is needed than those simple metrics, which give a basic estimate of “how different” species can be. The precise interpretation of these abstract parameters will become clearer, first from seeing what role they play in model results, and later in more in-depth discussion.

All higher-order correlations, that capture more precise structural information about the pool, are neglected. If more order is needed (e.g. for trophic systems, given that trophic roles are an essential feature that is not represented here), one could consistently push the formalism further, incorporating more and more structural properties in the form of correlations⁵.

2.2.3 The special case of competitive communities

Note that in competitive communities, where species have carrying capacities $K_i > 0$ which can follow a wide distribution (e.g. lognormal), it is often difficult to generate communities where many species coexist if interactions A_{ij} are drawn arbitrarily. It has been found that a particular scaling of interactions does allow for species coexistence: they must scale with the ratios of carrying capacities, more precisely

$$A_{ij} = \frac{R_i K_i}{K_j} B_{ij} \quad (5)$$

⁵The reader may recognize this idea as similar to that of moment equations, see e.g. [8].

where B_{ij} has a narrower distribution (e.g. normal) [17, 26]. In this case, we can divide the equilibrium condition (3) by R_i to get

$$0 = 1 - \eta_i - \sum_{j \neq i}^S B_{ij} \eta_j, \quad \eta_i = \frac{N_i}{K_i}. \quad (6)$$

In other words, all the calculations and results presented here apply to this particular competitive setting by replacing the N_i with the effective variables η_i , the real interactions A_{ij} by their “core” B_{ij} , and setting $R_i = D_i = 1$.

Our framework in fact gives a very clear mathematical explanation for how the scaling s(5) allows for coexistence: imagine instead that we had started from the Lotka-Volterra equations rewritten as

$$\frac{d}{dt} N_i = R_i N_i \left(1 - \frac{N_i - \sum_{j \neq i}^S \tilde{A}_{ij} N_j}{K_i} \right)$$

with some randomly drawn $\tilde{A}_{ij}$, then the equilibrium condition would be

$$0 = 1 - \eta_i - \sum_{j \neq i}^S \tilde{A}_{ij} K_j \eta_j \quad (7)$$

for the effective variables $\eta_i = N_i/K_i$. The effective interactions $B_{ij} = \tilde{A}_{ij} K_j$ could be very widely distributed, since K_j themselves are drawn from some wide (e.g. lognormal) distribution. This would correspond to a large value of interaction heterogeneity σ , which we show below has a very negative effect on coexistence. We give precise conditions on the statistics of B_{ij} which allow to generate communities of given diversity.

3 What can we predict?

While the following sections discuss the methods that we will use to solve the equations above, this one gives an overview of some of the main system properties that can be computed in this framework. These results fall roughly in four categories of important ecological properties: functioning, diversity, stability, and features selected by the assembly process.

We present these results in the abstract parameter space defined in (4). The interpretation of the parameters comes in Sec. 4. Importantly, the results that we show are exact for the region of parameters where there exists a single, globally stable, uninvadable final assembled state. However, most results hold approximately past the boundary of this region, and we later discuss in Sec. 4.1 when systems are likely to be found in this region. The complex region that lays beyond is still being investigated.

3.1 Functioning

For functioning, we can for instance compute the total biomass

$$N_{tot} = \sum_i N_i \quad (8)$$

and total productivity or growth from external resources

$$G_{tot} = \sum_i N_i R_i \Theta(R_i) \quad (9)$$

with the Heaviside function $\Theta(x) = 1$ if $x > 0$, $\Theta(x) = 0$ otherwise. In fact, since G_{tot} scales with the total biomass, it is more practical to look at the productivity ratio G_{tot}/N_{tot} , which increases when faster-growing species are favored.

Note that throughout, we make no difference between biomass and abundance, as we are not considering individual mass as a particular species trait. Our choice is equivalent to assuming that biomass has no systematic correlation with either interactions or carrying capacities in the pool of potential invaders, so that converting total abundance to total biomass only requires multiplying by the average individual's mass in the community. However the framework could easily be extended to account for individual mass as a trait controlling other dynamical properties (e.g. via metabolic scaling).

In addition, let us recall here that, as suggested in Sec. 2.2.3, the calculations made here can be extended to competitive communities with wide distributions of carrying capacity K_i , but all our predictions for N_i will instead be about the rescaled abundances N_i/K_i . Hence, for instance, predictions of $\langle N \rangle > 1$ will signal overyielding. Since these discussion points are specific to competitive communities and we strive to be as general as possible, we will leave the analysis of all these consequences to future work.

3.2 Diversity and coexistence

For diversity we will consider two main metrics. The first is the number of species surviving in the assembled state $S^* = S\phi$, see Fig. 2. As one would expect, coexistence decreases with increasing σ and ζ , since more variance in growth rates and interactions entails that some species are more likely to be overall stronger competitors than some others. Interestingly, as we discuss below, decreasing γ tends to increase coexistence. The second is Simpson's index, which gives a measure of biomass concentration,

$$I = \sum_i \left(\frac{N_i}{N_{tot}} \right)^2 \quad (10)$$

Its inverse, Simpson diversity D , tells us how effectively diverse the community is, in terms of the biomass being equitably distributed among all species. The index is related to the second moment of the abundance distribution. In fact, the full distribution can be computed within this framework, with some caveats discussed in Sec. 5.

3.3 Stability

We discuss three aspects of stability: the existence of multiple equilibria, variability around an equilibrium under stochastic perturbations, and the consequences of permanent community shifts – be they due to environmental changes or species extinction. All three can be shown to be related via simple expressions to pool parameters and to the fraction of surviving species ϕ .

First, we can find an expression for *invariability* as defined in [2], i.e. the inverse of the response of the system to demographic noise.

$$\mathcal{I} = 1 - \frac{\mu}{s} - \frac{\sigma^2}{2} (\gamma + 1) \frac{\phi}{1 - \mu/s} \quad (11)$$

It can formally be related to the response to other types of noise, and is also the basis for deriving a prediction for Taylor’s law.

Second, we can derive the number of secondary extinctions resulting from the extinction of one species in the system. In a disordered system, secondary extinctions are more or less numerous depending on parameters, but never make up a significant fraction of the community, due to species abundances being controlled by the community as a whole rather than by individual partners.

The final important result is χ the order parameter for multistability. Let us imagine that we perturb all the carrying capacities

$$K_i \rightarrow K_i + \xi_i \quad (12)$$

with a perturbation that is on average $\langle \xi_i \rangle = 0$, representing a small change in the environment that leads to different species being favored, without making the environment overall better or worse for the community. Then we can ask how much the abundances of the species are changed in response,

$$N_i \rightarrow N_i + x_i \quad (13)$$

and this change (with zero mean) is characterized by its variance, which we compute in Sec. 5.6 below:

$$\chi^2 = \left\langle \frac{x_i^2}{\xi_i^2} \right\rangle = \frac{1}{\frac{1}{4} \left(\sqrt{u^2 - 4\gamma\sigma^2\phi} + u \right)^2 - \sigma^2\phi}. \quad (14)$$

where $u = 1 - \mu/s$. This expression reflects the deep relationship between coexistence ϕ and sensitivity to the environment χ .

The divergence of χ^2 when the denominator of (14) vanishes, i.e. when

$$\sigma^2 \rightarrow \frac{1}{\phi} \left(\frac{1 - \mu/s}{1 + \gamma} \right)^2 \quad (15)$$

signals the boundary of the region where multiple equilibria exist: precisely at the boundary, an infinitesimal change in intrinsic fitness can lead to different species being present in the final assembled state. Of course, while the quantity

Figure 3: Average effect of a random perturbation of carrying capacities, calculated as χ (solid line) or directly from the assembled community in simulations (dots), depending on coupling heterogeneity σ . Left: $\gamma = 1$, with divergence of χ at finite σ , revealing the loss of stability of the single global attractor and the start of the multistable phase. Right: $\gamma = -1$, no divergence (no multistable phase).

formally diverges, the community’s response is not truly infinitely steep – it only signals that the approximations we use to compute the single globally attractive equilibrium break down. Indeed, as one gets close to the divergence, χ stops predicting the simulated response to a press perturbation, which plateaus as shown in Fig. 3.

Hence, it is possible to construct the *phase diagram* shown in Fig. 4, which indicates where in parameter space one will find a single, multiple or no equilibria. As we discuss more below, lowering γ has a stabilizing effect, as it extends the single equilibrium phase – in fact, it is apparent from (15) above that, in the limit $\gamma = -1$, the threshold to the multiple equilibria phase becomes infinite, and there can only be one attractor which is globally stable.

3.4 Selected features

Previous work [6] has shown that it is possible to use this framework to infer not only the distributions of abundances, growth rates and interactions in the assembled state, but also correlations between those, including the frequency of various small network motifs [15]. How these distributions differ from the pool baseline tells us how non-random properties have been selected from a random pool. This gives us worthwhile information on which factors are likely to enable a species to survive into the assembled state, or conversely, how ecosystems may adapt to ensure coexistence. We demonstrate this idea in this article with the selection of carrying capacities or growth rates, but see the Appendices in [6] for details on the selection of interaction patterns.

Figure 4: Phase diagram in the (μ, σ) plane for various γ and $\zeta = 0$, based on phase parameter χ^2 from Sec. 3.3. There are three regions depending on parameters. In the first region, there is a single global attractor, meaning that – after a long enough sequence of invasions and extinctions – a given community will reach a single uninvadable state, no matter what sequence it went through. In the second region however (up of the dashed line, signalling a divergence of χ^2), there are multiple such attractors, and the final assembled state depends on the history of the system. As $\gamma \rightarrow -1$, the transition to the second region disappears. Finally, the third region (left of the solid line, signaling a Rdivergence of $\langle N \rangle$) has no attractor, as some species have mutualistic interactions that are stronger than any negative feedback and their abundances grow to infinity. This signals a breakdown of the Lotka-Volterra approximation for more realistic dynamics.

3.5 Discussion

Most striking is the role of γ the reciprocity of interactions: as γ becomes negative, sensitivity to perturbations decreases, the single attractor region expands and coexistence increases. This suggests that asymmetrical interactions – including commensalism, amensalism and trophic interactions – lead to easier assembly than symmetrical ones, echoing the conclusions of other authors [1] who noted “remarkable differences between predator-prey interactions, which increase stability, and mutualistic and competitive, which are destabilizing”. Yet, we can also show that $\gamma < 0$ may very well be achieved by competitive interactions. It is in fact possible to construct trophic and competitive communities so that the latter have more negative γ than the former. Hence, this statement should not be taken as generally valid, and the meaning of γ should be discussed much more carefully, as we do in Sec. 4.2.2.

In the single attractor region, many different ecological scenarios become indistinguishable as long as one is only interested in the eventual assembled state. Whether invaders from the pool come from migrations or mutations, frequent or rare, correlated or not with the species dominating the community at any given time, interspersed with extinctions and recoveries or not – these distinctions become irrelevant in the long run, as any dynamical trajectory that does not involve irreversible changes in the pool of potential invaders (e.g. a permanent extinction) will end at the same point.

If we look at a pool that lies on the boundary between the two regions, multiple attractors first emerge as small deviations from the single attractor that existed below that boundary - most species remain, and only a few rare ones distinguish the multiple attractors by their absence or presence. Therefore, even though the results presented in this work focus on the single attractor region, where our analytical results hold exactly, we will demonstrate multiple times that our framework still provides good approximations some distance into the multiple attractor region. When our framework ceases to be predictive altogether, it is likely that the various attractors are different enough that perturbations may engender significant regime shifts.

If $\zeta \neq 0$, the model predictions seen on Fig. 5 are still quite similar to those shown in Fig. 2 for $\zeta = 0$, except for the fact that average coupling μ plays a much more important role: instead of controlling only the average biomass $\langle N \rangle$, it affects all the metrics, favoring the existence of a global attractor while reducing coexistence.

4 Mathematical setting

This section is meant to be an qualitative explanation of the mathematical concepts that are crucial for a better understanding of the methods and results shown in this paper. Sec. 5 below gives a detailed explanation of the derivation, of interest to the more mathematically-inclined reader.

We recall that the central equation we investigate is (3) which defines the

Figure 5: Coexistence, abundance and stability properties in the space (μ, σ) for $\gamma = 1, 0, -1$ (top to bottom) and $\zeta = 0.3$. See Fig. 2 for results at $\zeta = 0$.

equilibrium for Lotka-Volterra dynamics:

$$0 = R_i - D_i N_i - \sum_{j \neq i}^S A_{ij} N_j.$$

Before solving this equation, we show that it is possible to better isolate the role of the various parameters, thanks to two simple tricks of rescaling. First, it is easily seen in (1) that R_i , D_i and A_{ij} could all be multiplied by the same number, without changing the dynamics (only making them globally faster or slower). Hence we can arbitrarily choose the time scale so that, for instance, growth rates have unit mean:

$$\langle R_i \rangle = 1. \quad (16)$$

Now, we can look more specifically at equilibria, i.e. equation (3). Clearly, we could multiply all D_i and A_{ij} by some number, and divide all N_i and N_j by the same number, without changing the equilibrium: the average intensity of interactions, inter- and intra-species, simply changes the average abundance at equilibrium. Hence we can also arbitrarily take

$$\langle D_i \rangle = 1 \quad (17)$$

so that now, interactions A_{ij} are measured in units of the average self-interaction, and abundances are measured in units of “carrying capacity” – indeed, if it were alone with positive growth rate, the average equilibrium abundance of the species would be around 1 (or equal to 1 if $D_i = 1$). Neither of these choices subtract anything from the generality of the equations. An important consequence is that any change of the absolute value of growth rates or interactions, for instance via nutrient enrichment, will affect only the total amount of biomass and no other property of the assembled state.

The second trick is using the fact that interactions can always be rewritten as

$$A_{ij} = A + \sigma_A a_{ij} \quad (18)$$

with two parameters: the average interaction $A = \langle A_{ij} \rangle$ and the variance of interactions $\sigma_A^2 = \langle A_{ij}^2 \rangle - A^2$, so that the rescaled interactions a_{ij} have zero mean and unit variance

$$\langle a_{ij} \rangle = 0, \quad \langle a_{ij}^2 \rangle = 1 \quad (19)$$

For now, we make no further assumption on what the distribution of the A_{ij} and a_{ij} may be, beyond the fact that it has a well-defined mean and variance. We simply suggest that is important to know both how intense interactions are on average, i.e. A , and how much they differ between pairs of species, σ_A .

Hence, we transform (3) into the following equation

$$0 = R_i - D_i N_i - A \sum_j N_j - \sigma_A \sum_j a_{ij} N_j \quad (20)$$

and using

$$\sum_{j \neq i}^S N_j = S \langle N \rangle - N_i. \quad (21)$$

we finally get

$$0 = (R_i - SA \langle N \rangle) - N_i(D_i - A) - \sigma_A \sum_j a_{ij} N_j \quad (22)$$

The first term is an effective growth rate including the overall biomass loss due to competition ($A > 0$) or gain due to mutualism ($A < 0$). The second term contains the effective intra-species competition, and the third the relative gain or loss of fitness of species i due to the *heterogeneity* of inter-species interactions. Multiple insights already arise from this simple reformulation, without even solving the equation, and we discuss them now.

4.1 Scaling of interactions: when are large disordered communities not neutral?

4.1.1 Inter-species versus intra-species effects

It is well-known that intra-species competition plays a large role in stabilizing the dynamics: interactions with other species become important only if they are not negligible compared to this self-interaction. Strongly interacting communities can exhibit complex dynamics with multiple equilibria, limit cycles or other attractors. If interactions with others are negligible, each species simply follows its own logistic equation, and the system is sure to evolve toward a globally attractive equilibrium.

The equation (22) above shows that – beyond a simple translation of every species’ growth rate – the effective importance of interactions compared to self-interaction is controlled by the ratio

$$\frac{\sigma_A}{D_i - A} \quad (23)$$

This gives a clear intuition of why species heterogeneity, as encapsulated by σ_A , can be important in the dynamics. If species are too similar, $\sigma_A \rightarrow 0$, then the abundances N_i decouple from each other: species are only coupled collectively through the average value $\langle N \rangle$ in the effective growth rate, except in the very special case of the Hubbell point $D_i = A$ where all inter- and intra-specific interactions are equal [16, 18]. Away from that singular point, we see that what truly destabilizes the system is not only how *intense* interactions are, i.e. A at the denominator, but also how *diverse* they are, i.e. σ_A at the numerator.

4.1.2 Comparing communities of different sizes and connectivities

A second central insight comes from asking: what does it mean to say that two pools with different sizes S have the *same amount* of heterogeneity? Let

us assume that each species is connected to s_i other species in the interaction matrix, on average

$$\langle s_i \rangle = s = cS \quad (24)$$

with $c \in [0, 1]$ the connectivity of the network. If s is the same in the two pools⁶, then species from either pool have the same average number of interactions. In that case, it makes sense to directly compare the properties of these interactions, i.e. the values of A and σ_A , to decide which pool has stronger or more diverse interactions.

However, let us assume that c is the same for both pools, e.g. $c = 1$. In that case, the dynamical influence on one species from all the others will be proportional to the quantity

$$\mathcal{S}_i = \sum_{j \neq i}^s A_{ij} = sA + \sigma_A \sum_{j \neq i}^s a_{ij} \quad (25)$$

(this gives a general trend; of course, the precise dynamical effect will depend on species abundances). This quantity is on average

$$\langle \mathcal{S}_i \rangle = sA = ScA \quad (26)$$

which increases with the size of the community. This term is seen in (22) to intervene in the effective growth rate. As for how much $\mathcal{S}_i$ varies between species, that comes from the sum

$$\sum_{j \neq i}^s a_{ij} \sim \pm \sqrt{s} \quad (27)$$

By assumption, this sum is an unbiased random walk of mean 0 and typical step size 1; hence, after s steps (with s large enough), no matter the distribution of a_{ij} , the random walk will typically have gone a distance of order $\sqrt{s}$ toward either positive or negative values. There are two important assumptions here: one of large numbers, which can only be checked numerically (in fact, we can often find good agreement even for $s \sim 10$), and one of randomness which is discussed more carefully in Sec. 5. If one accepts these assumptions for now, the final interaction strength perceived by any species falls roughly in the range

$$\mathcal{S}_i \sim sA \pm \sqrt{s} \sigma_A \quad (28)$$

This is a crucial point: if we consider different systems with the same values of A and σ_A , but different values of S or c , then we can see from (28) that, the larger and more connected the pool, the smaller the relative variation in $\mathcal{S}_i$ compared to its mean. In fact, for very large communities, there will be practically no difference in the interaction strength perceived by different species, making

⁶Note that keeping the same s for different pool sizes S means that connectivity $c = s/S$ must decrease with pool size.

them *mean-field* (all species are effectively identical from the point of view of interactions) rather than disordered.

Instead, two pools of different size S have “equal” interaction strength and heterogeneity if their pools exhibit the same values of

$$\mu = sA = s\langle A_{ij} \rangle, \quad \sigma^2 = s\sigma_A^2 = s(\langle A_{ij}^2 \rangle - A^2) \quad (29)$$

rather than the same values of A and σ_A . To understand this, we can see that, no matter the pool size or connectivity, a species’ total interaction with others is given by

$$\mathcal{S}_i \sim \mu \pm \sigma. \quad (30)$$

We can indeed show that μ and σ (in addition to a few other parameters) are what controls the properties of the assembled state – systems with different S but the same μ and σ will obey the same trends in terms of diversity, stability, abundance distribution, and so on. This rescaling means that systems with more connections between species (i.e. larger S or larger c) must have proportionally more varied interactions to retain the same σ , so that heterogeneity still matters and is not “averaged out”.

4.1.3 Tradeoffs and global stability in large systems

We now further investigate this inverse perspective: instead of asking how *community parameters* μ and σ change with diversity S if we keep the same distribution of *individual parameters* A_{ij} , we ask how these A_{ij} should change with S to achieve a given μ and σ . This simple question actually contains an important insight on the properties of large assembled communities. It is easy to see that if one wants to fix connectivity c as well as μ and σ , but increase S , then necessarily interactions should vary like

$$\langle A_{ij} \rangle = \frac{\mu}{s} \propto \frac{1}{S}, \quad \sigma_A = \frac{\sigma}{\sqrt{s}} \propto \frac{1}{\sqrt{S}}. \quad (31)$$

For large enough systems $S \gg 1$, if we maintain the same values of μ and σ , we thus expect $\sigma_A \gg \langle A_{ij} \rangle$. Having a standard deviation much larger than the mean is not achievable with “reasonable” distributions⁷ unless we allow the interactions to take any sign (the mean can thus go to zero while the variance remains large).

Hence, due to simple mathematical constraints, large systems must have at least one of these four properties: large μ , small σ , low connectivity c , or mixed interactions allowing for all signs. For instance, a large competitive community where all species can potentially interact ($c = 1$) must have either large μ , meaning very strong competition and low biomass, or small σ , meaning that

⁷It can be achieved with a power-law tail whose exponent is around or below 2, meaning that the variance of the theoretical distribution is arbitrarily large or even diverges, but in effect, a few massive outliers will be responsible for this variance, while most interactions will be negligible.

heterogeneities in species interactions cease to matter and the system is effectively mean-field. This “biomass–heterogeneity” tradeoff in large competitive communities, which we emphasize comes from simple algebra independent of any model, is indeed very apparent in our analysis of resource consumer models in [4].

This tradeoff has important consequences: as shown in recent work [40], small σ is a condition under which a complex ecosystem might be reduced to a single dimension, leading to predictions of “universal” behavior [9]. Furthermore, small σ and small ζ taken together constitute the basic hypothesis of neutral theory – that we can take all species as identical in large communities, despite their very real heterogeneities. Hence, the fact that simple algebraic constraints could limit σ in large competitive communities provides some formal grounding for the intuitions behind “universal” and neutral dynamics.

4.2 Couplings rather than interactions

4.2.1 Allowing heterogeneous self-interactions

At multiple points in the discussion above, we have suggested that interactions are not, by themselves, the quantities that control the equilibria: instead, that role is taken by *couplings*, i.e. the ratios of interactions to self-interaction, which represents how much the influence of another species could cause one to deviate from its independent logistic trajectory.

In Sec. 2, we presented results assuming that all self-interactions were equal, $D_i = 1$. But in fact, we can easily see that we can transfer our conclusions to the Lotka-Volterra equilibrium equation (3) *divided by* D_i ⁸:

$$\begin{aligned} 0 &= \frac{R_i}{D_i} - N_i - \sum_{j \neq i}^S \frac{A_{ij}}{D_i} N_j \\ &= K_i - N_i - \sum_{j \neq i}^S \alpha_{ij} N_j. \end{aligned} \quad (32)$$

Here, we recall $K_i = R_i/D_i$ the carrying capacity of species i , and $\alpha_{ij} = A_{ij}/D_i$ the coupling of species i to species j . Thus, all the results we presented in Sec. 2 still hold if, instead of using the variance of growth rates, we use that of carrying capacities, and instead of using the statistics of A_{ij} , we use those of α_{ij} . In other words, we formally redefine the parameters of (4) as follows:

$$\begin{aligned} \zeta^2 &= \langle K_i^2 \rangle - \langle K_i \rangle^2, & \mu &= s \langle \alpha_{ij} \rangle, \\ \sigma^2 &= s(\langle \alpha_{ij}^2 \rangle - \langle \alpha_{ij} \rangle^2), & \gamma &= \frac{s}{\sigma^2} \left(\langle \alpha_{ij} \alpha_{ji} \rangle - \langle \alpha_{ij} \rangle^2 \right). \end{aligned} \quad (33)$$

⁸In fact, this is a useful parametrization, since the Lotka-Volterra approximation to more realistic dynamics may not hold if couplings are too large $|\alpha_{ij}| > 1$ (e.g. mutualistic interactions would then cause abundances to diverge). Hence couplings also give an idea of the range of validity of the model.

We can relabel the parameters and retain the same mathematical structure, and their interpretation does not change too dramatically – except in the important case of γ , as we soon discuss. Previously in this section, we used $A = \langle A_{ij} \rangle$ and σ_A extensively to discuss simple and general properties of the dynamics – the distribution of the interactions A_{ij} is otherwise entirely unspecified, and could take arbitrary forms (Gaussian, exponential, uniform, gamma, multimodal...) except for the pathological. We saw in examples that almost all other properties of this distribution $P(A_{ij})$ were irrelevant to assembly for large enough communities, one of the very counter-intuitive and attractive traits of disordered systems. In fact, we see now that, if self-interactions differ, we should use the α_{ij} instead; but thankfully, results remain insensitive to the details of $P(\alpha_{ij})$. The subtlety is found in the fact that σ now includes not only the heterogeneity of interactions, but also that of self-interaction – in other words, species with identical interactions but very different growth rates and carrying capacities will have large σ , with all the consequences shown in Sec. 2 in terms of coexistence and multistability.

This of course requires that $D_i \neq 0$ – and in fact that $A < D_i$, as the case of very strong interactions $A > D_i$ tends to lead to very unstable dynamics that we wish to avoid entirely in this approach⁹, as discussed in 1.2.1.

4.2.2 Reinterpreting γ

However, there is one crucial change in interpretation. The parameter $\gamma \in [-1, 1]$ represents the average degree of reciprocity between the dynamical coupling of species i to j , and that of j to i . The extreme case $\gamma = 1$ indicates fully symmetrical couplings $\alpha_{ij} = \alpha_{ji}$, while $\gamma = -1$ means anticorrelated couplings. Unfortunately, while this parameter is the one that truly encapsulates how asymmetry between interaction partners affects the assembly process, its interpretation is very subtle and counter-intuitive.

Let us first discuss γ in the previous case where all self-interactions are equal, $D_i = 1$, as even that simpler situation has some difficulties. If $A = 0$, then $\gamma < 0$ suggests trophic interactions, since A_{ij} and A_{ji} will typically be of opposite signs. However, if $A \gg \sigma_A$, then even with $\gamma = -1$, $A_{ij} \sim A + \sigma_A$ and $A_{ji} \sim A - \sigma_A$ will have the same sign. They will simply be maximally anticorrelated: for every pair of species, one will affect the other much more than it is affected by it. This can for instance give rise to “rock-paper-scissors” intransitive competition.

If we now return to varied self-interactions $D_i \neq 1$ so that couplings are different from bare interactions, even that refined notion of reciprocity becomes more complex. Consider for example the classic resource competition models [21, 10, 39]. By construction, interactions come from two species harvesting the same resource and are therefore perfectly symmetrical, $A_{ij} = A_{ji}$. However, couplings are *not* – in fact, depending on parameters, the model can give any γ

⁹A classic case is Levins and Culvers’s colonization-competition tradeoff [19] where only a very special and asymmetrical structure ensures the survival of more than a single species. We discuss it further in [4].

between -1 and 1 , and most plausible parameter values correspond to $\gamma < 0$. In any pair of competitors, each reduces the other's growth equally, yet this will on average have a much larger impact on the dynamical trajectory of one species than the other.

Hence, the proper way to understand γ is in relative terms: no matter whether the interaction is positive or negative or even symmetrical, if $\gamma = -1$, one species will respond much more to the interaction (in the sense that it will be more affected compared to the trajectory it would have if alone), while for $\gamma = 1$ they will respond equally. Given that assembly is a process of selecting species that are most compatible among a larger pool, it can be understood that what really matters is not the net outcome of each interaction, but how differently it affects the partners' dynamics. But despite its predictive power, any intuitive interpretation of this parameter can only be made in the context of a more concrete model.

4.3 Computing the assembled state

At the start of Sec. 4, we showed that the equilibrium of the Lotka-Volterra equations (3), which we recall is given by

$$0 = R_i - D_i N_i - \sum_{j \neq i}^S A_{ij} N_j,$$

can be rewritten in all generality as

$$0 = (R_i - S A \langle N \rangle) - N_i (D_i - A) - \sigma_A \sum_j a_{ij} N_j. \quad (34)$$

given the mean $A = \langle A_{ij} \rangle$ and the variance of interactions $\sigma_A^2 = \langle A_{ij}^2 \rangle - A^2$, with the rescaled interactions $a_{ij} = (A_{ij} - A)/\sigma_A$.

As we argued in Sec. 4.2.1, the quantities that truly control the dynamics are not “bare” interactions A_{ij} , but couplings

$$\alpha_{ij} = \frac{A_{ij}}{D_i}$$

which represent how much species j can cause species i to deviate from its baseline (logistic) trajectory controlled only by self-interaction D_i . We now rewrite

$$\alpha_{ij} = \frac{\mu}{s} + \frac{\sigma}{\sqrt{s}} a_{ij} \quad (35)$$

with

$$\langle a_{ij} \rangle = 0, \quad \langle a_{ij}^2 \rangle = 1. \quad (36)$$

Then, we have

$$0 = K_i - N_i - \sum_{j \neq i}^S \alpha_{ij} N_j.$$

and instead of (34) we get

$$0 = k_i - uN_i - \frac{\sigma}{\sqrt{s}} \sum_j a_{ij} N_j \quad (37)$$

where we have defined the effective self-interaction

$$u = 1 - \frac{\mu}{s} \quad (38)$$

and the effective carrying capacities

$$k_i = K_i - \mu \langle N \rangle \quad (39)$$

Following previous work [6], we use the cavity method from physics [29] to find the properties of the equilibria of equation (37). Technical details are given in Sec. 5. To understand the results and their limitations, it is enough to grasp the following simplified discussion, introducing some important concepts.

Let us take an assembled state created from a pool with $S-1$ species. When the final species i invades, if S is large, the properties of the assembled state should not change significantly¹⁰. Hence, we imagine that we can move from $S-1$ to S species holding everything constant, except for the fact that the other species have their final abundance N_j shifted a small amount, in the direction suggested by their couplings with i ,

$$N_j \rightarrow N_j - v a_{ji} N_i \quad (40)$$

with v a positive “response coefficient” that we can compute. Inserting this shifted abundance in (37), a feedback is created on species N_i . This is where the parameter γ enters the picture: if $\gamma > 0$ the feedback is positive (e.g. a species reduces its competitors’ abundance, which in turn allows it to grow more), while if $\gamma < 0$ the feedback is negative (e.g. a predator reduces its preys’ abundance, which stifles its growth). This explains how negative γ can play a stabilizing role in the dynamics, as we have seen above starting in Sec. 3.3.

The finite trick is to see that this equation expresses N_i as a random variable obeying:

$$N_i = \frac{k_i - \sigma a n}{1 - \frac{\mu}{S} - \sigma^2 \gamma v} \quad (41)$$

where a is a Gaussian random variable of mean 0, derived from the random walk (27) on the a_{ij} ’s, and n is a random variable drawn from the distribution $P(N)$ of surviving species. This simplification can be made technically precise, and relies on a crucial property of disordered systems: there is no correlation between the a_{ij} , which represent direct couplings, and the N_j before

¹⁰This assumption is correct if there is only one globally stable assembled state; the method should later be extended to settings with alternative stable states, where the addition of a new species could cause a sudden shift between different equilibria.

the shift (40), which are shaped by all the other indirect couplings in the system. These two could not be taken as independent parameters if each species depended strongly on a few others, but the disordered nature of interactions means that species abundances are shaped mostly by collective dynamics involving the whole community. A failure of this approach to correctly predict the behavior of some model could signal that local ordered dynamics prevail over collective disordered dynamics.

The second central assumption that is made here is the existence of a single global attractor, meaning that no matter which initial condition is taken, the same assembled state is reached provided all species are given a chance to invade when they can. This premise is hidden in the fact that we assume that equilibrium species abundances do not shift drastically with the addition of a new species. As we have mentioned in Sec. 3.3, we can compute a quantity that becomes infinite precisely at parameter values for which there ceases to be a single attractor, thus signalling the appearance of potential tipping points. Methods in physics have been derived for the express purpose of investigating such multistable regimes, and quantifying the number and stability of attractors. In the case of Lotka-Volterra dynamics, however, this task is ongoing.

The assembled state must be feasible and uninvadable, and therefore, studying it reveals what controls species coexistence. But it also contains information on the ecosystem’s sensitivity to external change: the quantities that arise naturally when computing the assembled state are long-term responses to a press perturbation, i.e. how the whole community’s equilibrium shifts in response to some species being pushed toward an abnormal abundance. This can come from durable environmental changes that affect the fitness of some species, but also from the introduction of a new invader into the system.

5 Detailed calculations

Readers who are only interested in the general spirit of the method can stop here, as we turn to the practicalities of using it and the dynamical intuitions it brings. In this section, we follow in closer detail the derivation of the equilibrium solution for a large Lotka-Volterra system first presented in [7]. For the mathematically minded reader, let us note that related calculations, using more technical methods from disordered systems, had previously been achieved for replicator equations (which are closely related to Lotka-Volterra) and mostly discussed within physics [32, 41]. Earlier results relied on the “replica method”, limited to symmetrical interactions (hinging on the fact that there then exists a Lyapunov function, as noted long ago by MacArthur for resource competition [21]). Later results employed a “dynamical generating functional method” which allows to find not only equilibria but a single equation for the effective dynamics of the entire system [32], although that equation is partially implicit and can only be solved by ansatz.

Bunin [7] proposed the first derivation for the equilibria of Lotka-Volterra systems, using a technically simple implementation of the “cavity method” [27].

This calculation is comparatively straightforward while retaining many important intuitions of disordered systems, and its application to ecological models is immediate, which makes it especially valuable for introductory purposes.

5.1 The cavity trick

The assembly process shapes the community by changing the abundances of species until an equilibrium is reached. From the imposed properties of the pool, assembly uses the abundances as its tool to select and reweigh species features, crafting a state that is feasible, stable and uninvadable by any other species in the pool. Hence, all the properties of the solution can be derived from the distribution of abundances and their correlations with species features (such as carrying capacities and interactions).

The logic of the cavity method is remarkably simple: 1) we compute what happens after a single invasion in an already large community, and 2) we then require that the equilibrium features before and after the invasion be unchanged. The second part means that we are computing the fixed point of the invasion process, i.e. an equilibrium which cannot be invaded nor destabilized, rather than a single-step invasion as is often done in population genetics (testing whether a given mutant species can invade).

As shown in Fig. 6, the calculation of the invasion is made simpler by assuming that, even though the invader itself retains its nonlinear dynamics, the effect of the invader on each already present species is small enough to be approximated by a *linear* change of their abundance. This assumption may seem intuitive if the final abundance of each species is the result of many interactions, so adding one more should not cause a dramatic change. A more formal justification is given in Appendix A.

Let us introduce species 0 into a community of $S - 1$ species. Its equilibrium abundance N_0 changes other N_i a little, as they now obey:

$$0 = k_i - uN_i - \sigma_A \sum_j a_{ij}N_j - \sigma_A a_{i0}N_0 \quad (42)$$

(where we recall that $\sigma_A = \sigma/\sqrt{s}$). Formally, this looks like changing their carrying capacities k_i a little

$$0 = \underbrace{(k_i - \sigma_A a_{i0}N_0)}_{\tilde{k}_i} - uN_i - \sigma_A \sum_j a_{ij}N_j \quad (43)$$

where $\tilde{k}_i$ is a slightly displaced carrying capacity. We now assume that N_i responds linearly to this small perturbation: using a Taylor expansion around the previous equilibrium abundance N_i^* , calculated without species 0, we can write

$$N_i \approx N_i^* + \frac{dN_i}{dk_i}(\tilde{k}_i - k_i) = N_i^* - v_{ii}\sigma_A a_{i0}N_0 \quad (44)$$

where we define the response coefficient,

$$v_{ii} = \frac{dN_i}{dk_i}. \quad (45)$$

This seems unhelpful at first, since v_{ii} is just another unknown in the problem, but we will show that creating this new set of unknowns actually helps us further down. To interpret v , let us remark that, if species were independent, then

$$N_0 = k_0, \quad v_{00} = \frac{dN_0}{dk_0} = 1 \quad (\text{independent species}). \quad (46)$$

Hence, $v \neq 1$ indicates that species abundances are not defined solely by their intrinsic fitness, but by their effective fitness including interactions – it is possible to change their carrying capacity without affecting their abundance proportionally.

In fact, all other species in the community perceive the effect of adding species 0, so a better approximation of the full linear response would include “cross-terms” i.e.

$$N_i \approx N_i^* - \sigma_A \sum_j v_{ij} a_{j0} N_0. \quad (47)$$

with

$$v_{ij} = \frac{dN_i}{dk_j} \quad (48)$$

the long-term response of species i to a change in the carrying capacity of species j , including all the indirect effects from other species. In fact we show later that these indirect effects tend to cancel out and disappear, but let us include them for now, for the sake of completeness.

Now, this shift in others’ abundance feeds back into the equation for N_0 , which tells us how it depends on N_i^* the abundances before its invasion

$$0 = k_0 - uN_0 - \sigma_A \sum_i a_{0i} N_i \quad (49)$$

$$= k_0 - \left(u - \frac{\sigma^2}{s} \sum_{i,j} a_{0i} v_{ij} a_{j0} \right) N_0 - \sigma_A \sum_i a_{0i} N_i^* \quad (50)$$

The issue is that it contains a double sum over all the unknown v_{ij} , representing the overall feedback of species 0 on itself through the rest of the community. A crucial simplification is that, due to the disorder of interactions a_{ij} , these many unknowns collapse into a single one that can easily be determined later: this large sum tends to an average value, as we explain very roughly here and further detail in Appendix A. Since $\langle a_{0i} \rangle = \langle a_{j0} \rangle = 0$, all terms that have $i \neq j$ are fluctuations around 0, while each term $\langle a_{i0} a_{0i} \rangle = \gamma$ has a significant contribution to the sum. Hence, as we noted before, the main contribution comes from the response of N_i to a change of k_i , rather than from cross-terms:

$$\frac{\sigma^2}{s} \sum_{i,j} a_{0i} v_{ij} a_{j0} \approx \frac{\sigma^2}{s} \sum_{i=1}^{s^*} a_{0i} v_{ii} a_{i0} \quad (51)$$

Figure 6: Cartoon of the cavity method.

where $s^* = s\phi$ is the number of species that interact with species 0 and are still alive in the assembled state (we recall that ϕ is the overall fraction of surviving species). This sum is an average over all the interaction partners of species 0, and it is in fact “self-averaging”. This is physics jargon to say that, given a large number of species to interact with, it tends to become a single well-defined number for all species in the community, with fluctuations that become small as s becomes large,

$$\frac{\sigma^2}{s} \sum_{i,j} a_{0i} v_{ij} a_{j0} \approx \sigma^2 \phi \gamma v + O\left(\frac{1}{\sqrt{s}}\right) \quad (52)$$

where

$$v = \langle v_{ii} \rangle. \quad (53)$$

Finally, we obtain the effective equation for N_0 which will be solved below.

$$(u - \phi \sigma^2 \gamma v) N_0 = k_0 - \frac{\sigma}{\sqrt{s}} \sum_i a_{0i} N_i^* \quad (54)$$

5.2 Universality of cavity results

Isolating N_0 on the left-hand side,

$$N_0 = \frac{k_0 - \frac{\sigma}{\sqrt{s}} \sum_i a_{0i} N_i^*}{\hat{u}} = \frac{k_0}{\hat{u}} - \frac{\Sigma_0}{\hat{u}} \quad (55)$$

where

$$\hat{u} = u - \phi\sigma^2\gamma v \quad (56)$$

and we have defined

$$\Sigma_0 = \frac{\sigma}{\sqrt{s}} \sum_i a_{0i} N_i^* \quad (57)$$

We can see in N_0 two main contributions (stretched by a factor $1/\hat{u}$ that is common to all species in the community): the effective carrying capacity k_0 , and a fitness shift Σ_0 arising only from to the *heterogeneity* of interspecies interactions. The former is model-dependent: we have no reason to constrain the distribution of carrying capacities K_0 which is a property of the species pool (we recall that $k_0 = K_0 - \mu \langle N \rangle$). But the latter becomes universal in the limit of many interactions. This relies on the sum $\sum_i a_{0i} N_i^*$ being effectively a random walk.

Here, we are going to assume that interactions between species are not correlated *in the pool* (we discuss a weaker assumption that is still tractable in Sec. 7.4). This lack of correlation has two consequences. First, note that a_{0i} is uncorrelated with N_i^* : the abundance of species i before species 0 was introduced depends on a_{ji} for $j \neq 0$, but bears no relationship to a_{0i} . More importantly, a_{0i} and a_{0j} are uncorrelated, which means that all terms in the sum above are uncorrelated variables, and if they have finite variance, then the sum must tend to a Gaussian law by the Central Limit Theorem. This Gaussian law has mean 0 and variance $s\phi \langle N^2 \rangle^*$, where the mean is taken here over $P(N^*)$ the equilibrium distribution of values for N_i^* .

Hence, for any model where interactions have finite variance, the final abundance distribution will be composed from the model-specific distribution of carrying capacities, and from a generic Gaussian term emerging from interactions¹¹. And even if $P(N^*)$ had a model-specific shape due to $P(K)$, the interaction term, being Gaussian, is always entirely parametrized by the first two moments of $P(N^*)$, meaning that one never needs a lot of information on $P(K)$ to correctly evaluate the effect of interactions.

This entire reasoning however weakens if $P(K)$ – and therefore $P(N^*)$ – have fat tails; any other structure (e.g. multimodal $P(K)$) will somewhat increase the necessary number of species and interactions to converge toward the Gaussian distribution of σ_0 , but fat tails may prevent this convergence altogether (if variance is formally infinite) or make it dramatically slow. This possibility would call for an extreme value calculation instead: what would matter in Σ_0 would not be its variance anymore, but its extreme values, which obey a different but still universal distribution.

We will now assume that $P(K)$ is “well-behaved” (e.g. Gaussian or uniform) so that it is enough to know its mean and variance to derive $P(N^*)$. We discuss in Sec. 6 the sensitivity of results to choices of $P(K)$.

¹¹Further clarifications on how randomness intervenes in interactions can be found in Appendix A.

5.3 Abundance distribution

Let us recall equation (55) which defines the equilibrium abundance N_0 of a representative species in the community as a sum of two random contributions:

$$N_0 = \frac{k_0}{\hat{u}} - \frac{\Sigma_0}{\hat{u}}.$$

This expression allows for $N_0 < 0$, hence, N_0 is a “virtual abundance” that signals that some species could invade and attain abundance N_0 if $N_0 > 0$, or that the species could not invade durably if $N_0 \leq 0$. (Note that the more traditional approach to assembly is to check whether a species can grow immediately at the start of its invasion, without predicting the rest of its trajectory; here on the contrary we check whether it could remain at infinitely long times, after all interactions with the rest of the community have played out).

This virtual abundance has been computed from the carrying capacities of invaders $P(K_0)$ and the abundances of surviving species $P(N^*)$. This means that, if species 0 itself survives, it becomes representative of survivors. If we can compute the probability distribution $P_0(N_0)$, then $P(N^*)$, must be the same distribution restricted to positive values.

An important conceptual trick in this approach is that, instead of computing species abundances for a given community, i.e. a fixed choice of interactions and carrying capacities, we suppose here that all abundances are fixed (even if we do not know their values yet) *except* for that of species 0, and try to compute the typical distribution of N_0 over the whole range of possibilities for a_{0i} , a_{i0} and K_0 . From it we will compute the “typical” abundance distribution $P(N^*)$, i.e. the one that we will see on average over all possible communities whose parameters are drawn from these probability distributions.

Since we know N_0 as a function of N_i^* in a given system, and $P(N^*)$ as a function of the distribution $P_0(N_0)$ over all possible choices for interactions and carrying capacities, we have a closed system that can be solved through the method of moments. First, we compute the first two moments of $P_0(N_0)$ as a function of those of $P(N^*)$. We can take the mean over $P(a_{0i})$ and $P(K_0)$ of both terms in (55), we find

$$\langle k_0 \rangle = \langle K \rangle - \phi \mu \langle N \rangle^*, \quad \langle \Sigma_0 \rangle = 0 \quad (58)$$

where we used $\langle N \rangle = \phi \langle N \rangle^*$ (since when we introduced $\langle N \rangle$ earlier, we took the mean over all S species, including those that are dead), and the facts that $\langle a_{0j} N_j \rangle = \langle a_{0j} \rangle \langle N \rangle$ and $\langle a_{0j} \rangle = 0$. Hence

$$\langle N_0 \rangle = \frac{\langle k_0 \rangle}{\hat{u}} = \frac{\langle K \rangle - \phi \mu \langle N \rangle^*}{\hat{u}}. \quad (59)$$

For variances, we can similarly find

$$\text{Var}(k_0) = \zeta^2, \quad \text{Var}(\Sigma_0) = \phi \sigma^2 \langle N^2 \rangle^*. \quad (60)$$

Hence, if k_0 and Σ_0 are independent variables (see below otherwise) the variance of the sum is the sum of variances and

$$\sigma_0^2 = \text{Var}(N_0) = \frac{\zeta^2 + \sigma^2 \phi \langle N^2 \rangle^*}{\hat{u}^2} \quad (61)$$

Following our assumption that $P(K)$ is well-behaved, these two moments are enough and we approximate $P_0(N_0)$ by a Gaussian law with the mean and variance just given,

$$P_0(N_0) = \frac{e^{-(N_0 - \langle N \rangle)^2 / \sigma_0^2}}{\sqrt{2\pi} \sigma_0}. \quad (62)$$

From there, we can give a definition to $\langle N \rangle^*$, $\langle N^2 \rangle^*$ and ϕ by restricting $P_0(N_0)$ to its positive values

$$\phi = \int_0^\infty dN_0 P_0(N_0) \quad (63)$$

$$\langle N \rangle^* = \frac{1}{\phi} \int_0^\infty dN_0 P_0(N_0) N_0 \quad (64)$$

$$\langle N^2 \rangle^* = \frac{1}{\phi} \int_0^\infty dN_0 P_0(N_0) N_0^2 \quad (65)$$

$$(66)$$

where the $1/\phi$ factor ensures normalization. On the other hand, v is given by taking the derivative (48) in (55)

$$v = \langle v_{00} \rangle = \left\langle \frac{dN_0}{dk_0} \right\rangle = \frac{1}{u - \phi \sigma^2 \gamma v} \quad (67)$$

which is thus coupled to the rest through ϕ (and by appearing in P_0). We now have a closed set of four equations for four unknowns: v , ϕ , $\langle N \rangle^*$, $\langle N^2 \rangle^*$ which appear on the left-hand side of each equation and in $P_0(N_0)$. These equations can be solved together numerically. It is then immediate to compute quantities such as total biomass

$$N_{tot} = S\phi \langle N \rangle^* \quad (68)$$

and Simpson index

$$I = \sum_i \left(\frac{N_i}{N_{tot}} \right)^2 = \frac{\langle N^2 \rangle^*}{S\phi (\langle N \rangle^*)^2} \quad (69)$$

5.3.1 Correlations between K_i and A_{ij}

Equation (61) assumed k_0 and Σ_0 , and hence K_i and A_{ij} , to be independent variables. If there is a correlation between them, then

$$\sigma_0^2 = \frac{\zeta^2 + \sigma^2 \phi \langle N^2 \rangle^* - 2\phi \langle k_0 a_{0i} \rangle}{(u - \sigma \gamma v)^2} \quad (70)$$

where

$$\langle k_0 a_{0i} \rangle = \langle K_0 A_{0i} \rangle - \frac{\mu}{S}. \quad (71)$$

Hence, such correlations can be accounted for, at the cost of adding a new parameter. Other simple extensions are discussed in Sec. 7.

5.3.2 Relationship to empirics

The fact that the theoretical abundance distribution is (approximately or exactly) a truncated Gaussian contradicts much evidence on empirical abundance distributions being fat-tailed – possibly log-normal or even power-law. No fault lies with the mathematical analysis, as it is in perfect quantitative agreement with many simulation models, see [4]. Hence the issue will occur for any system with Lotka-Volterra dynamics (1) with random interactions and carrying capacities drawn from many usual distributions, and without noise.

There are at least three solutions to this problem. The first is a wide distribution of carrying capacities: let us recall that in Sec. 2.2.3, we noted that many competitive models [17, 26] assume¹² that interactions scale like

$$A_{ij} = \frac{R_i K_i B_{ij}}{K_j} \quad (72)$$

with some narrowly distributed B_{ij} (e.g. uniform or Gaussian). In this case, we noted that we could write

$$0 = 1 - \frac{N_i}{K_i} - \sum_{j \neq i}^S B_{ij} \frac{N_j}{K_j}.$$

Hence, by having μ , σ and γ be the properties of the B_{ij} , the results above predict that the distribution $P(N_i/K_i)$ is a truncated Gaussian. If K_i follows a lognormal distribution for instance, as seen in plant monocultures [34], then N_i at equilibrium will have an even wider distribution.

This explanation however assumes that the wide spread of species abundances is due to some intrinsic species properties, which define K_i in a given environment. Under this explanatory scenario, our analysis is compatible with fat-tailed abundance distributions, but does not explain their emergence. However, there are at least two ways in which the dynamics themselves could produce fatter tails: first, the addition of noise, which is a central ingredient in neutral models that successfully predict abundance distributions. Second, the role of spatial aggregation in the data, as it is well known that aggregating over Gaussian distributions with different variances can create the appearance of more complex distributions and even fat tails - indeed, this is seen in an extension of this work to metacommunities (to be published).

¹²This assumption was found to be essential to allow high coexistence.

5.4 Selection by assembly: the case of carrying capacities

Starting from this section, the core of the method is established and we simply show how various quantities of interest can be calculated using the results above.

We have seen above that equilibrium abundances are made up of two independent contributions: one coming from the effective carrying capacity, the other a shift in abundance due to the heterogeneity of interactions

$$N_0 = \frac{k_0}{\hat{u}} - \frac{\Sigma_0}{\hat{u}}$$

We also recall that

$$k_0 = K_0 - \mu \langle N \rangle$$

If we want to know the equilibrium distribution of carrying capacities K_i^* , we start from the joint distribution of abundance N_0 and effective carrying capacity k_0 , which we will then restrict only to surviving species $N_0 > 0$.

$$P(N_0, k_0) dN_0 dk_0 = P(\Sigma_0, k_0) d\Sigma_0 dk_0 \quad (73)$$

hence

$$P(N_0, k_0) = \frac{d\Sigma_0}{dN_0} P(\Sigma_0, k_0) = \hat{u} P(\Sigma_0, k_0) \quad (74)$$

Now, since Σ_0 and k_0 are independent, $P(\Sigma_0, k_0) = P(\Sigma_0)P(k_0)$, so

$$P(N_0, k_0) = \hat{u} P(\Sigma_0 = k_0 - \hat{u} N_0) P(k_0) \quad (75)$$

We argued above that $P(\Sigma_0)$ is Gaussian with mean 0 and variance σ_Σ^2

$$P(\Sigma_0) = \mathcal{N}(\Sigma_0, 0, \sigma_\Sigma^2) \quad (76)$$

with

$$\sigma_\Sigma^2 = \sigma^2 \phi \langle N^2 \rangle^* \quad (77)$$

Now, we want to isolate $P_0(N_0)$ within $P(N_0, k_0)$, so that the integral over $N_0 > 0$ causes the moments we have already computed, ϕ , $\langle N \rangle$ and $\langle N^2 \rangle$, to appear. We recall

$$P_0(N_0) = \frac{e^{-(\hat{u} N_0 - \langle k_0 \rangle)^2 / (\zeta^2 + \sigma_\Sigma^2)}}{\sqrt{2\pi} \sigma_0}, \quad \sigma_0^2 = \frac{\zeta^2 + \sigma^2 \phi \langle N^2 \rangle^*}{\hat{u}^2} = \frac{\zeta^2 + \sigma_\Sigma^2}{\hat{u}^2}$$

With some manipulations (assuming a Gaussian distribution for K_i), we can see that

$$P(N_0, k_0) = P_0(N_0) \mathcal{N}\left(k_0, \frac{\zeta^2 \hat{u} N_0 + \sigma_\Sigma^2 \langle k_0 \rangle}{\zeta^2 + \sigma_\Sigma^2}, \frac{\zeta^2 \sigma_\Sigma^2}{\zeta^2 + \sigma_\Sigma^2}\right). \quad (78)$$

In particular, we find the averages over surviving species

$$\begin{aligned}\langle K_i \rangle^* &= \int_0^\infty dN_0 P_0(N_0) \int_{-\infty}^\infty dk_0 \frac{P(N_0, k_0)}{P_0(N_0)} (k_0 + \mu \langle N \rangle) \\ &= \langle K \rangle + \frac{(\hat{u} + \mu) \langle N \rangle + \langle K \rangle}{1 + \sigma^2 \langle N^2 \rangle / \zeta^2}\end{aligned}\quad (79)$$

and

$$\begin{aligned}\langle K_i N_i \rangle^* &= \int_0^\infty dN_0 P_0(N_0) N_0 \int_{-\infty}^\infty dk_0 \frac{P(N_0, k_0)}{P_0(N_0)} (k_0 + \mu \langle N \rangle) \\ &= \frac{\hat{u} \langle N^2 \rangle + \mu \langle N \rangle^2 + \frac{\sigma^2}{\zeta^2} \langle N^2 \rangle \langle K \rangle \langle N \rangle}{1 + \frac{\sigma^2}{\zeta^2} \langle N^2 \rangle}\end{aligned}\quad (80)$$

where we recall $\langle N \rangle = \phi \langle N \rangle^*$, $\langle N^2 \rangle = \phi \langle N^2 \rangle^*$.

5.5 Productivity

Given the abundance distribution $P(N_i)$, it is possible to compute the productivity distribution – or at least the distribution of growth on external resources,

$$g_i = R_i N_i \quad (81)$$

which acts as a proxy for productivity. If the quantities that are drawn independently are K_i , α_{ij} and D_i , then straightforwardly

$$P(g_i) = \int dN_i P(g_i, N_i) \quad (82)$$

$$= \int_0^\infty dN_i dR_i P(R_i, N_i) \delta(R_i - g_i/N_i) \quad (83)$$

$$= \int_0^\infty dN_i dK_i dD_i P(K_i, N_i) P(D_i) \delta(K_i D_i - g_i/N_i) \quad (84)$$

and we get $P(K_i, N_i)$ from the previous section. In particular, taking the average over surviving species, we find (if all $K_i > 0$)

$$\langle g_i \rangle^* = \int dg_i P(g_i) g_i = \langle K_i N_i \rangle^* \quad (85)$$

where we used $\langle D_i \rangle = 1$. If however the independent quantities are R_i , A_{ij} and D_i , the calculation is less straightforward: we must now start from

$$N_0 = \frac{1}{D_0 \hat{u}} \left(R_0 - S^* A \langle N \rangle^* - \hat{\Sigma}_0 \right)$$

where

$$\hat{\Sigma}_0 = \sum_i (A_{0i} - A) N_i^*, \quad \langle \hat{\Sigma}_0 \rangle = 0, \quad \text{Var}(\hat{\Sigma}_0) = S^* \sigma_A^2 \langle N^2 \rangle^* \quad (86)$$

and we now have $P(R_0, \hat{\Sigma}_0, D_0) = P(R_0) P(\hat{\Sigma}_0) P(D_0)$, hence

$$P(N_0|R_0) = \int d\hat{\Sigma}_0 dD_0 P(\hat{\Sigma}_0) P(D_0) \delta \left(N_0 - \frac{1}{D_0 \hat{u}} \left(R_0 - S^* A \langle N \rangle^* - \hat{\Sigma}_0 \right) \right) \quad (87)$$

and we get the total productivity

$$G_{tot} = S \int_0^\infty dN_0 \int_0^\infty dR_0 R_0 N_0 P(N_0|R_0) P(R_0) \quad (88)$$

5.6 Response to perturbations

Let us now assume that some small permanent (press) perturbation ξ_i is added to the system at equilibrium (acting only on the surviving species)

$$0 = k_i + \xi_i - uN_i - \frac{\sigma}{\sqrt{s}} \sum_j a_{ij} N_j \quad (89)$$

We can write $N_i = N_i^* + x_i$, with x_i the displacement around equilibrium, which obeys

$$0 = \xi_i - \mu \langle x \rangle - ux_i - \frac{\sigma}{\sqrt{s}} \sum_j a_{ij} x_j \quad (90)$$

a very similar equation to that obeyed by equilibrium abundances, except for the fact that the carrying capacity K_i has been replaced by the perturbation ξ_i , and of course that x_i is not constrained to be positive. Following the same logic as before, we isolate the response of species 0

$$(u - \phi \sigma^2 \gamma v) x_0 = \xi_0 - \mu \langle x \rangle - \frac{\sigma}{\sqrt{s}} \sum_j a_{0j} x_j^* \quad (91)$$

If there are no constraints on x_0 (see Sec. 5.6.2 below), we can simply assume

$$\langle x_0 \rangle = \langle x \rangle = \langle x_i^* \rangle \quad (92)$$

and thus get the average displacement:

$$(u - \phi \sigma^2 \gamma v + \mu) \langle x \rangle = \langle \xi \rangle \quad (93)$$

5.6.1 Multistability order parameter

To test whether the equilibrium we are obtaining as the fixed point of the invasion process is globally stable, or might be exited through some small perturbation in the appropriate direction, we must consider the response to a random perturbation with zero mean,

$$\langle \xi \rangle = 0 \quad (94)$$

representing random shifts in relative fitness that advantage some species over others, without changing the mean species fitness. Since the mean is zero, the response to the perturbation is characterized by the variance,

$$\chi^2 = \left\langle \frac{x_i^2}{\xi^2} \right\rangle \quad (95)$$

Now the equation above, using $\langle x \rangle = 0$, can be taken to the square to get:

$$(u - \phi\sigma^2\gamma v)^2 \langle x_0^2 \rangle = \xi_0^2 + \phi\sigma^2 \langle x_j^* \rangle^2 \quad (96)$$

hence

$$\chi^2 = \frac{1}{(u - \phi\sigma^2\gamma v)^2 - \phi\sigma^2} \quad (97)$$

which is seen to diverge when $\sigma = \sigma_c$ such that

$$\phi\sigma_c = u - \phi\sigma_c^2\gamma v \quad (98)$$

and by using $(u - \phi\sigma^2\gamma v)v = 1$ we find as above in (15)

$$\sigma_c = \frac{1}{\phi} \left(\frac{u}{1 + \gamma} \right)^2 \quad (99)$$

5.6.2 Secondary extinctions

In the case of secondary extinctions following the extinction of species 0, the perturbation is

$$\xi_i = \frac{\sigma}{\sqrt{s}} a_{i0} N_0^* \quad (100)$$

Now however, the perturbation is potentially large enough that we may need to impose $x_i > -N_i^*$, meaning that the response cannot cause a species' abundance to become negative. Hence, in a way similar to our previous calculation on abundances, the average change in abundance due to the extinction is

$$\langle x \rangle = \int dN_i P(N_i) \int_{-N_i}^{\infty} dx_0 P_0(x_0) x_0 \quad (101)$$

while the fraction of secondary extinctions is given by

$$\Delta\phi = \int dN_i P(N_i) \int_{-\infty}^{-N_i} dx_0 P_0(x_0) \quad (102)$$

5.6.3 Variability

Calculations for the variability are more technical and left in Appendix.

Figure 7: Sensitivity of results to distribution of growth rates $P(R_i)$.

6 Sensitivity testing

6.1 Distributions

In this section, we test how sensitive the results above are to the choice of distributions for growth rates and interactions. Throughout, we have assumed that these distributions had finite mean and variance, which is not the case for, e.g., a power-law tailed distribution $P(x) \sim x^{-a}$ with exponent $a \leq 3$.

For reasons explained in Sec. 5.2 above, we expect results to be more sensitive to the distribution of growth rates than to that of interactions. Indeed, Fig. 7 shows that, while Gaussian or uniform distributions lead to very good agreement with predictions, distributions with fatter tails (e.g. exponential tails) – and a fortiori multimodal distributions – lead to significant deviations from our baseline predictions. It is however not difficult to adapt the theory to fit any distribution $P(R_i)$ as long as it is known.

Figure 8: Sensitivity of results to interaction strength. By varying the connectivity from 0.1 to 1, we change the average interaction strength needed to achieve the same value of μ . Each curve is labelled here by the average coupling strength needed to obtain $\mu = 10$.

6.2 Interaction strength

We now test how small the interaction strength must be compared to self-interaction for cavity predictions to work. Fig. 8 shows that when average coupling strength is significantly larger than 0.1, the method stops being as accurate, in particular with fewer species surviving than expected.

7 Extending the model

There are two main types of limitations to the model so far. The first is dynamical: we started from Lotka-Volterra dynamics, which might work as an approximation in many limits, but do not always enable us to find the equilibria reached by more complex dynamics. Since the most usual variant to Lotka-Volterra dynamics is having a nonlinear functional response, especially a saturating one, we treat that case below to show that it is still largely amenable

(at least approximately or qualitatively) to our approach.

The second limitation is structural. Even though species are heterogeneous, we have treated them as if they all belong to the same functional group: their traits are all drawn from the same distributions. These distributions may be multimodal, but that does not convey the fact that different positions in the community may be truly different in many respects at once. There are two simple ways to address this concern, one discrete and one continuous.

The continuous approximation, which we call “ordered species”, supposes that we can characterize the network by a simple axis of possible niches or roles, knowing only how various traits (e.g. carrying capacity, interaction mean and variance) correlate to linear order with position on that axis.

The discrete approximation, which we call “species groups”, instead divides the network into distinct groups, each disordered inside, but with deterministic patterns within and between groups. In essence, this is extending simple reductionist descriptions (e.g. one predator and one prey) by replacing each species by a group of heterogeneous but functionally similar species, treated like a disordered community.

7.1 Functional response

The coefficients A_{ij} come from two distinct factors: first, how much an *individual* of species i can affect an individual of species j , which may be a property of their behavior and biology, and second, how often these individuals may meet and interact. For simplicity, let us say that the individual-level effects are fixed. Then, the frequency of pair encounters is what controls interaction strength. The interaction between species i and j is independent from any other interaction in relatively “sparse” environments where any pair can meet, and adding new individuals or new species always increases the number of such meetings. On the other hand, interactions may saturate if there is a limiting factor such as plants competing only with their neighbors, or a predator taking time to handle prey, preventing it from consuming more than a given number of prey among those available in a given timespan.

The same reasoning led Holling [14] to postulate a saturating functional response in the interaction between two species: instead of being a number, A_{ij} becomes a function of abundances such as

$$\tilde{A}_{ij}(N_1, \dots, N_S) N_i N_j = \frac{A_{ij} N_i N_j}{1 + \frac{1}{\langle A \rangle N_c} \sum_k A_{ik} N_k} \quad (103)$$

with constants A_{ij} , their average $\langle A \rangle$ and the population threshold N_c . This means that the interaction between species i and j cannot keep growing indefinitely with $N_i N_j$ the product of their abundances, as in Lotka-Volterra equa-

Figure 9: Effect of saturating functional response. Symbols are simulations with a Holling Type II functional response, while solid lines are analytical predictions using the approximate linear-saturating functional response (107).

tions. Hence, the sum is bounded

$$N_i \sum_j \tilde{A}_{ij} N_j = N_i \frac{\sum_k A_{ik} N_k}{1 + \frac{1}{AN_c} \sum_k A_{ik} N_k} \quad (104)$$

$$\approx AN_i N_c \quad \text{if } \left| \sum_k A_{ik} N_k \right| \gg |AN_c| \quad (105)$$

Analytically, it is possible to deal with the saturating functional response by a simple approximation. For each species, we replace it with a piecewise linear function:

$$\sum_j \tilde{A}_{ij} N_j \approx F \left(\sum_k A_{ik} N_k \right), \quad F(x) = \begin{cases} x & |x| < |AN_c| \\ AN_c & |x| \geq |AN_c| \end{cases} \quad (106)$$

then at the community level, we use the real saturating function to evaluate the expected fraction f of species whose interactions are saturated.

For simplicity, we will “cheat” and instead apply the piecewise functional response to the *couplings* (which differ from interactions if species have heterogeneous self-interaction D_i), i.e.

$$0 = K_0 - N_0 - \tilde{F} \left(\sum_i \alpha_{0i} N_i \right), \quad \tilde{F}(x) = \begin{cases} x & |x| < |\mu N_c / s| \\ \mu N_c / s & |x| \geq |\mu N_c / s| \end{cases} \quad (107)$$

Avoiding this simplification is possible but requires keeping track of more parameters, and does not make much of a qualitative change.

Going through the same steps as before, then for each species

$$N_0 = \frac{1}{u - (1-f)\phi\sigma^2\gamma v} \begin{cases} K_0 - \mu\phi\langle N \rangle^* - \frac{\sigma}{\sqrt{s}} \sum_i a_{0i} N_i^* & \text{with probability } 1-f \\ K_0 - \mu N_c / s & \text{with probability } f \end{cases} \quad (108)$$

Notice the $1-f$ factor in front of the feedback term in the denominator: only species whose interactions have not saturated will contribute to the feedback of the invader on itself through its interaction partners. The mean is now

$$\langle N_0 \rangle = \frac{\langle K \rangle - \mu(fN_c/s + (1-f)\phi\langle N \rangle^*)}{u - (1-f)\phi\sigma^2\gamma v} \quad (109)$$

while taking the variance, we find

$$\sigma_0^2 = \frac{\zeta^2 + (1-f)^2\sigma^2\phi\langle N^2 \rangle^*}{(u - (1-f)\phi\sigma^2\gamma v)^2} \quad (110)$$

The rest of the calculations follow as in Sec. 5.3, except we additionally need to know the fraction f of species with saturated interactions. This is where the

real saturation function reappears (slightly rearranged):

$$f = \int_{-\infty}^{\infty} dx P \left(x = \sum_k \alpha_{0k} N_k \right) \frac{1}{1 + \left| \frac{\mu N_c}{sx} \right|} \quad (111)$$

where, as usual, we can approximate $P(x)$ by a Gaussian distribution with

$$\langle x \rangle = |\mu| \phi \quad (112)$$

$$\text{Var}(x) = \sigma^2 \phi \langle N^2 \rangle^* \quad (113)$$

Note that if we want to make an equivalent model without saturation, we can apply the equivalence

$$\tilde{\mu} = \mu \left(f \frac{N_c}{s \langle N \rangle^*} + (1-f)\phi \right), \quad \tilde{\sigma} = (1-f)\sigma, \quad \tilde{\gamma} = \frac{\gamma}{1-f} \quad (114)$$

which is possible only if $\tilde{\gamma} \in [-1, -1]$. Then, a Lotka-Volterra system with parameters $\tilde{\mu}$, $\tilde{\sigma}$ and $\tilde{\gamma}$ will have identical equilibrium properties to the system with a saturating functional response studied here. Of course, this equivalence can only be computed once we know f and $\langle N \rangle^*$, i.e. after having solved the calculation above for this system.

In Fig. 9, we show that predictions from our calculations match simulation results quite well, despite many simplifications which tend to create a more abrupt transition between the saturated and unsaturated regimes. By the same reasoning, we can easily extend the approach to any functional response, including an approximation of Type-III sigmoid response, with two thresholds instead of a single one.

7.2 Heterogeneous means

Let us come back to the cavity method:

$$0 = K_0 - N_0 - \sum_i^{S^*} A_{0i} N_i \quad (115)$$

Now let us assume that

$$A_{0i} = \langle A \rangle_0 + \frac{\sigma_0}{\sqrt{S}} a_{0i} \quad (116)$$

where the mean and the variance both depend on species 0.

$$0 = K_0 - S^* \langle A \rangle_0 \langle N \rangle^* - N_0 - \frac{\sigma_0}{\sqrt{S}} \sum_i^{S^*} a_{0i} N_i \quad (117)$$

where we still assume that $\langle N \rangle$ over partners of 0 does not differ from the community average. Then,

$$N_i \approx N_i^* - \frac{dN_i}{dK_i} \frac{\sigma_i}{\sqrt{S}} a_{i0} N_0 \quad (118)$$

hence

$$\left(1 - \frac{1}{S} \sum_i \sigma_i \sigma_0 a_{i0} a_{0i} \frac{dN_i}{dK_i}\right) N_0 = K_0 - S^* \langle A \rangle_0 \langle N \rangle^* - \frac{\sigma_0}{\sqrt{S}} \sum_i^{S^*} a_{0i} N_i^* \quad (119)$$

$$\approx (1 - \phi \sigma_0 \sigma \gamma v) N_0 \quad (120)$$

where

$$v = \left\langle \frac{1}{1 - \phi \sigma_0 \sigma \gamma v} \right\rangle \approx \frac{1}{1 - \phi \sigma^2 \gamma v} \quad (121)$$

and finally

$$\langle N_0 \rangle = v(\langle K \rangle - \phi \mu \langle N \rangle^*) \quad (122)$$

$$\langle N_0^2 \rangle = \langle N_0 \rangle^2 + v^2(\sigma_K^2 + S^{*2}(\sigma_A^2) \langle N \rangle^{*2} - 2S^*(\langle K_0 A_0 \rangle - \langle K \rangle \langle A \rangle) + \sigma^2 \phi \langle N^2 \rangle^*) \quad (123)$$

7.3 Species groups

The equations can be extended to any structure comprised of discrete groups, with disordered interactions within and between groups, but different statistics for each set of interactions. Coming back to the equilibrium equation, we can write for species i in group x (which contains S^x species) as

$$0 = K_i^x - N_i^x - \sum_{j,y} \alpha_{ij}^{xy} N_j^y \quad (124)$$

Thus, we now have vector $\mathbf{K}$ and $\boldsymbol{\zeta}$ and matrices $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$, $\boldsymbol{\gamma}$, defined by

$$\alpha_{ij}^{xy} = \frac{\mu^{xy}}{S} + \frac{\sigma^{xy}}{\sqrt{S}} a_{ij}^{xy} \quad (125)$$

with

$$\langle a \rangle = 0, \quad \langle a^2 \rangle = 1, \quad \langle a_{ij}^{xy} a_{ji}^{yx} \rangle = \gamma^{xy}. \quad (126)$$

The equations to solve are the same as above, except there are now four equations per group, all coupled: for each group we solve

$$\phi^x = \int_0^\infty dN_0 P_0^x(N_0) \quad (127)$$

$$\langle N \rangle^x = \int_0^\infty dN_0 P_0^x(N_0) N_0 \quad (128)$$

$$\langle N^2 \rangle^x = \int_0^\infty dN_0 P_0^x(N_0) N_0^2 \quad (129)$$

$$V^x = \frac{\phi^x}{\tilde{u}^x} \quad (130)$$

with

$$\langle N_0 \rangle^x = \frac{\langle K \rangle^x - \sum_y \mu^{xy} \rho^y \langle N \rangle^y}{u^x}, \quad (131)$$

$$(\sigma_0^2)^x = \frac{(\zeta^x)^2 + \sum_y (\sigma^{xy})^2 \rho^y \langle N^2 \rangle^y}{(u^x)^2}, \quad (132)$$

$$\tilde{u}^x = 1 - \sum_y \rho^y V^y \gamma^{xy} \sigma^{xy} \sigma^{yx}. \quad (133)$$

where

$$\rho^x = \frac{S^x}{S}. \quad (134)$$

7.4 Ordered species

We can take the continuous limit of the mixture model:

$$\phi(x) = \int_0^\infty dN_0 P_0(N_0, x) \quad (135)$$

$$\langle N \rangle(x) = \int_0^\infty dN_0 P_0(N_0, x) N_0 \quad (136)$$

$$\langle N^2 \rangle(x) = \int_0^\infty dN_0 P_0(N_0, x) N_0^2 \quad (137)$$

$$V(x) = \frac{\phi(x)}{\tilde{u}(x)} \quad (138)$$

with $P_0(n, x)$ a Gaussian with mean $\langle N_0 \rangle(x)$ and variance $\sigma_0^2(x)$, given by

$$\langle N_0 \rangle(x) = \frac{\langle K \rangle(x) - \int dy \mu(x, y) \rho(y) \langle N \rangle(y)}{\tilde{u}(x)}, \quad (139)$$

$$\sigma_0^2(x) = \frac{\zeta^2(x) + \int dy \sigma^2(x, y) \rho(y) \langle N^2 \rangle(y)}{\tilde{u}^2(x)}, \quad (140)$$

$$\tilde{u}(x) = 1 - \gamma \int dy \rho(y) V(y) \sigma(x, y) \sigma(y, x). \quad (141)$$

where $\rho(x)$ is the density of species with rank x .

There are two cases that allow for a simple solution: first, if all functions depend only on x . Second, if all functions of x and y depend only on the difference $x - y$, and these functions are either exponential or linear, see results in Fig. 10. By construction, in all these situation $\gamma(x, y) = \gamma(y, x)$ must be a constant.

7.5 One-sided order

We call one-sided order the situation where functions of x and y depend only on x .

$$\mu(x, y) = \mu(x) \quad (142)$$

$$\sigma^2(x, y) = \sigma^2(x) \quad (143)$$

$$\langle K \rangle(x) = K(x) \quad (144)$$

$$\zeta^2(x) = \zeta^2(x) \quad (145)$$

Then, we have

$$\langle N_0 \rangle(x) = \frac{\langle K \rangle(x) - \mu(x) \langle \varphi N \rangle}{\tilde{u}(x)}, \quad (146)$$

$$\sigma_0^2(x) = \frac{\zeta^2(x) + \sigma^2(x) \langle \varphi N^2 \rangle}{\tilde{u}^2(x)}, \quad (147)$$

$$\tilde{u}(x) = 1 - \gamma \sigma(x) \langle \varphi V \sigma \rangle. \quad (148)$$

and we can compute

$$\langle \varphi N \rangle = \int dx \rho(x) \int_0^\infty dN_0 P_0(N_0, x) N_0 \quad (149)$$

$$\langle \varphi N^2 \rangle = \int dx \rho(x) \int_0^\infty dN_0 P_0(N_0, x) N_0^2 \quad (150)$$

$$\langle \varphi V \sigma \rangle = \int dx \rho(x) \frac{\phi(x) \sigma(x)}{\tilde{u}(x)} \quad (151)$$

For instance, we can build a tradeoff reminiscent of competition-colonization models by having

$$\mu(x, y) = 2\mu x \quad (152)$$

$$\langle K \rangle(x) = 2Kx \quad (153)$$

with $\mu > 0$ and σ and ζ constant. In this case, low rank x corresponds to high competitive ability (no competition experienced) but low colonization ability.

7.5.1 Exponential order

We can choose exponential functions

$$\mu(x, y) = \mu c_\mu^{x-y} \quad (154)$$

$$\sigma^2(x, y) = \sigma^2 c_\sigma^{x-y} \quad (155)$$

$$\langle K \rangle(x) = K c_K^x \quad (156)$$

$$\zeta^2(x) = \zeta^2 c_\zeta^x \quad (157)$$

Figure 10: Results for ordered species with exponential and linear scaling. The bottom-right panel is $\langle x \rangle^*$ the average value of the niche value for surviving species, given that niche values are initially distributed uniformly over $[0, 1]$.

Then, we have

$$\langle N_0 \rangle(x) = \frac{K c_K^x - \mu c_\mu^x \langle N c_\mu^{-x} \rangle}{\tilde{u}(x)}, \quad (158)$$

$$\sigma_0^2(x) = \frac{\zeta^2 c_\zeta^x + \sigma^2 c_\sigma^x \langle N^2 c_\sigma^{-x} \rangle}{\tilde{u}^2(x)}, \quad (159)$$

$$\tilde{u}(x) = 1 - \gamma \langle V \rangle \sigma^2. \quad (160)$$

and therefore we must compute

$$\langle N c_\mu^{-x} \rangle = \int dx \rho(x) c_\mu^{-x} \int_0^\infty dN_0 P_0(N_0, x) N_0 \quad (161)$$

$$\langle N^2 c_\sigma^{-x} \rangle = \int dx \rho(x) c_\sigma^{-x} \int_0^\infty dN_0 P_0(N_0, x) N_0^2 \quad (162)$$

$$\langle V \rangle = \int dx \rho(x) \frac{\phi(x)}{\tilde{u}(x)} \quad (163)$$

to close the equations. A typical example is metabolic scaling in trophic interactions: given metabolic rate m_i and its scaling with body mass M_i , $m_i \propto M_i^{1/4}$, we expect self-interaction $D_i \propto m_i$ and interaction $A_{ji} \propto m_j$, hence $\alpha_{ji} \propto m_j/m_i = (1/4)^{x-y}$ with the niche values given by the logarithm of body mass $x = \log M_i$ and $y = \log M_j$.

7.5.2 Linear order

If we choose linear functions instead

$$\mu(x, y) = \mu + c_\mu(x - y) \quad (164)$$

$$\sigma^2(x, y) = \sigma^2 + c_\sigma(x - y) \quad (165)$$

$$\langle K \rangle(x) = K + c_K x \quad (166)$$

$$\zeta^2(x) = \zeta^2 + c_\zeta x \quad (167)$$

Then we have

$$\langle N_0 \rangle(x) = \frac{K + c_K x - (\mu + c_\mu x) \langle \varphi N \rangle + c_\mu \langle \varphi N x \rangle}{\tilde{u}(x)}, \quad (168)$$

$$\sigma_0^2(x) = \frac{\zeta^2 + c_\zeta x + (\sigma^2 + c_\sigma x) \langle \varphi N^2 \rangle - c_\sigma \langle \varphi N^2 x \rangle}{\tilde{u}^2(x)}, \quad (169)$$

$$\tilde{u}(x) \approx 1 - \gamma \langle \varphi v \rangle \sigma^2. \quad (170)$$

(in fact $\sigma(x, y)\sigma(y, x) = \sigma^2 \sqrt{1 - c_\sigma^2(x - y)^2/\sigma^4}$ but this makes separation impossible, so we must assume that c_σ is small enough). Hence there are two sets of coupled unknowns: for $\Psi = N, N^2, V$ we must compute

$$\langle \varphi \Psi \rangle = \int dx P(x) \int_0^\infty dN_0 P_0(N_0, x) \Psi \quad (171)$$

and (except for V),

$$\langle \varphi \Psi x \rangle = \int dx \rho(x) x \int_0^\infty dN_0 P_0(N_0, x) \Psi. \quad (172)$$

References

- [1] Stefano Allesina and Si Tang. Stability criteria for complex ecosystems. *Nature*, 483(7388):205–208, 2012.
- [2] Jean-François Arnoldi, Michel Loreau, and Bart Haegeman. Resilience, reactivity and variability: A mathematical comparison of ecological stability measures. *Journal of theoretical biology*, 389:47–59, 2016.
- [3] Matthieu Barbier, Jeff Arnoldi, and Guy Bunin. Generic assembly patterns in large ecological communities. *bioRxiv preprint 145862*, 2017.
- [4] Matthieu Barbier, Jeff Arnoldi, Michel Loreau, and Guy Bunin. Generic assembly patterns in large ecological communities. 2017.
- [5] Jordi Bascompte and Daniel B Stouffer. The assembly and disassembly of ecological networks. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1524):1781–1787, 2009.
- [6] Guy Bunin. Interaction patterns and diversity in assembled ecological communities. *arXiv preprint arXiv:1607.04734*, 2016.
- [7] Guy Bunin. Ecological communities with lotka-volterra dynamics. *Physical Review E*, 95(4):042414, 2017.
- [8] Amar Gandhi, Simon Levin, and Steven Orszag. Moment expansions in spatial ecological models and moment closure through gaussian approximation. *Bulletin of mathematical biology*, 62(4):595–632, 2000.
- [9] Jianxi Gao, Baruch Barzel, and Albert-László Barabási. Universal resilience patterns in complex networks. *Nature*, 530(7590):307–312, 2016.
- [10] M Gatto. A general minimum principle for competing populations: Some ecological and evolutionary consequences. *Theoretical Population Biology*, 37(3):369–388, 1990.
- [11] Dominique Gravel, François Massol, Elsa Canard, David Mouillot, and Nicolas Mouquet. Trophic theory of island biogeography. *Ecology letters*, 14(10):1010–1016, 2011.
- [12] Jacopo Grilli, Tim Rogers, and Stefano Allesina. Modularity and stability in ecological communities. *Nature communications*, 7, 2016.
- [13] J Nathaniel Holland, Donald L DeAngelis, and Judith L Bronstein. Population dynamics and mutualism: functional responses of benefits and costs. *The American Naturalist*, 159(3):231–244, 2002.

- [14] Crawford Stanley Holling. The functional response of predators to prey density and its role in mimicry and population regulation. *Memoirs of the Entomological Society of Canada*, 97(S45):5–60, 1965.
- [15] Robert D Holt. Community modules. In *Multitrophic interactions in terrestrial ecosystems, 36th Symposium of the British Ecological Society*, pages 333–349. Blackwell Science Oxford, 1997.
- [16] Stephen P Hubbell. *The unified neutral theory of species abundance and diversity*. Princeton University Press, Princeton, 2001.
- [17] Vincent AA Jansen and Giorgos D Kokkoris. Complexity and stability revisited. *Ecology Letters*, 6(6):498–502, 2003.
- [18] David A Kessler and Nadav M Shnerb. Generalized model of island biodiversity. *Physical Review E*, 91(4):042705, 2015.
- [19] Richard Levins and David Culver. Regional coexistence of species and competition between rare species. *Proceedings of the National Academy of Sciences*, 68(6):1246–1248, 1971.
- [20] Michel Loreau. *From populations to ecosystems: Theoretical foundations for a new ecological synthesis*. Princeton University Press, 2010.
- [21] Robert Mac Arthur. Species packing, and what competition minimizes. *Proceedings of the National Academy of Sciences*, 64(4):1369–1371, 1969.
- [22] Robert H MacArthur and Edward O Wilson. *Theory of Island Biogeography*. Princeton University Press, 2015.
- [23] Ramon Margalef. Diversity, stability and maturity in natural ecosystems. In *Unifying concepts in ecology*, pages 151–160. Springer, 1975.
- [24] Robert M May. Will a large complex system be stable? *Nature*, 238:413–414, 1972.
- [25] Robert M May and Roy M Anderson. Parasite—host coevolution. *Parasitology*, 100(S1):S89–S101, 1990.
- [26] Claire Mazancourt, Forest Isbell, Allen Larocque, Frank Berendse, Enrica Luca, James B Grace, Bart Haegeman, H Wayne Polley, Christiane Roscher, Bernhard Schmid, et al. Predicting ecosystem stability from community composition and biodiversity. *Ecology letters*, 16(5):617–625, 2013.
- [27] M Mézard. The space of interactions in neural networks: Gardner’s computation with the cavity method. *Journal of Physics A: Mathematical and General*, 22(12):2181, 1989.
- [28] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

- [29] Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Co Inc, 1987.
- [30] Andrea Montanari, M Müller, and Marc Mézard. Phase diagram of random heteropolymers. *Physical review letters*, 92(18):185509, 2004.
- [31] Nicolas Mouquet and Michel Loreau. Community patterns in source-sink metacommunities. *The american naturalist*, 162(5):544–557, 2003.
- [32] Manfred Oppen and Sigurd Diederich. Phase transition and $1/f$ noise in a game dynamical model. *Physical review letters*, 69(10):1616, 1992.
- [33] Pierre Paga and Reimer Kühn. Contagion in an interacting economy. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(3):P03008, 2015.
- [34] Bernhard Schmid, Andy Hector, Prasenjit Saha, and Michel Loreau. Bio-diversity effects and transgressive overyielding. *Journal of Plant Ecology*, 1(2):95–102, 2008.
- [35] Alessandro Seganti, Andrea De Martino, and Federico Ricci-Tersenghi. Boolean constraint satisfaction problems for reaction networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(09):P09009, 2013.
- [36] Torsten Thomas, Lucas Moitinho-Silva, Miguel Lurgi, Johannes R Björk, Cole Easson, Carmen Astudillo-García, Julie B Olson, Patrick M Erwin, Susanna López-Legentil, Heidi Luter, et al. Diversity, structure and convergent evolution of the global sponge microbiome. *Nature Communications*, 7, 2016.
- [37] Mikhail Tikhonov. Community-level cohesion without cooperation. *Elife*, 5:e15747, 2016.
- [38] Mikhail Tikhonov and Remi Monasson. A collective phase in resource competition in a highly diverse ecosystem. *arXiv preprint arXiv:1609.01270*, 2016.
- [39] David Tilman. *Resource competition and community structure*. Princeton university press, 1982.
- [40] Chengyi Tu, Jacopo Grilli, and Samir Suweis. Universality of resilience patterns in generalized lotka volterra dynamics and beyond. *arXiv preprint arXiv:1606.09630*, 2016.
- [41] Yoshimi Yoshino, Tobias Galla, and Kei Tokita. Statistical mechanics and stability of a model eco-system. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(09):P09003, 2007.

8 Appendices

A. The role of disorder

Computing the equilibrium state in effect requires inverting the matrix a_{ij} in the equation:

$$0 = k_i - uN_i - \sigma_A \sum_j a_{ij} N_j \quad (173)$$

In principle we could expand this equation, substituting

$$N_j = \frac{k_j - \sigma_A \sum_k a_{jk} N_k}{u} \quad (174)$$

into the sum, and so on recursively, causing all paths $a_{ij}a_{jk}\dots$ ($i \rightarrow j \rightarrow k \dots$) in the community matrix to appear in the first equation above. The crucial idea of disordered systems methods is that, for any reasonably complex system, these paths soon lose coherence and start to interfere with each other. These interferences are so complex as to produce seemingly-random results.

One way to see it is that each term in the sum is a (possibly very long) product of coefficients a_{kl} and is weakly correlated with other such terms – unless the same coefficients appear again and again in the same combinations in many terms. That can happen in a very structured community, or one with only small loops (the extreme is a collection of independent few-species motifs: in each motif, there are not enough paths for them to interfere and create a seemingly random result). This enables the cavity trick, which retains only the first-order feedback of a species on itself, and for the rest needs only the mean and variance of interactions, as the effect of interactions is essentially random.

Let us see a precise example of these ideas. In the main text, we noted that the feedback of a species on itself through the rest of the network is a self-averaging quantity. Consider the following sum:

$$\sum_{i,j}^{s^*} a_{0i} v_{ij} a_{j0} \quad (175)$$

which occurs in (50) (we recall $s^* = \phi s$ is the average number of surviving interaction partners for species 0). We separate the sum into two parts:

$$\sum_i^{s^*} a_{0i}^2 v_{ii} + \sum_{i \neq j}^{s^*} a_{0i} v_{ij} a_{j0} \quad (176)$$

Recall that $\langle a_{0i} a_{i0} \rangle = \gamma$ while $\langle a_{0i} \rangle = 0$. Hence, the second sum is a random walk with $s^*(s^* - 1)$ terms, and its total scales like s^* . Now, we should note that

$$v_{ij} = \frac{dN_i}{dk_j}$$

for $i \neq j$ is created by the interaction between species i and j , which is at best proportional to $\sigma/\sqrt{s}$ (if the species are linked directly) or even weaker (if they interact only indirectly).

Finally, we get the respective scalings:

$$\sum_i^{s^*} a_{0i} a_{i0} v_{ii} \propto s^* \gamma v \quad (177)$$

$$\sum_{i \neq j}^{s^*} a_{0i} v_{ij} a_{j0} \propto s^* \frac{\sigma}{\sqrt{s}} v \quad (178)$$

In other words, even though the sum over $i \neq j$ has many more terms, it ends up being of order $1/\sqrt{s}$ compared to the first sum. That sum itself has variations of order $1/\sqrt{s}$ compared to its mean, depending on which species is isolated as species 0. As s becomes larger, the feedback thus tends to become a single number, no matter which species we consider.

B. Linear response

Typical cavity method (Press)

Notation: $a_i^j x_j = \sum_j a_{ij} x_j$. $\langle y. \rangle = \frac{1}{S} \sum_i y_i$. Let ξ be a random press-perturbation. Its long term effect on the equilibrium associated to the Jacobian B is

$$x = -B^{-1} \xi$$

$$u x_i = \xi_i - \mu \langle x. \rangle - \sigma a_i^j x_j - \sigma a_{i0} x_0; \quad u = 1 - \mu/S$$

Where we isolated the effect of a generic species 0 (this is the starting point of the cavity method). Assuming a small effect of that species we may write

$$x_j = x_{j/0} - \sigma x_0 \mathcal{V}_j^k a_{k0}; \quad \mathcal{V}_j^k := \frac{\partial x_j}{\partial \xi_k}$$

Under the replica symmetry, species 0 is generic and its addition or removal does not change any statistical properties of community, in particular $\langle x_{./0} \rangle = \langle x. \rangle$. Then

$$u x_0 = \xi_0 - \mu \langle x. \rangle - \sigma a_0^j x_{j/0} + \sigma^2 x_0 \underbrace{\mathcal{V}_j^k a_0^j a_{k0}}_{\approx \gamma \langle \mathcal{V}_{ii} \rangle}$$

$$(u - \sigma^2 \gamma \mathcal{V}) x_0 \approx \xi_0 - \mu \langle x. \rangle - \sigma a_0^j x_{j/0}; \quad \mathcal{V} = \langle \mathcal{V}_{ii} \rangle$$

Taking the square

$$(u - \sigma^2 \gamma \mathcal{V})^2 x_0^2 \approx \xi_0^2 + \underbrace{\sigma^2 a_0^j a_0^i x_{j/0} x_{i/0}}_{\approx \sigma^2 \langle x^2 \rangle} + \mu^2 \langle x. \rangle^2 + \mu \sigma \langle x. \rangle a_0^j x_{j/0} + \mathcal{O}(\xi_0)$$

taking the mean over ξ_0 , knowing that $\mathbb{E}\xi_i = 0$; $\mathbb{E}\xi_i\xi_j = \delta_{ij}$, and by linearity, $\mathbb{E}_\xi x_i = 0$, gives

$$(u - \sigma^2\gamma\mathcal{V})^2 \mathbb{E}x_0^2 \approx 1 + \sigma^2 \langle \mathbb{E}x^2 \rangle$$

Taking the mean over all species gives

$$\chi^2 := \langle \mathbb{E}x^2 \rangle \approx \frac{1}{(u - \sigma^2\gamma\mathcal{V})^2 - \sigma^2}$$

To conclude we need $\mathcal{V}$. For that we write

$$(u - \sigma^2\gamma\mathcal{V}) \mathcal{V} \approx 1 - \underbrace{\sigma \langle a^j \mathcal{V}_j^j \rangle}_{\approx 0}$$

Hence it is solution to (knowing that at $\sigma = 0$, $\mathcal{V} = u$)

$$\sigma^2\gamma\mathcal{V}^2 - u\mathcal{V} + 1 = 0$$

we thus have

$$\chi^2 \approx \frac{1}{u^2(1-v)^2 - \sigma^2}; \quad v = \frac{\sigma^2\gamma\mathcal{V}}{u} = \frac{1 - \sqrt{1 - \frac{4\sigma^2\gamma}{u^2}}}{2}$$

In the Lotka-Volterra assembly perspective, B is the interaction matrix and only a fraction ϕ of species survive. We still have $u = 1 - \mu/S$ but

$$\chi^2 \approx \frac{1}{u^2(1-v)^2 - \phi\sigma^2}$$

Externally applied press

In the assembly perspective, the Jacobian matrix reads

$$J_{ij} = -\frac{r_i}{k_i} N_i^* B_{ij}$$

where r_i is the growth rate of species i and k_i its carrying capacity. A press acting proportionally to $N_i^{*\nu}$ will induce a displacement

$$x_i = \frac{k_i}{r_i} N_i^{*\nu-1} \zeta_i - a_i^j x_j$$

so we can simply adapt the above calculations by replacing ξ_i by $\frac{k_i}{r_i} N_i^{*\nu-1} \zeta_i$. This gives us

$$\langle \mathbb{E}_\zeta x^2 \rangle \approx \chi^2 \left\langle \left(\frac{k}{r}\right)^2 N^{*2\nu-2} \right\rangle$$

To compute (mean case) variability we must solve the Lyapunov equation

$$BC + CB^T = -\mathbb{I} \Leftrightarrow 2C + AC + (AC)^T = \mathbb{I}$$

Off diagonal terms, if $\varphi = \frac{S-1}{S}$, read

$$2C_{kj} = -\frac{1-u}{u} (C_{kk} + C_{jj} + \varphi \langle C_{\cdot \neq k} \rangle + \varphi \langle C_{\cdot \neq j} \rangle) - \frac{\sigma}{u} (a_k^l C_{lj} + a_j^l C_{lk})$$

taking the mean, the self-averaging of the terms proportional to σ makes the problem reduce to the mean-field case, which gives us the relationship

$$\langle C_{\neq} \rangle \approx -\frac{S-\mu}{S+\mu(\varphi S-1)} \langle C_{=} \rangle$$

The interaction diversity (σ) enters on the diagonal (variances):

$$2C_{jj} = 1 - \sigma a_j^k 2C_{kj} - 2\mu\varphi \langle C_{\neq j} \rangle$$

in replacing C_{kj} in $\sigma a_j^k 2C_{kj}$ by its expression written above, we only need to track terms proportional to σ^2 , the other will self average. So:

$$2C_{jj} \approx 1 + \frac{\sigma^2}{u} \underbrace{a_j^k a_k^l C_{lj}}_{\approx \gamma C_{jj}} + \frac{\sigma^2}{u} \underbrace{a_j^k a_j^l C_{lk}}_{\approx \langle C_{=} \rangle} - 2\mu\varphi \langle C_{\neq j} \rangle$$

$$2 \left(1 - \frac{\sigma^2 (\gamma + 1)}{2u} \right) \langle C_{=} \rangle \approx 1 - 2\mu\varphi \langle C_{\neq} \rangle$$

so that

$$\langle C_{=} \rangle \approx \mathcal{V} = \frac{1}{2\mathcal{I}}; \mathcal{I} = \mathcal{I}_{\text{MF}} - \frac{\sigma^2}{2} (\gamma + 1) \frac{S}{S-\mu}; \mathcal{I}_{\text{MF}} = \frac{1+\mu\varphi}{1+\mu\varphi/u} = 1 - \mu/S + \mathcal{O}(S^{-1})$$

when $\mu = 0$, $u = 1$, $\mathcal{I}_{\text{MF}} = 1$. When $\sigma = 0$, the expression for the mean variance is exact (mean-field case).

Demographic noise in the assembly perspective for identical carrying capacities and growth rates

B must be replaced with the Jacobian matrix

$$J = -D_* B = -D_* - D_* A; D_* = \text{diag}(N^*)$$

Where the N_i^* are the equilibrium abundances of surviving species. The Lyapunov equation becomes

$$D_* BC + CB^T D_* = -D_* \Leftrightarrow D_* C + CD_* + D_* AC + (AC)^T D_* = D_*$$

Off diagonal terms if $\varphi = \frac{S^*-1}{S} = \phi - 1/S \approx \phi$ read

$$u(N_k^* + N_j^*)C_{kj} + N_k^* \left(\mu\varphi \langle C_{\neq j} \rangle + \frac{\mu}{S} C_{jj} + \sigma a_k^l C_{lj} \right) + N_j^* \left(\mu\varphi \langle C_{\neq k} \rangle + \frac{\mu}{S} C_{kk} + \sigma a_j^l C_{lk} \right) = 0$$

let us write

$$\mathcal{N}_k^{(j)} = \frac{2N_k^*}{N_k^* + N_j^*}; \quad \mathcal{N}_k^{(k)} = 1 \text{ and } \langle \mathcal{N}^{(\cdot)} \rangle = 1$$

then

$$2C_{kj} = -\frac{\mu/S}{u} \left(C_{kk} \mathcal{N}_j^{(k)} + \mathcal{N}_k^{(j)} C_{jj} + \varphi \langle C_{\neq k} \rangle \mathcal{N}_j^{(k)} + \varphi \mathcal{N}_k^{(j)} \langle C_{\neq j} \rangle \right) - \frac{\sigma}{u} \left(\mathcal{N}_k^{(j)} a_k^l C_{lj} + \mathcal{N}_j^{(k)} a_j^l C_{lk} \right)$$

taking the mean over $k \neq j$ gives the same approximate relationship between $\langle C_{\neq} \rangle$ and $\langle C_{=} \rangle$ as for the mean field case (and is exact when $\sigma = 0$). On the diagonal:

$$2C_{jj} = 1 - \sigma a_j^k 2C_{kj} - 2\mu\varphi \langle C_{\neq j} \rangle$$

The abundances only enter via the expression of C_{kj} but always as a ratio $\mathcal{N}_k^{(j)}$ or $\mathcal{N}_j^{(k)}$ whose average is 1. We get

$$2C_{jj} \approx 1 + \frac{\sigma^2}{u} \underbrace{\mathcal{N}_j^{(k)} a_j^k a_k^l C_{lj}}_{\approx \phi \gamma \langle \mathcal{N}^{(\cdot)} \rangle C_{jj}} + \frac{\sigma^2}{u} \underbrace{\mathcal{N}_k^{(j)} a_j^k a_j^l C_{lk}}_{\approx \phi \gamma \langle \mathcal{N}^{(j)} \rangle \langle C_{=} \rangle} - 2\mu\varphi \langle C_{\neq j} \rangle$$

taking the mean over j gives

$$2 \langle C_{=} \rangle \approx 1 + \frac{\sigma^2}{u} \phi (\gamma + 1) \langle C_{=} \rangle - 2\mu\varphi \langle C_{\neq} \rangle$$

so that

$$\langle C_{=} \rangle \approx \mathcal{V} = \frac{1}{2 \mathcal{I}_{\text{MF}} - \frac{\sigma^2 \phi}{2u} (\gamma + 1)}$$

when $\mu = 0$, $u = 1$, $\mathcal{I}_{\text{MF}} = 1$. All in all the approximation for (adimensional) invariability is

$$\mathcal{I} \approx \mathcal{I}_{\text{MF}} - \frac{\sigma^2}{2} (\gamma + 1) \frac{\phi S}{S - \mu}; \quad \mathcal{I}_{\text{MF}} = \frac{1 + \mu\phi}{1 + \mu\phi/u} = 1 - \mu/S + \mathcal{O}(S^{-1})$$

More general case in the assembly perspective

In general B must be replaced with the Jacobian matrix

$$J = -DB = -D - DA; \quad D = \text{diag}(N^* r/k)$$

Under stochastic perturbations scaling as $N_i^{*\nu/2}$ ($\nu = 1$ demographic noise, $\nu = 2$ environmental noise) the Lyapunov equation becomes

$$DBC + CB^\top D = -D_*^\nu \Leftrightarrow DC + CD + DAC + (AC)^\top D = D_*^\nu$$

For off-diagonal terms we may reproduce the previous calculations by replacing N_i^* by $\mathcal{K}_i = N_i^* r_i / k_i$ and

$$\mathcal{K}_k^{(j)} = \frac{2\mathcal{K}_k}{\mathcal{K}_k + \mathcal{K}_j}; \quad \mathcal{K}_k^{(k)} = 1 \text{ and } \langle \mathcal{K}^{(\cdot)} \rangle = 1$$

On the diagonal however we now have:

$$2C_{jj} + A_j^l C_{lj} + A_j^l C_{lj} = \frac{k_j}{r_j} N_j^{*\nu-1}$$

which, as above, leads us to

$$\left(1 - \frac{\sigma^2}{2u} \phi(\gamma + 1)\right) \langle C_{=} \rangle \approx \left\langle \frac{k}{2r} N^{*\nu-1} \right\rangle - \mu \phi \langle C_{\neq} \rangle$$

so that

$$\langle C_{=} \rangle \approx \mathcal{V}_\nu = \left\langle \frac{k}{r} N^{*\nu-1} \right\rangle \mathcal{V}$$

In particular, for $\nu = 0, 1, 2$ (immigration, demographic and environmental noise, respectively):

$$\mathcal{V}_0 = \left\langle \frac{k/r}{N^*} \right\rangle \mathcal{V}; \quad \mathcal{V}_1 = \left\langle \frac{k}{r} \right\rangle \mathcal{V}; \quad \mathcal{V}_2 = \left\langle \frac{k}{r} N^* \right\rangle \mathcal{V}$$

Taylor's law

relates species mean abundance and their variance. In our context, under some stochastic noise acting independently on species (mean case scenario) we have that

$$C_{jj} \approx \frac{1}{2} \frac{k_j}{r_j} N_j^{*\nu-1} + \frac{\sigma^2 \gamma}{2} \frac{\phi S}{S - \mu} \left\{ \langle \mathcal{K}_j^{(\cdot)} \rangle C_{jj} + \langle \mathcal{K}_j^{(j)} \rangle \langle C_{=} \rangle \right\} - \mu \phi \langle C_{\neq j} \rangle$$

We roughly approximate $\langle \mathcal{K}_j^{(\cdot)} \rangle$ and $\langle \mathcal{K}_j^{(j)} \rangle$ by 1. Using that

$$2 \langle C_{\neq j} \rangle \approx -\frac{1-u}{u} (\langle C_{=} \rangle + C_{jj} + \phi \langle C_{\neq} \rangle + \phi \langle C_{\neq j} \rangle)$$

so

$$-\mu \phi \langle C_{\neq j} \rangle \approx \frac{\mu \phi \mu}{2(S - \mu) + \phi \mu} (\langle C_{=} \rangle + C_{jj} + \phi \langle C_{\neq} \rangle)$$

If we define

$$\frac{\mu \phi \mu}{2(S - \mu) + \phi \mu} = f_S(\mu, \phi)$$

This leads us to

$$\left(1 - \frac{\sigma^2 \gamma}{2} \frac{\phi S}{S - \mu} - f_S(\mu, \phi)\right) C_{jj} \approx \frac{1}{2} \frac{k_j}{r_j} N_j^{*\nu-1} + \left(\frac{\sigma^2 \gamma}{2} \frac{\phi S}{S - \mu} + f_S(\mu, \phi)\right) \mathcal{V}_\nu$$

if μ is of order zero in $1/S$ then f_S is of order one and for S large enough

$$\left(1 - \frac{\sigma^2 \phi \gamma}{2}\right) C_{jj} \approx \frac{1}{2} \frac{k_j}{r_j} N_j^{*\nu-1} + \frac{\sigma^2 \phi \gamma}{2} \mathcal{V}_\nu$$

We thus expect an exact power law $\text{Var}(N_i) = C_{jj} \sim N_i^{\nu-1}$ when $\gamma = 0$ and σ^2 large enough (so that N_j and k_j become uncorrelated).

C. Other calculations

Generating A with correlation to K

Here is how to generate A_{ij} with prescribed correlation to K_i :

$$A_{ij} = \frac{\mu}{S} + f(K_i) + \sigma a_{ij} \quad (179)$$

$$\langle A_{ij} \rangle_j = \frac{\mu}{S} + f(K_i) \quad (180)$$

$$\langle A_{ij} \rangle = \frac{\mu}{S} + \langle f(K_i) \rangle \quad (181)$$

Let's set $\langle K_i \rangle = 1$, then:

$$\langle K_i A_{ij} \rangle = \frac{\mu}{S} + \langle f(K_i) K_i \rangle \quad (182)$$

$$\langle K_i A_{ij} \rangle - \langle K_i \rangle \langle A_{ij} \rangle = \langle f(K_i) K_i \rangle - \langle f(K_i) \rangle \quad (183)$$

$$\langle A_{ij}^2 \rangle = \frac{\mu^2}{S^2} + 2 \langle f(K_i) \rangle \frac{\mu}{S} + \langle f(K_i)^2 \rangle + \frac{\sigma^2}{S} \quad (184)$$

$$\langle A_{ij} A_{ji} \rangle = \frac{\mu^2}{S^2} + 2 \langle f(K_i) \rangle \frac{\mu}{S} + \langle f(K_i)^2 \rangle + \frac{\sigma^2 \gamma}{S} \quad (185)$$

$$\langle A_{ij}^2 \rangle - \langle A_{ij} \rangle^2 = \langle f(K_i)^2 \rangle - \langle f(K_i) \rangle^2 + \frac{\sigma^2}{S} \quad (186)$$

$$\langle A_{ij} A_{ji} \rangle - \langle A_{ij} \rangle^2 = \frac{\sigma^2 \gamma}{S} \quad (187)$$

Now, let's set $f(K_i) = c(K_i - 1)$ so that

$$\langle f(K_i) \rangle = 0 \quad (188)$$

$$\langle f(K_i)^2 \rangle = c^2 (\langle K_i^2 \rangle - 2 \langle K_i \rangle + 1) \quad (189)$$

$$= c^2 \sigma_K^2 \quad (190)$$

Then

$$\langle A_{ij} \rangle = \frac{\mu}{S} \quad (191)$$

$$\langle K_i A_{ij} \rangle - \langle K_i \rangle \langle A_{ij} \rangle = c \langle K_i^2 \rangle - c \langle K_i \rangle = c \sigma_K^2 \quad (192)$$

$$\langle A_{ij}^2 \rangle - \langle A_{ij} \rangle^2 = c^2 \sigma_K^2 + \frac{\sigma^2}{S} \quad (193)$$

$$\frac{\langle A_{ij} A_{ji} \rangle - \langle A_{ij} \rangle^2}{\langle A_{ij}^2 \rangle - \langle A_{ij} \rangle^2} = \frac{\gamma}{1 + S c^2 \frac{\sigma_K^2}{\sigma^2}} \quad (194)$$

Hence, achieving a prescribed correlation $C = c \sigma_K^2$ requires to generate a matrix a_{ij} with parameters

$$\tilde{\sigma}^2 = \sigma^2 - S c^2 \sigma_K^2, \quad \tilde{\gamma} = \gamma(1 + S c^2 \sigma_K^2 / \sigma^2) \quad (195)$$

before inserting it into $A_{ij} = \frac{\mu}{S} + c(K_i - 1) + \sigma a_{ij}$.