Local equations for the generalized Lotka-Volterra model on sparse asymmetric graphs

D. Machado[1, 2, 3, ⋆], **P. Valigi**[1], **T. Tonolo**[4, 5] **and M. C. Angelini**[1, 6, †]

1 Dipartimento di Fisica, Sapienza Università di Roma, P.le Aldo Moro 5, 00185 Rome, Italy
2 Group of Complex Systems and Statistical Physics and Department of Theoretical Physics. Physics Faculty, University of Havana. CP10400, La Habana, Cuba
3 CNR - Nanotec, unità di Roma, P.le Aldo Moro 5, 00185 Rome, Italy **4** Gran Sasso Science Institute, Viale F. Crispi 7, 67100 L'Aquila, Italy
5 INFN-Laboratori Nazionali del Gran Sasso, Via G. Acitelli 22, 67100 Assergi (AQ), Italy
6 Istituto Nazionale di Fisica Nucleare, Sezione di Roma I, P.le A. Moro 5, 00185 Rome, Italy

⋆ david.machado@uniroma1.it , † mariachiara.angelini@uniroma1.it

Abstract

Real ecosystems are characterized by sparse and asymmetric interactions, posing a major challenge to theoretical analysis. We introduce a new method to study the generalized Lotka-Volterra model with stochastic dynamics on sparse graphs. By deriving local Fokker-Planck equations and employing a mean-field closure, we can efficiently compute stationary states for both symmetric and asymmetric interactions. We validate our approach by comparing the results with the direct integration of the dynamical equations and by reproducing known results and, for the first time, we map the phase diagram for sparse asymmetric networks. Our framework provides a versatile tool for exploring stability in realistic ecological communities and can be generalized to applications in different contexts, such as economics and evolutionary game theory.

Received Date
Accepted Date
Published Date

Contents

1 Introduction **2**

2 The model **4**

3 Local Fokker-Planck equations **5**
3.1 Individual Based Mean Field 6
3.2 Continuous Belief Propagation 8

4 Numerical results **9**
4.1 Undirected graphs with asymmetric interactions 12
4.2 Directed graphs 14
4.3 Including thermal noise 16

5 Conclusions **19**

6 Acknowledgments 19
A From Ito's rule to the Fokker-Planck equation 20
B Solution for the isolated variable 21
C Limits of IBMF with parallel updates at zero temperature 22
D Unbounded growth for asymmetric mutualistic interactions 23
E Use of damping to improve convergence 24
F IBMF for graphs with correlated couplings 25
G Convergence of IBMF in directed graphs 26
H Finite size effects of IBMF on directed graphs 28
I Comparing runtimes of IBMF and simulations in directed graphs 29
J Zero-temperature limit of BP 29
References 32

1 Introduction

The stability of complex ecosystems and the rules governing species coexistence present a central puzzle in theoretical ecology. The generalized Lotka-Volterra (gLV) model has been a cornerstone of this inquiry, also related to models used in evolutionary game theory and in economic theory [1–3]. Recently, the random, symmetric, and dense version of the gLV model has been analyzed through the lens of equilibrium statistical mechanics. This approach has yielded profound insights, revealing how the phases of single equilibrium, unbounded growth, and multiple equilibria are dictated by a handful of macroscopic parameters: the mean μ and variance σ of inter-species interactions, and the intensity of demographic noise T. A key finding is that densely connected networks with random symmetric interactions can exhibit a phase of multiple equilibria at high interaction heterogeneity [4].

However, real ecological networks are generally neither dense nor symmetric. In real ecosystems, the interactions between two species are almost always asymmetric, the existence of predator-prey couples of species being just an example. For this reason, some other works have tried to go beyond the assumption of symmetry [5–14]. Moreover, in real ecosystems, a species typically interacts only with a few others [15, 16]. In the last years, Random Matrix Theory (RMT) results [17, 18] have shown that the spectra of sparse random graphs exhibit qualitative differences from the dense case, suggesting that this may have implications for the stability of ecological models defined on sparse networks. Accordingly, it would be interesting to study the properties of the gLV model on graphs with finite connectivity. Unfortunately, in this case, the methods used for fully-connected systems [4, 7] cannot be applied, in particular because no central-limit-type arguments hold.

In a very recent paper, for the first time, the equilibrium properties of the symmetric gLV model on a sparse graph were analyzed [19] using the so-called Belief-Propagation (BP) cavity

method. This can, however, be used only when an equilibrium measure exists, preventing its implementation in the case of asymmetric interactions. In general, the description of out-of-equilibrium systems has been even more elusive, and the available techniques deal only with specific limits. When species interact through a fully-connected network [20], we can use Dynamical Mean-Field Theory (DMFT) to describe the temporal evolution of dynamic observables, which has also been extended to non-Gaussian disorder in the interactions [21]. Other generalizations include the limit of very small connectivity [21], the limit of a large connectivity that grows sublinearly with the number of species in the graph [22], and the case of unidirectional interactions [23]. However, until now, there has been no way to systematically analyze gLV models with sparse asymmetric interactions.

In this work, we bridge this gap by analyzing both symmetric and asymmetric, quenched-disordered sparse interactions. We go beyond equilibrium statistical mechanics, introducing a new method for the evaluation of the stationary probability distribution for the stochastic differential equations (SDE) that describe the evolution of the species abundance. We start from the usual formulation of the stochastic dynamics in terms of an SDE and derive the equivalent Fokker-Planck equation for the associated time-dependent probability densities. However, solving the full system of partial differential equations defined on the whole graph in a high-dimensional space is a cumbersome task. To overcome this difficulty, we derive local closures that allow us to obtain tractable relations for the stationary distributions.

The main idea of a local closure is to propose an *ansatz* for the probability densities, in general involving some suitable factorizations that exploit the properties of the interaction graph. It has been used successfully in several contexts, such as the study of epidemics spreading on networks [24–27], algorithmic dynamics in hard combinatorial optimization problems [28, 29], spin-glass dynamics in random graphs [30], or the dynamics of the voter model [31]. As far as we know, this work constitutes the first application of local closures to Fokker-Planck equations in sparse graphs. The new approximate descriptions derived here are what we call local Fokker-Planck equations.

We will validate our method by applying it to different situations, comparing its prediction with the results obtained from the direct numerical integration of the SDE and recovering known results from previous literature. We also show how, starting from the general equations in the asymmetric case, one can recover the BP equations when only symmetric couplings are considered.

The rest of the manuscript is organized as follows. In Section 2 we introduce the generalized Lotka-Volterra equations and the underlying networks that we will analyze. In Section 3, we present the Fokker-Planck equations that describe the evolution of the probability distributions of the species abundances in time. These are complicated global equations, which in general are not solvable. For this reason, in Section 3.1 we introduce a new local closure, corresponding to a mean-field approximation for the dynamics, called Individual Based Mean Field (IBMF). Our procedure leads to the main *local* solvable Fokker-Planck equations that we discuss in this article. In Section 3.2 we go beyond IBMF, introducing a more refined closed local Fokker-Planck equation that we call Pair Based Mean Field (PBMF), showing that the BP equations introduced in Ref. [19] correspond to the stationary solution of PBMF in the symmetric case.

We validate our new methods in Section 4, where we give details on the numerical implementation of IBMF both at null and finite temperatures. We compare the stationary abundances obtained with IBMF with those obtained from simulations in the presence of thermal noise for a single random graph. The results are a useful example to emphasize the strengths of IBMF, and to also point out its limitations. Then, we proceed to apply IBMF to three different scenarios. In Subsection 4.1, we study undirected graphs with asymmetric interactions at null temperature, obtaining the corresponding phase diagram in the plane (μ, σ) for the first

time, as far as we know. We thus generalize the results for the fully-connected asymmetric case in Ref. [5] and for the sparse case with symmetric interactions in Ref. [19]. In Subsection 4.2, we apply the IBMF closure to directed graphs with null variance in the couplings and null temperature, confirming and extending the results of Ref. [32], which were obtained there with a completely different method. In Subsection 4.3 we move back to undirected graphs, but this time with symmetric couplings at finite T. This setting helps us study the performance of IBMF in the presence of thermal noise in a systematic way. We compare the results with Ref. [19], where the BP method is used to exactly solve the model. We identify the limitations of IBMF, which stops converging to a single equilibrium as soon as the exact species abundances found by BP start developing a non-Gaussian distribution tilted towards extinctions. Finally, in Section 5, we draw our conclusions.

2 The model

Let us introduce the generalized Lotka-Volterra (gLV) equations that we will study in the rest of this article. They describe the dynamics of an ecosystem with N interacting species. To each of them, we associate a positive real variable n_i, interpreted as the abundance of the i-th species, with $i = 1, \dots, N$. In general, a single species will not interact with all the others, but instead with a subset of the species known as the neighborhood of i. The interactions occur in a graph $G(V, E)$, where V is the set of vertices, each representing a species, and E is the set of edges.

To keep the definitions as general as needed, for now the reader should think of G as a directed graph. If the presence of species i influences the growth of species j, we add the directed edge $i \to j$. It is possible to have $i \to j$ in the graph without having the edge in the opposite direction ($j \to i$). We define the in-neighborhood ∂i^- of i as the set of in-neighbors j such that the edge $j \to i$ exists in the graph. For simplicity, graphs G with self-loops will not be considered here

The gLV equation for the abundance of the i-th species can be written as:

$$\frac{dn_i}{dt} = \frac{r_i}{K_i} n_i (K_i - n_i - \sum_{j \in \partial i^-} \alpha_{ij} n_j) + \xi_i(t) + \lambda \,, \tag{1}$$

where $n_i \geq 0$ is the abundance of the i-th species, and the real parameters r_i and K_i are known as the intrinsic growth rate and carrying capacity, respectively. To simplify the setting, we will take $r_i = K_i = 1$ in what follows, but the reader will find no difficulties in generalizing our results to consider other values of these constants.

The term $\xi_i(t)$ in Eq. (1) is a noise term, which has average $\langle \xi_i(t) \rangle = 0$ and second moments $\langle \xi_i(t_1)\, \xi_j(t_2) \rangle = 2\, T\, n_i\, \delta_{i,j}\, \delta(t_1 - t_2)$, where T is known as temperature of the noise. This thermal noise is referred to as *demographic* [4, 33–36] and accounts for death and birth processes. The parameter λ, known as *immigration rate*, acts as a small source term that allows extinct species to come back should conditions become favorable to them [20, 37]. Its effect will be clarified later.

The couplings α_{ij} are real numbers that set the type and strength of the interactions. The value of α_{ij} encodes the way that species j affects the evolution of species i, and therefore corresponds to the edge $j \to i$ on the graph. In the case where for all $j \in \partial i^-$ we also have the edge $i \to j$, the graph is known as undirected. Having α_{ij} and α_{ji} simultaneously positive means that the two species i and j have a competitive interaction, where the presence of individuals of species j is prejudicial for the individuals of species i, and vice versa. When they are both negative, we have a mutualistic interaction, and the species are beneficial to each other. On the other hand, when the interaction is positive for one species and negative

for the other, we have a predator-prey or antagonistic interaction. Finally, in the presence of directed interactions, we can also have commensalism and amensalism, whereby one species benefits or is harmed by the interaction, while the other is unaffected.

Given that the edge $j \to i$ already exists in the graph, we can add noise to the interactions by drawing α_{ij} at random from some probability distribution. Following many other works [19, 20, 38], we choose the Gaussian distribution $\alpha_{ij} \sim \mathcal{N}(\mu, \sigma)$, with mean μ and variance σ^2. The reader should note that, as a particular case, we can set the interaction strengths α_{ij} to be homogeneous by choosing $\sigma = 0$. In that case, we get $\alpha_{ij} = \mu$ for all the edges $j \to i$.

Our methodology, derived below in Section 3, applies to graphs with directed and/or undirected interactions. We demonstrate this by including results for three different scenarios in Section 4. In Subsection 4.2 we study a case where, with high probability, the edges $j \to i$ and $i \to j$ are not simultaneously present. We follow the same model used in Ref. [39] to study the gLV dynamics with asymmetric interactions. To construct the network, for each species i we select the incoming edges by going over all possible $j \neq i$, and adding the edge $j \neq i$ with probability c/N, where $c > 0$. The in-neighbors of i are chosen independently of the in-neighbors of j, and in the limit when the number of species is large, it is highly improbable that we find $i \to j$ and $j \to i$ simultaneously in the graph. The result is a graph where the degree follows a Poisson distribution with mean c, and where most interactions are directed. Finally, for each edge $j \to i$ we draw α_{ij} from the Gaussian distribution $\mathcal{N}(\mu, \sigma)$. In Subsection 4.2 below, we include results for different values of μ but only two values of σ ($\sigma = 0$ and $\sigma = 0.15$).

In Fig. 1, and in Subsections 4.1 and 4.3, the species interact over an undirected random regular graph, whose edges are randomly selected such that each species (vertex) has the same number of neighbors, denoted by c and called connectivity. After the graph is built, we need to choose the interaction strengths α_{ij} and α_{ji} for each edge. In Subsection 4.1, we draw α_{ij} independently of α_{ji} using the Gaussian distribution $\mathcal{N}(\mu, \sigma)$, for different values of μ and σ. Notice that this creates asymmetric interactions where, in general, we have $\alpha_{ij} \neq \alpha_{ji}$ whenever $\sigma \neq 0$. In Subsection 4.3 we study the case $\sigma = 0$, where we always get $\alpha_{ij} = \alpha_{ji} = \mu$ and the interactions are symmetric.

3 Local Fokker-Planck equations

Given a graph of interactions, Eq. (1) gives the temporal evolution of the abundances in a stochastic process with thermal noise $\xi(t)$. Sampling different realizations of the initial conditions and of $\xi(t)$, one gets the probability distribution $P_t(\vec{n})$ of the vector $\vec{n} = (n_1, \ldots, n_N)$ at time t. This quantity obeys a Fokker-Planck equation that can be derived from Ito's rule [38]. However, contrary to the usual case where the variables are defined in the interval $(-\infty, +\infty)$, each abundance n_i is defined in the interval $[0, +\infty)$. As a natural consequence of this fact (see Appendix A for more details), one needs to impose the proper boundary conditions at $n_i = 0$, guaranteeing that the current of probability density through the border is always zero [40]. The resulting Fokker-Planck equation is:

$$\frac{\partial P_t(\vec{n})}{\partial t} = T \sum_{i=1}^{N} \frac{\partial^2}{\partial n_i^2} \{ n_i P_t(\vec{n}) \} - \sum_{i=1}^{N} \frac{\partial}{\partial n_i} \Big\{ \Big[n_i (1 - n_i - \sum_{j \in \partial i^-} \alpha_{ij} n_j) + \lambda \Big] P_t(\vec{n}) \Big\}, \quad (2)$$

The first and second terms on the right-hand side of Eq. (2) are the usual diffusion and drift terms of the Fokker-Planck equation, respectively. They encode the evolution of a species subject to Eq. (1). The deterministic growth ratio $n_i(1 - n_i - \sum_{j \in \partial i^-} \alpha_{ij} n_j) + \lambda$ experienced by species i goes into the drift term. The thermal noise with temperature T gives birth to the diffusion term.

In any case, solving Eq. (2) is a cumbersome task mainly because $P_t(\vec{n})$ is a highly dimensional object. The abundances are defined on the space $(0,+\infty)^N$ and the time can be in general defined in the space $(-\infty,\infty)$. Even when we consider a single species ($N=1$), finding $P_t(n)$ at any time t is not simple. However, we can obtain its stationary solution (see Appendix B), which will be useful for us later. It reads:

$$P_\infty(n) = \frac{1}{Z}\, n^{\beta\lambda-1}\, \exp\Big\{-\frac{\beta}{2}(n-1)^2\Big\}, \tag{3}$$

where Z is a normalization constant and $\beta \equiv 1/T$.

Eq. (3) clarifies the role of the parameter λ in the model. The integral $\int_0^\infty dn\, P_\infty(n)$ is finite if and only if $\lambda > 0$. Otherwise, the divergence at $n=0$ dominates the integral, which would be divergent. In other words, the existence of $\lambda > 0$ allows the density $P_\infty(n)$ to be normalizable. On the other hand, when $\lambda = 0$ and $T > 0$ the species are doomed to go extinct for large times.

3.1 Individual Based Mean Field

To find solvable equations, we need to simplify Eq. (2). In this subsection, we obtain the first local Fokker-Planck equation for the gLV model. Let us marginalize Eq. (2) over all the abundances except n_i to obtain the differential equations for the local probabilities, which are defined as $P_t(n_i) = \int_0^\infty [\prod_{k\neq i} dn_k] P_t(\vec{n})$:

$$\frac{\partial P_t(n_i)}{\partial t} = T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i)\big\} - \frac{\partial}{\partial n_i}\Big\{\Big[n_i\Big(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\, m_{j\to i}(n_i,t)\Big)+\lambda\Big]P_t(n_i)\Big\}, \tag{4}$$

where $m_{j\to i}(n_i)$ is the conditional average

$$m_{j\to i}(n_i,t) \equiv \int_0^\infty dn_j\, n_j P_t(n_j \mid n_i). \tag{5}$$

As in the Fokker-Planck equation for the whole system, the local version in Eq. (4) has two different contributions. The first line in the equation shows the diffusion term. On the other hand, after averaging over the rest of species, the single species i senses an effective drift $n_i\big(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\, m_{j\to i}(n_i,t)\big)+\lambda$, where n_j is substituted by its conditional average $m_{j\to i}(n_i,t)$. For more details on the derivation of Eq. (4), the reader is referred to Section 1 of the Supplemental Materials (SM).

We have not introduced any approximation so far. To solve Eq. (4), one would also need to obtain all the functions $m_{j\to i}(n_i)$, but from its definition (Eq. (5)) it is evident that this is equivalent to getting the solution for the pair probabilities $P_t(n_i,n_j)$. Indeed, to compute the conditional probability density $P_t(n_j \mid n_i)$, we need the pair $P_t(n_i,n_j)$ and the single-site $P_t(n_i)$ probabilities. As we will show in the next section, the local Fokker-Planck equation for $P_t(n_i,n_j)$ depends, in turn, on probabilities $P_t(n_i,n_j,n_k)$ defined over three species. After iterating this process, we get a hierarchy of equations that never closes until we recover the full Eq. (2). Therefore, solving Eq. (4) has the same level of difficulty as solving Eq. (2).

To overcome this problem, we need to introduce an approximation that allows us to get a *closed* system of differential equations for the $P_t(n_i)$, *i.e*, one that can be solved without going up in the hierarchy. The first step that one could take in that direction is to assume $m_{j\to i}(n_i,t)$ is independent of n_i and write $m_{j\to i}(n_i,t) \approx m_j(t)$, where

$$m_j(t) \equiv \int_0^\infty dn_j\, n_j P_t(n_j). \tag{6}$$

Doing this is equivalent to assuming that the pair probabilities are all factorized such that $P_t(n_i, n_j) \approx P_t(n_i)P_t(n_j)$. Therefore, we are trivializing the correlations in the system. Nevertheless, this approximation allows us to close the system of differential equations since all the information that we need is in the single-species distributions $P_t(n_i)$. We get:

$$\frac{\partial P_t(n_i)}{\partial t} = T\frac{\partial^2\{n_i P_t(n_i)\}}{\partial n_i^2} - \frac{\partial}{\partial n_i}\Big\{\Big[n_i\Big(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(t)\Big)+\lambda\Big]P_t(n_i)\Big\}. \tag{7}$$

These local Fokker-Planck equations form a dynamic closure that can in principle be solved, and that we will call Individual Based Mean Field (IBMF) in what follows. This name has been used before in the literature, particularly in the study of epidemic spreading throughout a network [24,26], to identify an approximation that factorizes the pair probabilities distribution as explained above. Finding the stationary solution of Eq. (7) has the same level of difficulty as for an isolated variable (see Eq. (3) and Appendix B). The result is:

$$P_\infty(n_i) = \frac{1}{Z_i}\,n_i^{\beta\lambda-1}\,\exp\Big\{-\frac{\beta}{2}(n_i-M_i)^2\Big\}, \tag{8}$$

where

$$Z_i = \int_0^\infty dn_i\,n_i^{\beta\lambda-1}\,\exp\Big\{-\frac{\beta}{2}(n_i-M_i)^2\Big\} \tag{9}$$

$$M_i = 1-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(\infty)\ . \tag{10}$$

This is a mean-field solution to the problem derived for sparse graphs with any type of interactions. In fact, a similar probability density has been recently introduced in Ref. [41] for the case of fully-connected models, a scenario where mean-field assumptions like these are more commonly used. The presence of other species modifies the center M_i of the Gaussian in Eq. (8) to make n_i align with the average effect of its neighbors. To evaluate the stationary solution of IBMF, we need to design an algorithm capable of computing the averages $m_j(\infty)$, which we will denote by m_j for simplicity. Exploiting Eq. (8), we can write:

$$m_i = \frac{1}{Z_i}\int_0^\infty dn_i\,n_i^{\beta\lambda}\,\exp\Big\{-\frac{\beta}{2}(n_i-M_i)^2\Big\}. \tag{11}$$

As said above, Z_i and M_i are functions of the averages m_j, with $j\in\partial i^-$. After making an initial guess for the average abundances m_i, with $i=1,\dots,N$, we can use Eq. (11) to update their values. Then, we iterate until all m_i converge to the IBMF's prediction for the stationary average abundances. In practice, we employ numerical tricks, such as adding damping and using sequential updates, to aid this iterative process in reaching convergence (see Appendices E and C).

In the process of obtaining the stationary solution of IBMF, we did not assume any particular structure of the graph, nor any specific type of interaction. It is in principle applicable to any directed or undirected graph, with symmetric or asymmetric interactions. However, one should expect better results when correlations are weak enough for the factorization $P_\infty(n_i, n_j) \approx P_\infty(n_i)P_\infty(n_j)$ to approximately hold. The latter is intuitively more likely to happen in the case of asymmetric interactions.

3.2 Continuous Belief Propagation

To go beyond IBMF, we need to include non-trivial correlations between pairs of interacting species. Therefore, it is reasonable to go up one level in the hierarchy and write the local Fokker-Planck equations for the pair probabilities $P_t(n_i, n_j) = \int_0^\infty [\prod_{k\neq i,j} dn_k] P_t(\vec{n})$. After marginalizing Eq. (2) over the abundances of all the species (the details can be found in Section 1 of the SM), except for i and j, we obtain a local equation that depends on the conditional averages:

$$m_{k\to i,j}(n_i, n_j, t) \equiv \int_0^\infty dn_k\, n_k\, P_t(n_k \mid n_i, n_j). \tag{12}$$

Solving for $m_{k\to i,j}(n_i, n_j, t)$ implies knowing the three-species probabilities $P_t(n_i, n_j, n_k)$, but we need to close the hierarchy at some point. However, if one focuses on random sparse graphs like random regular or Erdős-Rényi, the tree-like structure of the interactions makes it very unlikely that species k interacts with i and j simultaneously, provided that i and j interact. If $k \in \partial i^-$, we can assume that $m_{k\to i,j}(n_i, n_j, t)$ strongly depends on n_i and only weakly depends on n_j. We then make the approximation $m_{k\to i,j}(n_i, n_j, t) \approx m_{k\to i}(n_i, t)$. The result (see Section 1 of the SM) is a closed local Fokker-Planck equation that we call Pair Based Mean Field (PBMF), also using a name that is popular in the context of epidemics spreading on networks when the dynamics is described using pair probabilities [25, 26].

Although the PBMF is simpler than the full Eq. (2), it is still difficult to obtain a general solution, even when we focus only on the stationary point. Only in the case of symmetric interactions in undirected graphs ($\alpha_{ij} = \alpha_{ji}$), as we show in Subsection 2.1 of the SM, the proper solution to the local Fokker-Planck equation for pair probabilities in sparse random graphs is Belief Propagation (BP). This technique was already introduced in Ref. [19] for the gLV model. However, while in that case the abundances of the species are considered as discrete variables with states $n_i = 1, 2, \ldots$, here we use the continuous version of the model. We solve this issue by proposing a new continuous implementation of BP equations for our model, which have the same structure as the ones in Ref. [19], but include the proper adjustments to consider continuous $n_i \in [0, +\infty)$. These are:

$$\eta_{i\to j}(n_i) = \frac{n_i^{\beta\lambda-1}}{z_{i\to j}} \exp\Big\{-\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \prod_{k\in\partial i^-\setminus j} \int_0^\infty dn_k\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i n_k}\ , \tag{13}$$

where $\eta_{i\to j}(n_i)$ is the cavity marginal, or message, that represents the marginal probability density of species i in a modified graph where the edge connecting i and j is removed. The constant $z_{i\to j}$ is a normalization factor.

Here, we can identify that the local field h_i, whose exponential $e^{-\beta h_i(n_i)}$ usually appears in front of BP equations, is simply $h_i = n_i^2/2 - n_i + (T-\lambda)\ln(n_i)$. This expression properly considers the immigration rate and the continuous nature of n_i. From it, we can obtain the stationary single-site and the pair probabilities as follows:

$$P_{BP}(n_i) = \frac{n_i^{\beta\lambda-1}}{Z_i} \exp\Big\{-\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \prod_{k\in\partial i^-} \int_0^\infty dn_k\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i n_k} \tag{14}$$

$$P_{BP}(n_i, n_j) = \frac{1}{Z_{ij}} \eta_{i\to j}(n_i)\, e^{-\beta\alpha_{ij} n_i n_j}\, \eta_{j\to i}(n_j)\,, \tag{15}$$

where Z_i and Z_{ij} are normalization factors and the messages $\eta_{i\to j}(n_i)$ are the fixed point solution of Eq. (13).

The reader could wonder what the relation is between BP and the local Fokker-Planck equations that we have been presenting here. Remarkably, it is possible to prove that the expression in Eq. (15) for the pair probabilities, together with BP equations (Eq. (13)), is a stationary solution of PBMF when the interactions are symmetric. The details of the proof are given in Subsection 2.1 of the SM. To obtain numerical results from BP, we use an iterative algorithm analogous to the one we introduced above for IBMF. Making an initial guess for the messages $\eta_{i\to j}(n_i)$, we can use Eq. (13) to update their values at each n_i. This procedure is iterated until all $\eta_{i\to j}(n_i)$ converge, and the final messages are used to compute the true marginals $P_{BP}(n_i)$ and $P_{BP}(n_i, n_j)$. For the interested reader, we include details about our specific implementation of BP in Section 4 of the SM.

4 Numerical results

As said above, to use IBMF to obtain the actual values of the averages m_i, we need to numerically compute integrals of the form:

$$I_k(\beta,\lambda,M) = \int_0^\infty dn\, n^{\beta\lambda-1+k}\, \exp\Big\{-\frac{\beta}{2}(n-M)^2\Big\}\,, \tag{16}$$

with the parameter k taking the value $k=0$ in Eq. (9), and $k=1$ in Eq. (11).

Luckily, the integral in Eq. (16) can be expressed in terms of known special functions, called parabolic cylinder functions (see 9.241 in Ref. [42]). This is very convenient because we can write the parabolic cylinder functions in terms of the more practical Kummer's confluent hypergeometric function, which can be found already tabulated in different programming languages. The interested reader can find the details in Section 3 of the SM. The code is available at Ref. [43].

When the temperature is zero or close to zero, the equations can be simplified even further. The exponential in Eq. (16) concentrates around its maximum, and provided that λ is small, IBMF equations reduce to:

$$n_i = \max\Big\{0,\, 1-\sum_{j\in\partial i^-}\alpha_{ij}\, n_j\Big\}. \tag{17}$$

At $T=0$, we simply need to iterate Eq. (17) until convergence. It is important to note that this is not the same as running the zero-temperature simulations of the dynamics, which in turn implies integrating the system of differential equations:

$$\frac{dn_i}{dt} = n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}n_j)+\lambda\ . \tag{18}$$

Such numerical integration of the gLV dynamical equations at $T=0$ can be performed using the Cash-Karp adaptive Runge-Kutta method [44]. The code is available at Ref. [43].

For small λ, the fixed points of Eq. (17) coincide with the stationary solutions of the exact dynamics in Eq. (18). Therefore, whenever IBMF converges, the fixed point represents an actual stationary configuration of the dynamics. On the other hand, the non-convergence of IBMF is not guaranteed to be reflected in the behavior of the simulated dynamics. One could think of them as two different algorithms trying to find the same fixed points. If one of the algorithms succeeds, the resulting abundances constitute a fixed point also for the other algorithm. If one of them does not succeed, the other still could.

However, we show in Subsections 4.1 and 4.2 (see below) that IBMF can nevertheless be used to predict the relevant phase transitions observed in the simulations at $T=0$. In

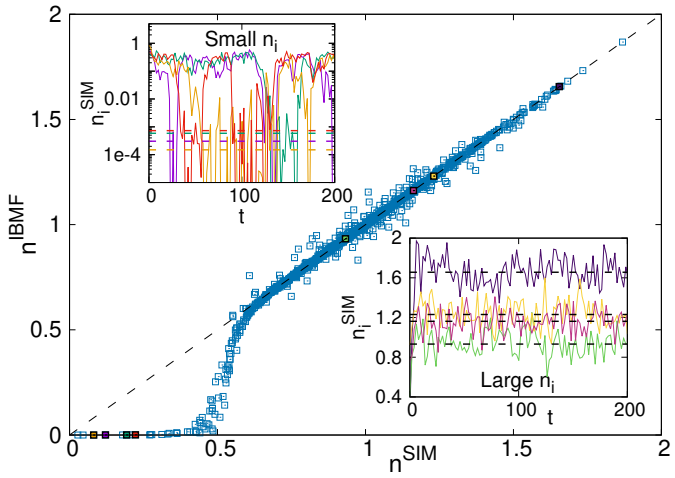

Figure 1: Comparing individual abundances from IBMF and simulations in a random regular graph at finite temperature $T = 0.015$. The connectivity is $c = 3$, and the immigration rate is $\lambda = 10^{-6}$. Each α_{ij} is independently drawn from the Gaussian $\mathcal{N}(0,\sigma)$, with $\sigma = 0.15$ (the interactions are asymmetric). Each point in the main graphic has coordinates $(n_i^{\rm SIM}, n_i^{\rm IBMF})$, where $n_i^{\rm SIM}$ is the average stationary abundance of species i obtained from 100 simulations of the dynamics, and $n_i^{\rm IBMF}$ is the prediction of IBMF for the same species. The black dashed line is just the linear function $f(x) = x$. The system has $N = 1024$ species, thus there are 1024 points in the main graphic. The inserted graphic in the top-left corner shows the temporal evolution of four species with small stationary abundances. The corresponding points are marked with the same colors in the main graphic. The horizontal lines are the predictions of IBMF for the same species. The graphic is in semi-log scale. Analogously, the inserted graphic in the bottom-right is done with four species whose abundances are not small. Colored lines show the results of simulations, and the horizontal black lines show the predictions made with IBMF.

Appendix C and for a random regular graph with homogeneous interactions, we show that the exact result for the transition from single-to-multiple equilibria [32] arises naturally from IBMF. Furthermore, in Section 2.2 of the SM, we recover the stationary solution of Dynamical Mean Field Theory [20] in the limit of large connectivity at zero temperature.

In the presence of thermal noise ($T > 0$), we still have a fast implementation of IBMF. It is important to note that, for finite temperatures, IBMF is always a factorized ansatz for the probability density of the abundances. With its stationary solution, we can predict the final average abundance $n_i^{\rm IBMF}$ for each of the species in a given graph. To illustrate how this works, Fig. 1 compares each $n_i^{\rm IBMF}$ with the average stationary abundances $n_i^{\rm SIM}$ obtained from simulations for a specific realisation of a random regular graph in the presence of thermal noise. The numerical integration of the SDE defining the gLV model in Eq. (1) are performed by means of the Milstein method [45, 46]. As with the other algorithms, the code is provided in Ref. [43]. The interactions are asymmetric, *i.e.*, we choose α_{ij} independently of α_{ji}, each from a Gaussian distribution $\mathcal{N}(0,\sigma)$. As far as we know, this is the first time a theoretical prediction of this kind has been made for sparse graphs.

The main graphic of Fig. 1 shows that IBMF accurately predicts the abundances of the species that are dominant in the ecosystem. For $n_i > 0.6$, the points $(n_i^{\rm SIM}, n_i^{\rm IBMF})$ lie around the line $f(x) = x$ of perfect agreement. Most species ($\sim 91\%$) are in this group. In the bottom-right corner of Fig. 1, the inserted graphic shows the temporal evolution of four of those species observed in a single simulation. The corresponding points in the main graphic

are marked using the same colors. The stationary abundances of those species, which were selected at random, oscillate around the corresponding predictions from IBMF (presented in black dashed lines). We observe almost perfect agreement between them.

On the other hand, IBMF consistently underestimates the stationary abundances obtained from simulations for species that are closer to extinction, with $n_i < 0.6$. Very few ($\sim 9\%$) species are in this group. The results become clearer after analyzing the inserted graphic in the top-left of Fig. 1. There, we show the temporal evolution of the abundances of four species that are very close to extinction according to IBMF, but whose average stationary abundance from simulations is not as small. The corresponding points in the main graphic are marked using the same colors.

The inserted graphic shows that these species in the lower bottom corner of Fig. 1 continuously switch between two time-persistent states. After spending some time oscillating around a value of the abundance n_i^{high} that is not small (from the figure we see that $n_i^{\text{high}} > 0.1$), the species suddenly drop down and oscillate for a while around a very small abundance $n_i^{\text{low}} \sim 0.001$. This small n_i^{low} corresponds well to the predictions of IBMF, marked with horizontal dashed lines in the graphic. The real stationary abundance measured in simulations by averaging n_i for long times, however, is somewhere in between n_i^{low} and n_i^{high}. Instead of mimicking this intermediate value without a clear physical meaning, IBMF gives only the smallest of the two true values n_i^{low} and n_i^{high}. Although it is only partially right, it definitely allows identifying the species that are going to exhibit this type of dynamics. A similar behavior has been recently found in fully-connected systems with asymmetric interactions and without thermal noise [47], where the species also switch between two time-persistent states, only one of which is close to extinction. Remarkably, while in Ref. [47] this is observed for every species, in our case we have only a few switching species. This species heterogeneity is probably related to the sparsity of interactions.

Both with IBMF and with simulations, we verified that the corresponding stationary abundances were independent of the initial conditions (and of the realization of the noise in simulations). Remarkably, computing the stationary abundances with IBMF is two orders of magnitude faster than running the simulations. By averaging over 100 different initial conditions, we get the average wall-clock times of 22.0 ± 0.3 ms for IBMF, and of 1390 ± 60 ms for simulations (ms stands for milliseconds).

The results in Fig. 1 clarify the meaning of IBMF and its predictions for a single graph, while also raising new questions on the link between the structure of the interaction graph and the observed non-trivial dynamics. This phenomenon can also depend on the model's parameters. Getting a clear picture will require further work, especially because of the difficulties involved in numerically analyzing the results from simulations in the presence of thermal noise.

In this introductory work, we provide a general and clearer picture of how IBMF works in more controlled scenarios. First, Subsections 4.1 and 4.2 compare IBMF with simulations at zero temperature, where the results from the latter are easier to interpret. In Subsection 4.1, we study the phase diagram of the model in graphs with asymmetric interactions in undirected random regular graphs. In Subsection 4.2, we revisit a model discussed in Ref. [39] to predict the probability of observing persistent fluctuations in the dynamics for any given system size. Subsection 4.3 is devoted, instead, to a case where we include thermal noise. Although IBMF with $T > 0$ can be applied to symmetric or asymmetric interactions, we chose to study the model with symmetric and homogeneous interactions. The reason is that, in this case, we can compare the output of IBMF with the results of BP, thereby avoiding the numerical complications associated with studying the phase transitions of the simulated dynamics in the presence of thermal noise. The latter is left for future work.

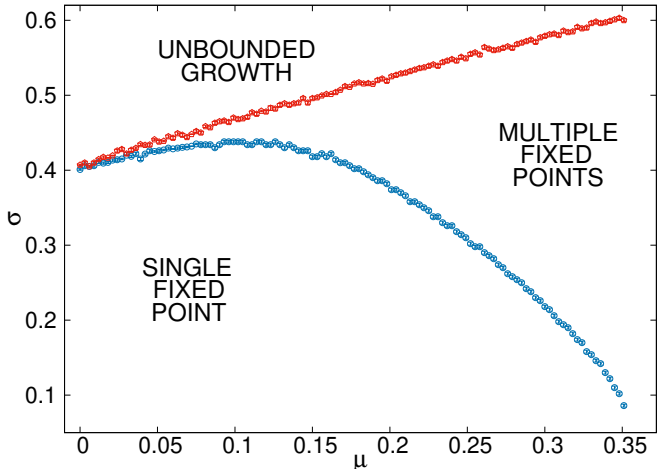

Figure 2: Transitions obtained simulating the gLV model for $T = 0$, asymmetric interactions (α_{ij} is chosen independently of α_{ji}), and $\lambda = 10^{-6}$. For several pairs (μ, σ), we run the dynamics for 10000 different random regular graphs with connectivity $c = 3$ and size $N = 1024$. The interaction strengths are drawn from the Gaussian distribution: $\alpha_{ij} \sim \mathcal{N}(\mu, \sigma)$. By repeating the simulation 10 times with different initial conditions for each graph, we identify one of three possible outcomes: i) all realizations converge to the same fixed point, ii) all the realizations converge but the fixed points are different, or iii) the abundances in at least one of the simulations grow and diverge for long times. The blue points mark, for each μ, the maximum value of σ at which more than 50% of the samples are of type i). The red points mark, for each μ, the minimum value of σ at which more than 50% of the samples are of type iii).

4.1 Undirected graphs with asymmetric interactions

In this Subsection, we apply IBMF to undirected graphs with Gaussian noise in the interactions at zero temperature. We take a random regular graph with a given connectivity c, and draw every α_{ij} from a Gaussian with mean μ and standard deviation σ ($\alpha_{ij} \sim \mathcal{N}(\mu, \sigma)$). This means that the coupling in the opposite direction, α_{ji}, is independently drawn from the same distribution. Thus, the interactions are generally asymmetric. The larger the standard deviation σ, the bigger the average difference between α_{ij} and α_{ji}.

Fig. 2 shows the phase diagram obtained by simulating the gLV dynamics at $T = 0$ and $\mu > 0$, with $\lambda = 10^{-6}$ (see Eq. (18)). We identify three distinct regions. When σ is small enough, the species reach a unique fixed point for long times, which corresponds to the single-fixed-point (SFP) phase. The first transition occurs at $\sigma_{SFP}(\mu)$, and is represented with blue points in Fig. 2. For $\sigma > \sigma_{SFP}(\mu)$, simulations with different initial conditions will not converge to the same fixed point in most interaction graphs. The unbounded growth (UG) transition is located at $\sigma_{UG}(\mu) \geq \sigma_{SFP}(\mu)$, and is represented by the red points in Fig. 2. Above this line, the abundance of at least one species grows and diverges in most simulations.

Note that the transition at $\sigma_{SFP}(\mu)$ is not purely between a single-fixed-point phase and a multiple-fixed-points phase. Although for large μ this is indeed the case, for $\mu \leq 0$ the system goes directly from reaching a single fixed point to showing unbounded growth (see Appendix D). In between, we have a crossover between these two types of transitions.

The crossover poses a problem in predicting them using IBMF. As said before, any fixed point of IBMF is also a stationary solution of the exact dynamics at $T = 0$. We need to design a procedure capable of detecting the presence of different fixed points if they exist. We then

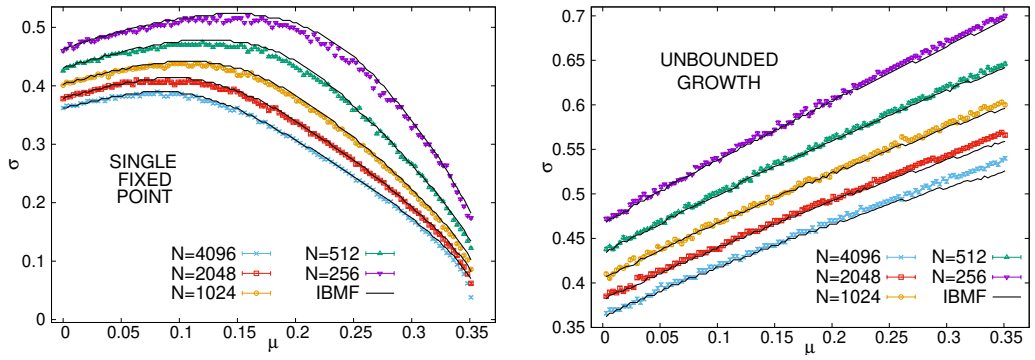

(a) (b)

Figure 3: Transitions of the gLV model for different system sizes at $T = 0$. The interactions are asymmetric (α_{ij} is chosen independently of α_{ji}) and defined on random regular graphs with connectivity $c = 3$. Points represent the results of simulations with immigration rate $\lambda = 10^{-6}$, and lines are the predictions made with IBMF for the same sizes. Each transition was determined using 10000 graphs. Simulations are repeated for 10 different initial conditions. IBMF was run with damping (see Appendix E) for 10 different random initial conditions. **(a)** For each system size N and average strength μ, points (lines) mark the maximum value of σ such that simulations (IBMF) converged to the same fixed point in more than 50% of the interaction graphs. **(b)** Points (lines) mark the minimum value of σ such that simulations (IBMF) displayed unbounded growth (not converged) more than 50% of the interaction graphs.

compare the results when we choose different initial conditions for the average abundances of IBMF.

In Fig. 3, we compare the results of the simulations with the predictions of IBMF. There are three possible outcomes of IBMF, and they are similar to those obtained from the simulations. For σ small enough, running IBMF in a specific graph many times with different initial conditions always gives the same fixed point. For σ large enough, two things can happen. Either IBMF converges to multiple fixed points for a given graph, provided that we change the initial conditions, or it does not converge at all. Interestingly, for IBMF we also observe a crossover between these behaviors. Fig. 3a shows the transitions between the single-fixed-point phase and the region where we either find multiple fixed points or no convergence. The lines, representing IBMF, accurately reproduce the results from the simulations.

Although we know that if IBMF reaches a fixed point, this is also a fixed point of the exact dynamics, it does not necessarily stop converging when the simulations do. We used IBMF to reproduce the transition to unbounded growth. Fig. 3b shows a very good agreement between simulations and IBMF. This also extends to negative values of μ, which correspond to interactions that are mutualistic on average (see Appendix D).

Our results indicate that IBMF is enough to independently describe both aspects of the crossover: the transition to the multiple-fixed-points phase and the transition to unbounded growth. Note that the finite-size effects are relevant in both panels of Fig. 3. The transitions obtained with the simulations and with the theory move downward when the number of species N increases. Nevertheless, IBMF is enough to capture these effects correctly, and its description is already accurate for finite systems.

As mentioned above, to obtain Fig. 3 we drew α_{ij} and α_{ji} independently from Gaussian distributions for every pair of interacting species. It is important to mention, however, that the quantitative agreement between IBMF and simulations also holds when α_{ij} and α_{ji} are correlated, as we show in Appendix F.

4.2 Directed graphs

As in Ref. [39], we study the emergence of fluctuating abundances n_i in graphs where the degree follows a Poisson distribution with mean c, and where most interactions are directed. If the edge $j \to i$ is present, with high probability the edge in the opposite direction does not exist. The graph can be seen as representing a directed flow between the species. We say that species j is *upstream* with respect to species i ($j \to i$). Conversely, species i is said to be *downstream* with respect to species j. We then independently draw each associated α_{ij} from the Gaussian distribution $\mathcal{N}(\mu, \sigma)$.

The authors of Ref. [39] carefully studied the case with homogeneous interactions ($\sigma = 0$), and demonstrated that the zero temperature dynamics in this *toy model* can have two distinct outcomes at long times. One possibility is that all abundances converge to a fixed stationary value, with the whole system reaching a fixed point. The second possibility is that not all species converge, resulting in a system with persistent fluctuations. The latter case can also be subdivided into two by taking into account the number of fluctuating species, with one regime with local fluctuations and another with global fluctuations.

One of the main objects to measure is the probability that we obtain persistent fluctuations p_{fluc} after running the dynamics in a graph extracted at random for some average connectivity c and interaction strength μ. Note that p_{fluc} does not distinguish between local and global fluctuations.

The system undergoes a transition around $\mu = 1$ [39]. For $c < \mathrm{e}$, where e is the Euler's constant, and in the limit when the number of species is large ($N \to \infty$), one gets $p_{\text{fluc}} = 0$ for $\mu < \mu_c = 1$ and $p_{\text{fluc}} > 0$ for $\mu > \mu_c = 1$. For $c > e$ and also in the limit $N \to \infty$, the same transition occurs but is displaced to smaller μ, and the critical $\mu_c \lesssim 1$ slowly decreases when the connectivity increases.

As in the case of asymmetric interactions in a random regular graph (Subsection 4.1), we numerically demonstrate that, for this toy model, the probability of having persistent fluctuations in the simulations can be well approximated by studying the probability that IBMF converges. Fig. 4a presents the results obtained in the toy model. In this case, we observe that adding damping to the iterations is particularly important (see Appendix G). For each graph of interactions generated with the rules described above, we also have two different outcomes. Either the abundances converge to a fixed point, or they continue to exhibit persistent fluctuations. As in simulations, the probability of convergence displays a qualitative change in its behavior around $\mu = 1.0$.

The probability P_{nc} that IBMF does not converge corresponds very well to the predictions made in Ref. [39] for the probability of having fluctuations, represented with dashed lines in the main graphic of Fig. 4a. When $\mu > 1$, the authors of Ref. [39] conclude that, in a given graph, the only fluctuating species are located in short cycles of odd length. The species in any cycle of length $n = 2k + 1$ will fluctuate if two conditions are met: i) all the species that are upstream of the species in the cycle are extinct, and ii) the cycle is unstable, which happens for $\mu > 1/\cos(\pi/n)$. The value of μ thus determines the minimum length n_{min} of the fluctuating cycles. For details on the computation, see Appendix G in this article or directly read Ref. [39].

From top to bottom in the figure, the dashed lines correspond to $\mu = 3.0$ ($n_{\text{min}} = 3$), $\mu = 1.5$ ($n_{\text{min}} = 5$), and $\mu = 1.1$ ($n_{\text{min}} = 9$). When n is large, the values of $\mu_c(n) = 1/\cos(\pi/n)$ are close to each other and to $\mu = 1$, and it is numerically harder to distinguish between two values of μ. However, in Fig. 4a, IBMF results for $\mu = 1.1$ (orange points) are not far apart from the corresponding dashed line.

Below $\mu = 1.0$, the results in Fig. 4a are qualitatively different, also in agreement with Ref. [39]. For $\mu = 0.9$, the probability that IBMF does not converge remains close to zero until it abruptly grows towards one around $c \sim 3.1$. If we decrease μ just a bit more to $\mu = 0.8$, we get $P_{\text{nc}} \sim 0$ for all $c < 4$. In Appendix H, we study the dependence of the results on the

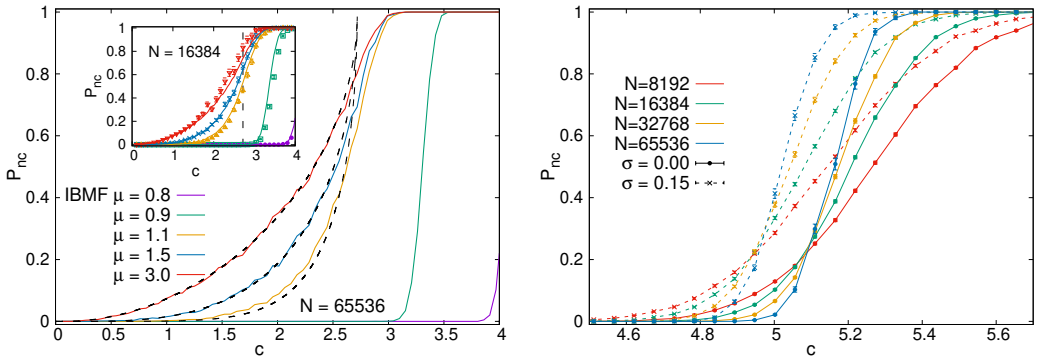

(a) (b)

Figure 4: Probability that IBMF does not converge ($P_{\rm nc}$) in directed graphs. IBMF is run over different realizations of the interaction graph with a given average connectivity c, size N, and interaction strength μ. There is no unique function for all $\mu > 1$, and dashed lines in the top panel are obtained exactly as in Ref. [39] (see the text for clarification). **(a)** Toy model without noise in the interactions ($\sigma = 0$). The colored lines in the main graphic represent the results of IBMF for $N = 65536$ and different values of μ. In the inserted graphic, IBMF (lines) is run instead for systems with $N = 16384$ species, and the points represent the results of simulations of the dynamics for the same system size. The vertical line marks the value $c = \mathrm{e}$. The error bars for IBMF predictions are small and are not included in the graphics. **(b)** The interaction strengths are drawn from the distribution $\mathcal{N}(\mu, \sigma)$ with $\mu = 0.7$ and two values of σ. The values of $P_{\rm nc}$ for different values of c are represented using points with their corresponding error bars. Lines are a guide to the eye.

system size to conclude that, effectively, the transition for $\mu < 1$ is qualitatively different from the one for $\mu > 1$.

Even for a large system with $N = 65536$ species, we observe that the predictions of IBMF deviate from the theory of Ref. [39], which is derived in the infinite size limit. In the inserted graphic of the same Fig. 4a, we show that this is not a particular problem of IBMF. With points, we represent the results obtained after simulating the dynamics by integrating Eq. (18). They are in very good agreement with the predictions of IBMF (lines), also above the critical connectivity $c = \mathrm{e}$, which is marked with a vertical dashed line. As in Section 4.1, IBMF allows us to capture the finite-size effects in simulations, which is an advantage with respect to previous theoretical predictions. It is important to note that, even when run in single graphs, obtaining results from IBMF is computationally much simpler than performing actual simulations. Remarkably, the average runtime of IBMF for different values of μ and c is consistently 10 times faster than the runtime of simulations, under equivalent conditions and for the same system size (see Appendix I).

Moreover, our IBMF equations are not restricted to this toy model, and some variations can also be studied. The authors of Ref. [39] try a modification to include noise in the interaction strengths. They take α_{ij} from the Gaussian distribution $\mathcal{N}(\mu, \sigma)$, again with probability c/N, and zero otherwise. According to the text of that article, for $\sigma = 0$ and $\mu = 0.7$ the transition occurs around $c \sim 5.3$, while for $\sigma = 0.15$ and $\mu = 0.7$ they get $c \sim 4.9$.

Fig. 4b shows the probability that IBMF does not converge when the interactions are drawn using this modified toy model. We include results for $\mu = 0.7$ and two values of σ. With $\sigma = 0$ (continuous lines), we recover the original toy model and use it as a reference. On the other hand, setting $\sigma = 0.15$ (dashed lines) adds noise to the interaction strengths, and the probability that IBMF does not converge increases.

In both cases, we run IBMF for different system sizes. The curves show crossing points at $c \sim 5.1$ and $c \sim 4.95$, for $\sigma = 0$ and $\sigma = 0.15$, respectively. When the number of species N increases, the probability $P_{\rm nc}$ has a sharper transition between $P_{\rm nc} \sim 0$ to the left and $P_{\rm nc} \sim 1$ to the right of the crossing point. If this trend continues as expected when the number of species is large, the values $c \sim 5.1$ and $c \sim 4.95$ are reliable estimates of the location of the transition between a single equilibrium phase and the phase with global fluctuations. These results are indeed close to the ones mentioned in the text of Ref. [39]. The small discrepancies should be investigated further by performing the same analysis with data from simulations of the dynamics, which is technically more difficult because simulations take more computational time (see Appendix I)

4.3 Including thermal noise

Subsections 4.1 and 4.2 show that at $T = 0$ the results of IBMF are in good agreement with simulations. When thermal noise is present ($T > 0$), the fixed points of IBMF can still be efficiently retrieved using Eqs. (9), (10), and (11), and taking advantage of the fact that these integrals can be expressed in terms of Kummer's confluent hypergeometric functions (see Section 3 in the SM). However, in this case, IBMF is an approximation that considers the probability distribution of the system to be factorized as $P(\vec{n}) = \prod_i P(n_i)$. Its predictions, accurate for low temperatures, are expected to fail when T is high enough.

In this section, we study how the results of IBMF depend on the temperature in the gLV model defined over random regular graphs with symmetric interactions. This is a controlled scenario where we have a reliable theoretical technique to compare with, which is BP. We further simplify the setting by eliminating any noise in the interaction strength. Provided that (i, j) is an edge in the random regular graph, we set $\alpha_{ij} = \alpha_{ji} = \mu$, which is equivalent to drawing all α_{ij} from the trivial Gaussian $\mathcal{N}(\mu, 0)$. A first version of the phase diagram T *vs.* μ is available in Ref. [19], where BP is run using discretized abundances. Interestingly, the authors of Ref. [19] note a re-entrant transition in their phase diagram: for low temperatures, the critical value of μ decreases when T increases, until it reaches a minimum. Then, it returns and starts increasing as the temperature continues to rise.

To compare BP with IBMF, which works directly with continuous variables, we use our new implementation of BP with continuous variables. After obtaining the messages by iterating Eq. (13) until convergence, we use Eqs. (14) and (15) to get the *true* marginals. The reader can find details about our implementation in Section 4 of the SM. From Eq. (14), we see that the stationary distribution for a single species has the form $P_{BP}(n_i) = n_i^{\beta\lambda-1}\tilde{P}_{BP}(n_i)/Z_i$, where:

$$\tilde{P}_{BP}(n_i) = \frac{1}{\tilde{Z}_i} \exp\Big\{-\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \prod_{k\in\partial i^-} \int_0^\infty dn_k\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i n_k} \tag{19}$$

can be interpreted as an auxiliary probability distribution if $\tilde{Z}_i$ is taken as the proper normalization factor. The messages $\eta_{k\to i}(n_k)$, necessary to compute $\tilde{P}_{BP}(n_i)$, are the fixed point solution of Eq. (13).

Since the abundance n_i must be positive, the distribution in Eq. (19) is defined only for $n_i \geq 0$. Whenever interactions are absent ($\alpha_{ik} = \mu = 0$ for all i and k), $\tilde{P}_{BP}(n_i)$ becomes a truncated Gaussian centered at $n_i = 1$. Letting μ increase away from zero, one gets a distribution $\tilde{P}_{BP}(n_i)$ that is no longer strictly a truncated Gaussian. Moreover, when the interactions are homogeneous (all $\alpha_{ik} = \mu$), BP converges to the same $\tilde{P}_{BP}(n_i)$ for all sites i. This index can be dropped, and the average distribution $\tilde{P}_{BP}(n) = \sum_i \tilde{P}_{BP}(n_i)/N$ is equal to $\tilde{P}_{BP}(n_i)$ itself.

Figs. 5a, 5b, and 5c show that, even with non-negligible thermal noise and non-zero values of μ, the distribution $\tilde{P}_{BP}(n)$ obtained with BP is not far from Gaussian. These three distributions are obtained for the same temperature $T = 0.03$, using $\mu = 0.04$, $\mu = 0.06$, and $\mu = 0.12$,

respectively. While $\tilde{P}_{BP}(n)$ is represented with colored points in the main graphics, the continuous black lines are the result of fitting truncated Gaussians $\tilde{P}_G(n)$ to the data. The inserted graphics show the relative deviation Δ_{BP-G} of the points with respect to the fits:

$$\Delta_{BP-G} = \frac{\tilde{P}_{BP}(n) - \tilde{P}_G(n)}{\tilde{P}_G(n)} \quad . \tag{20}$$

For $\mu = 0.04$ and $\mu = 0.06$ (Figs. 5a and 5b, respectively), the relative deviation is small for all the values of the abundance. In the corresponding main graphics, the points are indeed very close to the fits. The most significant difference occurs at the tails of the distribution. Close to $n = 0$, we get that $\tilde{P}_{BP}(n)$ is below the truncated Gaussian, while for n large it is above. In other words, the presence of thermal noise and interactions tilts the true distribution and gives slightly more weight to large abundances. The ecosystem is a bit more favorable for the species to thrive.

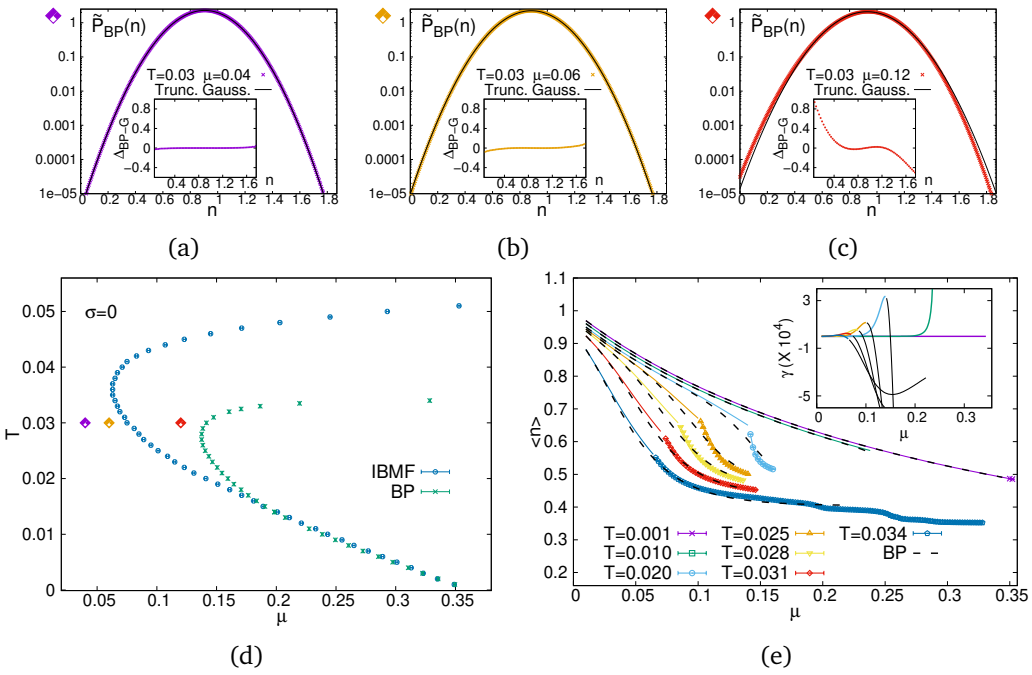

Figure 5: Predictions of IBMF and BP with $T > 0$ for symmetric and homogeneous interactions ($\sigma = 0$) in random regular graphs with connectivity $c = 3$. The immigration rate is $\lambda = 10^{-6}$. **(a), (b), and (c)** Distributions $\tilde{P}_{BP}(n)$ obtained with BP (see Eq. (19)) for $\mu = 0.04, 0.06, 0.12$, at temperature $T = 0.03$, and with system size $N = 128$. Black continuous lines are fits to the points via truncated Gaussians. The inserted graphics show the relative deviation Δ_{BP-G} of the points with respect to the fits (Eq. (20)). **(d)** For each temperature T, we mark the maximum value of μ where BP converges (green points). We also run IBMF on 10000 graphs, each with 10 different initial conditions, and mark the maximum value of μ where it converges to the same fixed point in at least 50% of the graphs (blue points). System sizes are $N = 128$ and $N = 1024$ for BP and IBMF, respectively. **(e)** Average abundance as a function of μ. The black dashed lines are obtained with BP where this algorithm converges. The continuous colored lines are obtained with IBMF where it converges to a single fixed point. The colored points represent the average over several IBMF's fixed points, sampled using 10000 distinct initial conditions. The inserted graphic shows the skewness γ of the distribution $\hat{P}_{BP}(n)$. We use colored lines in the region where IBMF finds a single equilibrium point, and continuous black lines elsewhere.

The difference with respect to a Gaussian increases with μ and is more evident at $\mu = 0.12$ (Fig. 5c). There, the scenario has been reversed. The main and inserted graphics show that now the true distribution is above the Gaussian for $n \sim 0$, and is below for large n. The non-Gaussianity of the distribution gives more weight to species that are close to extinction. In terms of the model, this corresponds to an ecosystem that can support fewer species. Eventually, the situation is no longer compatible with the existence of a fixed point of BP, and this algorithm stops converging around $\mu = 0.1405(5)$. In summary, we observe two distinct types of non-Gaussianity in BP at $T = 0.03$: one that is tilted towards larger abundances, which occurs at small values of μ, and another one tilted towards extinctions, which occurs close to the point where BP stops converging.

In the phase diagram in Fig. 5d, the points where BP stops converging at each temperature are marked with green crosses. The results confirm the re-entrant transition detected in Ref. [19]. When the temperature goes to zero, one recovers the exact result $\mu \approx 0.354$ of Ref. [32] (see Appendix J) for the transition from the single equilibrium phase to the phase with multiple equilibria. The blue circles, on the other hand, represent the prediction made with IBMF for the same transition. We run this approximation for several random regular graphs, in each case using different initial conditions for the average abundances. To the left of the circles, different runs of IBMF converge to the same average abundances in at least 50% of the graphs. To the right, we find instead distinct stationary values of the average abundances just by changing the initial conditions in at least 50% of the graphs. IBMF reproduces very closely the results of BP for low temperatures and, as expected, deviates from it for high temperatures. However, it maintains qualitative agreement with BP, also displaying a re-entrant transition.

It is also important to mention that the computational cost of running IBMF is considerably lower than the cost of BP. When an extensive use of computational resources is required, as in Fig. 5d, this advantage of IBMF plays an important role and one can use it to study larger systems. Nevertheless, we checked that the transition points presented in this figure do not change when the system size is increased, neither for IBMF nor for BP.

Since IBMF is a factorized *ansatz* for the stationary distribution $P_\infty(\vec{n}) = \prod_i P_\infty(n_i)$, and each of the factors $P_\infty(n_i)$ is a Gaussian multiplied by the factor $n_i^{\beta\lambda-1}$, it can be used to shed light on the discussion about the non-Gaussianity of BP's solution. The latter, illustrated here in Figs. 5a, 5b, and 5c, has already been noticed in Ref. [19]. To study it in more detail, we computed the stationary average abundance $\langle n \rangle$ with both techniques, IBMF and BP, for different temperatures. The black dashed lines in Fig. 5e are obtained with BP in the range of values of μ where this algorithm converges. In turn, the colored continuous lines represent the prediction of IBMF when it converges to a single fixed point. However, this is not the only way to estimate the average abundance with IBMF. Even in the region where the fixed point is not unique, we compute $\langle n \rangle$ by averaging over the different fixed points of IBMF (colored points in the figure). Interestingly, the predictions of IBMF closely follow the dashed lines of BP, also in the region where the fixed point of IBMF is not unique. In other words, this approximation describes the average abundance well at all temperatures under consideration.

Furthermore, IBMF's transition, marked with blue circles in Fig. 5d, is related to the type of non-Gaussianity displayed by BP. In the inserted graphic of Fig. 5e, we present the skewness γ of the distribution $\tilde{P}_{BP}(n)$ (see Eq. (19)) to quantify this non-Gaussianity. When the distribution is tilted towards large abundances, we get $\gamma > 0$. On the other hand, when $\tilde{P}_{BP}(n)$ is tilted towards extinctions, we get $\gamma < 0$.

We use colored lines in the inserted graphic to represent the skewness of $\tilde{P}_{BP}(n)$ in the region where IBMF converges to a single fixed point. Black continuous lines are used, in turn, in the region where IBMF converges to different fixed points. The results indicate that around the same value of μ where BP starts developing a distribution $\tilde{P}_{BP}(n)$ that is tilted towards extinctions, IBMF stops converging to a single fixed point.

5 Conclusions

In conclusion, our local closures for the global Fokker-Planck equations, and in particular the Individual Based Mean Field (IBMF) method, provide a powerful and versatile tool for analyzing the stationary states of the generalized Lotka-Volterra model on sparse graphs. We have demonstrated its efficacy across a range of scenarios, from asymmetric interactions on undirected graphs to directed networks and systems with thermal noise. IBMF faces its greatest challenge in the symmetric case, where its assumption of species independence breaks down due to correlations. These correlations are precisely what the Belief Propagation (BP) method captures, but BP is restricted to symmetric interactions. This highlights a key trade-off: BP offers higher accuracy for symmetric networks, while IBMF provides a versatile and effective solution for the more general and common case of asymmetric couplings.

Our analysis reveals that the phase diagram for sparse, asymmetric interactions (Fig. 2) is qualitatively distinct from its symmetric counterpart; notably, we observe a transition towards a multiple-equilibria phase at a positive σ even for $\mu \gtrsim 0$, whereas symmetric interactions always lead to a single fixed point at small μ [19]. This sparse topology also induces a different stability landscape compared to fully-connected systems [5], with a transition occurring at a positive μ for $\sigma = 0$, a phenomenon linked to the intrinsic instability of sparse competitive loops identified in [32].

The observed finite-size effects in the transition lines are correctly captured by IBMF. Notably, the transition lines to both multiple fixed points and unbounded growth progressively shift toward lower heterogeneity σ as the system size N increases. This suggests that the sparse ecological models under consideration may be intrinsically unstable in the infinite-species limit whenever heterogeneity is finite. This observation is consistent with previous RMT results [17, 18, 48], which show that, in general, sparse random matrices remain stable in the large size limit only if the interactions are purely antagonistic or unidirectional. However, for finite sizes we do observe a region where there is a single stable fixed point, a fact that is accurately predicted by our new method.

Looking forward to new applications, the computational efficiency and general applicability of IBMF make it a promising candidate for predicting stable states in real ecological networks, when direct data on interaction strengths is available [49–52] or, in its absence, when one has access to the relevant parameters from which the interaction strengths can be drawn [41, 53]. Furthermore, the methodological framework is not restricted to ecology and could be fruitfully generalized to analyze a wide class of models in economics, evolutionary game theory, and other fields defined on complex, sparse, and even asymmetric interaction networks.

6 Acknowledgments

We thank Giacomo Gradenigo and Chiara Cammarota for early discussions on this problem. Our special gratitude goes to Mattia Tarabolo, Roberto Mulet, and Luca Dall'Asta, with whom we maintained a productive scientific dialogue throughout the writing process, and who generously contributed their own novel insights. This project has been supported by the FIS 1 funding scheme (SMaC - Statistical Mechanics and Complexity) from Italian MUR (Ministry of University and Research) and from the project MIUR-PRIN2022, "Emergent Dynamical Patterns of Disordered Systems with Applications to Natural Communities", code 2022WPHMXK, funded by European Union — Next Generation EU, Mission 4 Component 1 CUP: B53D23005310001. This study was conducted using the DARIAH HPC cluster at CNR-NANOTEC in Lecce, funded by the "MUR PON Ricerca e Innovazione 2014-2020" project, code PIR01 00022.

A From Ito's rule to the Fokker-Planck equation

Let $A(\vec{n})$ be a generic observable that depends on the whole system $\vec{n} = (n_1, \ldots, n_N)$ at time t, but does not explicitly depend on time. For example, $A(\vec{n})$ could be the average abundance $A(\vec{n}) = \sum_i n_i(t)/N$. Following Ito's rule:

$$\frac{d}{dt}\mathbb{E}[A(\vec{n})] = \mathbb{E}\Big[\sum_{i=1}^{N} \frac{\partial A(\vec{n})}{\partial n_i}\frac{dn_i}{dt}\Big] + T\,\mathbb{E}\Big[\sum_{i=1}^{N}\frac{\partial^2 A(\vec{n})}{\partial n_i^2}\,n_i\Big], \tag{21}$$

where $\mathbb{E}[\,\cdot\,]$ is the average over the probability density $P_t(\vec{n})$ of having abundances $\vec{n}$ at time t, which is defined in a stochastic process where each trajectory is given by a specific realization of the thermal noise and a specific choice for the initial conditions. In other words, $\mathbb{E}[\,\cdot\,]$ is an average over the thermal noise and the initial conditions. Using this definition and Eq. (1) for dn_i/dt, one gets:

$$\int_0^\infty d\vec{n} A(\vec{n})\frac{\partial}{\partial t}P_t(\vec{n}) = \sum_{i=1}^{N}\int_0^\infty d\vec{n}\,P_t(\vec{n})\Big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}n_j)+\lambda\Big]\frac{\partial A(\vec{n})}{\partial n_i} + \\ +T\sum_{i=1}^{N}\int_0^\infty d\vec{n}\,P_t(\vec{n})\,n_i\frac{\partial^2 A(\vec{n})}{\partial n_i^2}\,. \tag{22}$$

To obtain Eq. (22) one needs to use the fact that $\langle\xi_i(t)\rangle = 0$, where $\xi_i(t)$ is the Gaussian noise that appears in Eq. (1). Integrating by parts and using that, to have finite moments, $P_t(\vec{n}) \to 0$ faster than n_i^{-2} when $n_i \to \infty$, leads to:

$$\int_0^\infty d\vec{n} A(\vec{n})\frac{\partial}{\partial t}P_t(\vec{n}) = -\sum_{i=1}^{N}\int_0^\infty d\vec{n}A(\vec{n})\,\frac{\partial}{\partial n_i}\Big\{\Big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}n_j)+\lambda\Big]P_t(\vec{n})\Big\} + \\ +T\sum_{i=1}^{N}\int_0^\infty d\vec{n}A(\vec{n})\frac{\partial^2}{\partial n_i^2}\big\{n_iP_t(\vec{n})\big\} + \sum_{i=1}^{N}(T-\lambda)\int_0^\infty\Big[\prod_{k\neq i}dn_k\Big]\lim_{n_i\to 0^+}\big[A(\vec{n})P_t(\vec{n})\big]. \tag{23}$$

In the usual diffusion problems where one follows this procedure, the variables are defined in the whole open space $x \in (-\infty, \infty)$. The property $\lim_{x\to\pm\infty} P(x) = 0$ kills all the terms that come from evaluating the integrands in $x \to \pm\infty$. However, now one has a variable n_i defined in $[0, +\infty)$, and $\lim_{n_i\to 0^+} P_t(\vec{n}) \neq 0$ in general. The last term in Eq. (23) highlights the role of the conditions at the border $n_i = 0$, and this is relevant to find the right local closures.

To continue from here, however, we should impose the proper boundary conditions for this Fokker-Planck equation. To guarantee that $P_t(\vec{n})$ keeps properly normalized, the current of probability density must be zero at every border $n_i = 0$, using what is called a reflecting boundary condition [40]. We must then enforce the relations $(T-\lambda)\lim_{n_i\to 0^+} P_t(n_i) = 0$ for all species i, where $P_t(n_i) = \int_0^\infty[\prod_{k\neq i} dn_k]P_t(\vec{n})$ is the single-site probability for the abundance of species i. Therefore, the last term in Eq. (23) vanishes. In other words, we can neglect the surface terms that arise after integrating by parts.

Since $A(\vec{n})$ is a generic function, the only way to fulfill this relation is to have

$$\frac{\partial P_t(\vec{n})}{\partial t} = T\sum_{i=1}^{N}\frac{\partial^2}{\partial n_i^2}\big\{n_iP_t(\vec{n})\big\} - \sum_{i=1}^{N}\frac{\partial}{\partial n_i}\Big\{\Big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}n_j)+\lambda\Big]P_t(\vec{n})\Big\}, \tag{24}$$

which is the right Fokker-Planck equation, valid for any graph $G(V, E)$.

B Solution for the isolated variable

For the rest of this article, it will be useful to obtain the stationary solution of Eq. (2) in the particular case where there is only one variable n. The equation is then:

$$\frac{\partial}{\partial t}P_t(n) = T\frac{\partial^2}{\partial n^2}\{nP_t(n)\} - \frac{\partial}{\partial n}\Big\{\big[n(1-n)+\lambda\big]P_t(n)\Big\}. \tag{25}$$

In the steady state:

$$0 = T\frac{\partial^2\{nP_\infty(n)\}}{\partial n^2} - \frac{\partial}{\partial n}\Big\{\big[n(1-n)+\lambda\big]P_\infty(n)\Big\}. \tag{26}$$

Integrating over n and remembering that $\lim_{n\to\infty} n^2 P_\infty(n) = 0$, the integration constant goes away. We get:

$$\frac{d}{dn}P_\infty(n) = \frac{1-n}{T}P_\infty(n) + \Big(\frac{\lambda}{T}-1\Big)\frac{P_\infty(n)}{n}. \tag{27}$$

Solving this differential equation with separable variables is simple. The result is:

$$P_\infty(n) = \frac{1}{Z}n^{\beta\lambda-1}\exp\Big\{-\frac{\beta}{2}(n-1)^2\Big\}, \tag{28}$$

where Z is a normalization constant and $\beta \equiv 1/T$.

Once one has the solution (Eq. (28)) to the Fokker-Planck equation for a single species (Eq. (25)), it is not hard to see what will be the solution for IBMF in the stationary regime. In the open space $n_i \in (0,+\infty)$, the equation to fulfill is:

$$0 = -\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(\infty))+\lambda\big]P_\infty(n_i)\Big\} + T\frac{\partial^2}{\partial n_i^2}\{n_iP_\infty(n_i)\}. \tag{29}$$

This is essentially the same Eq. (26), where one substitutes $1-n_i$ by the mean-field expression $1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(\infty)$. Thus, if the solution to the single variable was Eq. (28), the solution to Eq. (29) is:

$$P_\infty(n_i) = \frac{1}{Z_i}n_i^{\beta\lambda-1}\exp\Big\{-\frac{\beta}{2}(n_i-M_i)^2\Big\}, \tag{30}$$

where $M_i = 1-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(\infty)$.

The stationary solution in Eqs. (28) and (30) are normalizable functions for any $\lambda > 0$. Therefore, they give valid distributions $P_\infty(n_i)$ that solves the Fokker-Planck equation when $t\to\infty$ for any $\lambda > 0$. Nevertheless, we identify two distinct qualitative behaviors of $P_\infty(n_i)$, depending on the value of the immigration rate λ. When $\lambda > T$, we get that the probability density is zero at the border ($\lim_{n_i\to 0^+}P_\infty(n_i) = 0$), which is consistent with the boundary conditions imposed by us while deriving the Fokker-Planck equation. In this case, the immigration effectively counteracts the thermic noise and the species are strongly repelled from extinction ($n_i = 0$). When $0 < \lambda < T$, the stationary distribution $P_\infty(n_i)$ diverges at $n_i = 0$. This contradicts the boundary condition $(T-\lambda)\lim_{n_i\to 0^+}P_t(n_i) = 0$.

However, the mathematical inconsistency in the definition of the Fokker-Planck problem for $0 < \lambda < T$ does not impede $P_\infty(n_i)$ from being a valid solution of Eq. (26) also in this interval.

We can still give a physical interpretation to this case: when the effect of the immigration λ is small, a finite fraction of the species goes nearly extinct. Indeed, our simulations in Fig. 1, for a system of N interacting species in a random graph, remain stable even for $\lambda \ll T$. We observe, both from simulations and from IBMF, how a fraction of species spend long times close to extinction.

C Limits of IBMF with parallel updates at zero temperature

Eq. (17), which is the zero-temperature limit of IBMF, can be straightforwardly recast to matrix form as:

$$\vec{n} = \vec{1} - \hat{J} \cdot \vec{n}\,, \tag{31}$$

where $\vec{n} = (n_1, \dots n_N)$ is the vector of species abundances, $\vec{1}$ is a vector full of ones, and $\hat{J}$ is the interaction matrix, whose elements are $J_{ij} = \alpha_{ij}$. In the case with homogeneous interactions, we have that $\hat{J}$ can be expressed in a simple way in terms of the adjacency matrix $\hat{A}$ associated with the interaction graph. By its definition, the element A_{ij} of this matrix is zero if $\alpha_{ij} = 0$, and is one otherwise. Thus, when all nonzero α_{ij} are equal to the same number μ, Eq. (31) transforms into $\vec{n} = \vec{1} - \mu\hat{A} \cdot \vec{n}$.

Similarly to Ref. [32], we can use the properties of $\hat{A}$ to derive the exact single-to-multiple-equilibria transition in the case of random regular graphs with homogeneous interactions. The same equation can be rewritten as $(\mathbb{I} + \mu\hat{A}) \cdot \vec{n} = \vec{1}$, where $\mathbb{I}$ is the identity matrix. Then, the solution is obtained after a matrix inversion by making $\vec{n}^* = (\mathbb{I} + \mu\hat{A})^{-1} \cdot \vec{1}$. For the matrix $\mathbb{I} + \mu\hat{A}$ to be invertible, all its eigenvalues must be nonzero. When A is the adjacency matrix of a large random regular graph, we can use the fact that its smallest eigenvalue must be close to $\lambda^A_{\min} = -2\sqrt{c-1}$ [54]. Then, the smallest eigenvalue of $\mathbb{I} + \mu\hat{A}$ is $\lambda_{\min} = 1 - 2\mu\sqrt{c-1}$. To compute the fixed point $\vec{n}^*$, or equivalently, for the matrix $\mathbb{I} + \mu\hat{A}$ to be invertible, one should verify the relation $\lambda_{\min} > 0$. In other words, the interaction strength μ must satisfy the relation:

$$\mu < \mu^* \equiv \frac{1}{2\sqrt{c-1}}\,, \tag{32}$$

which is the same result presented in Ref. [32]. Since the eigenvalues of $\hat{A}$ cover the interval $\lambda^A \in [-2\sqrt{c-1}, 2\sqrt{c-1}]$, for $\mu \geq \mu^*$ one could always find an eigenvalue in this bulk of the distribution such that $1 + \lambda^A_{\min} = 0$, and the matrix will not be invertible for $\mu \geq \mu^*$. Furthermore, given the homogeneity of the interactions, Eq. (17) must admit the solution:

$$n = \frac{1}{1 + c\,\mu}\,, \tag{33}$$

which is also presented in Ref. [32] as the unique fixed point for $\mu < \mu^*$.

It is also interesting to understand what would happen if one uses Eq. (17) in an iterative way by setting $\vec{n}_{k+1} = \vec{1} - \hat{J} \cdot \vec{n}_k$, choosing a given initial $\vec{n}_0$. This, simply put, is a parallel update that gives a full vector $\vec{n}_{k+1}$ using information about the previous vector $\vec{n}_k$. The result of iterating k times can be expressed as:

$$\vec{n}_k = \left(\vec{1} - \hat{J} \cdot \left(\vec{1} - \hat{J} \cdot (\vec{1} - \hat{J} \cdot \ldots \cdot (1 - \hat{J} \cdot \vec{n}_0))\right)\right) \tag{34}$$

$$\vec{n}_k = (-\hat{J})^k \cdot \vec{n}_0 + \sum_{i=0}^{k-1} (-\hat{J})^i \cdot \vec{1}\ . \tag{35}$$

In Eq. (35), the power $(-\hat{J})^i$ must be interpreted as the product of the matrix $\hat{J}$ with itself i times, which returns a matrix. Now, we can use the expression for the geometric sum of matrices to explicitly write the result of the sum on the right-hand side and get:

$$\begin{aligned}\vec{n}_k &= (-\hat{J})^k \cdot \vec{n}_0 + (\mathbb{I}+\hat{J})^{-1} \cdot \left((\mathbb{I} - (-\hat{J})^k) \cdot \vec{1} \right) \\ \vec{n}_k &= (\mathbb{I}+\hat{J})^{-1} \cdot \vec{1} + (-\hat{J})^k \cdot \left(\vec{n}_0 - (\mathbb{I}+\hat{J})^{-1} \cdot \vec{1} \right) \ . \end{aligned} \tag{36}$$

This expression will converge to the right solution if $\lim_{k\to\infty}(-\hat{J})^k = \hat{0}$, or, equivalently, if the eigenvalue of $\hat{J}$ with the maximum absolute value $|\lambda^J|_{\max}$ is smaller than one. At this point, it is important to note that this conclusion is independent on the specific interaction graph. The convergence of IBMF with parallel updates at zero temperature can be determined by computing the eigenvalue of $\hat{J}$ with the largest absolute value.

Unfortunately, the maximum eigenvalue associated to the adjacency matrix of a random regular graph is not inside the bulk $[-2\sqrt{c-1}, 2\sqrt{c-1}]$. There is an eigenvalue outside of the bulk and its value is $\lambda^{\hat{A}}_{\max} = c$. Therefore, $|\lambda^{\hat{J}}|_{\max} = c\mu$ and the iterations will not converge for any $\mu \geq \mu^{\text{par}} = 1/c$. As μ^{par} is smaller than $\mu^* = 1/(2\sqrt{c-1})$ for any $c > 2$, in these cases there will be an interval $\mu \in [\mu^{\text{par}}, \mu^*)$ where the parallel iterations will not converge to the right solution, even if that solution exists.

To overcome this problem, the results in the main text are obtained using a sequential update. In practice, this means that each n_i^k is updated to n_i^{k+1} asynchronously. We choose a random order of the species, and one by one we apply Eq. (17). When n_i is recomputed, the new value is ready to be used in the next update. It is important to mention that the problem with parallel updates is already known and has been solved using sequential updates in other contexts [55].

D Unbounded growth for asymmetric mutualistic interactions

Fig. 2 in Subsection 4.1 shows that, in random regular graphs with Gaussian asymmetric interactions, the phase with multiple equilibria exists only for positive values of the average interaction strength μ. A positive value of μ corresponds to ecosystems in which most species develop competitive interactions. On the other hand, when μ is negative the interactions are mostly mutualistic (species abundances grow together). In this case, we have only one transition line $\sigma_c(\mu)$ that separates two phases. For $\sigma < \sigma_c(\mu)$ the abundances converge to a single equilibrium state, while for $\sigma > \sigma_c(\mu)$ at least one abundance grows indefinitely and diverges for long times. The latter is called unbounded growth.

The results in the right panel of Fig. 3 are easily extended to $\mu < 0$. Fig. 6 shows that IBMF (lines) maintains a good agreement with the results of the simulations (points) for random regular graphs with connectivity $c = 3$. For $\mu < -1/c \approx -0.333$ the abundances diverge for any value of σ, and the transition line goes to $\sigma(-1/3) = 0$. This is consistent with the fact that at $\sigma = 0$ the interaction strengths α_{ij} are homogeneous and all equal to μ. The solution in the single equilibria phase is then $n_i = 1/(1 + c\,\mu)$ for all species (for $i = 1, \ldots, N$). Given that n_i must be non-negative, when $\mu < -1/c$ we do not have a feasible solution anymore and the abundances diverge in any simulation.

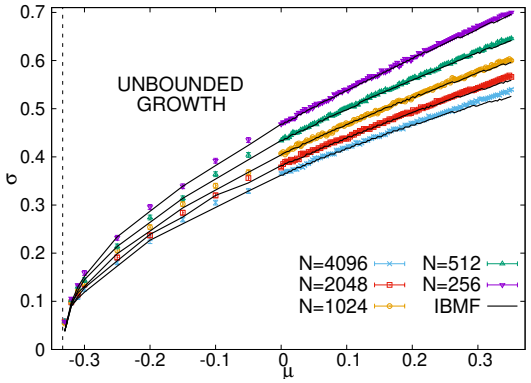

Figure 6: Transitions of the Generalized Lotka-Volterra model for different system sizes at $T = 0$. The interactions are asymmetric (α_{ij} is chosen independently of α_{ji}) and defined on random regular graphs with connectivity $c = 3$. Points represent the results of simulations with immigration rate $\lambda = 10^{-6}$, and lines are the predictions made with IBMF for the same sizes. Each transition was determined using 10000 graphs. Simulations are repeated for 10 different initial conditions. IBMF was run with sequential updates for 10 different random initial conditions. Points (lines) mark the minimum value of σ such that simulations (IBMF) displayed unbounded growth (not converged) in more than 50% of the interaction graphs. The vertical line marks the limit value $\mu = -1/c \approx -0.333$

E Use of damping to improve convergence

As usual in the scenario of iterating equations until the quantities reach a fixed point, some standard tricks can be used to help IBMF converge. Perhaps the most common mechanism is to add damping to the iterations. Given the update rule $m_i = f_i(\{m_k\}_{k\in\partial i^-})$, one chooses a parameter $d \in [0,1]$, and updates the vector $\{m_i^k\}_{i=1,\dots,N}$ of the average abundances by doing:

$$m_i^{k+1} = d \cdot f\left(\{m_k\}_{k\in\partial i^-}\right) + (1-d)\cdot m_i^k \quad . \tag{37}$$

The value $d = 1$ corresponds to the original case, where IBMF is iterated without damping. When $d = 0$ nothing happens to m_i^k. A quick study shows that, for IBMF, the probability of convergence is maximized for some intermediate d around $d = 0.2$. The results in Subsections 4.1 and 4.2 are obtained using precisely this value ($d = 0.2$). In the latter case, the impact of damping is explained in detail in Appendix G. In the first case, achieving convergence with IBMF is an important issue due to the crossover between the phase of multiple fixed points and the phase of unbounded growth.

For $T = 0$ any fixed point of IBMF is also a fixed point for simulations, and we can be sure that whenever we find different fixed points with IBMF this has implications also for simulations. However, the phase of unbounded growth is determined by the divergence of the abundances, and, as said in the main text, it could be that the iteration process of IBMF does not converge while the simulations do. Fig. 7a shows that the differences between IBMF without damping and simulations are indeed noticeable only for $\mu \leq 0.2$, where unbounded growth starts to dominate the crossover. Therefore, one needs to add damping to overcome the convergence problems that are not physical and are only related to the algorithmic dynamics of the iterations.

On the other hand, the phase diagram of Fig. 5d is produced without damping (using $d = 1$). Here, we do not find any problems in achieving convergence with IBMF. In fact, in that phase diagram the unbounded growth phase is not present. The iteration process reaches

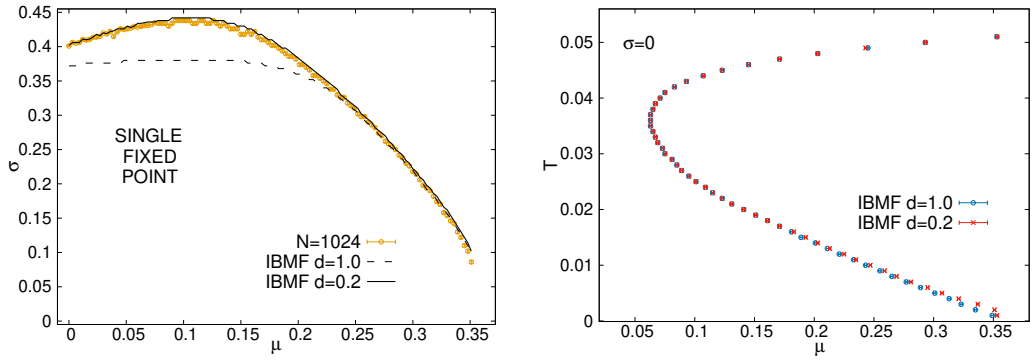

(a) (b)

Figure 7: Effects of damping in the results of IBMF for the phase diagrams of the gLV model. The equations are run on random regular graphs with size $N = 1024$ and connectivity $c = 3$. Each transition was determined using 10000 graphs. In panel **(a)** (panel **(b)**), we mark the maximum value of σ (resp. μ) such that simulations or IBMF converged to the same fixed point in more than 50% of the interaction graphs. IBMF was run on each graph with damping ($d = 0.2$) and without damping ($d = 1$) for 10 different random initial conditions. **(a)** Phase diagram at $T = 0$. The interactions are asymmetric (α_{ij} is chosen independently of α_{ji}). Points represent the results of simulations with immigration rate $\lambda = 10^{-6}$, and lines are obtained with IBMF. Simulations are repeated for 10 different initial conditions. **(b)** Results of IBMF in the presence of thermal noise for graphs with symmetric and homogeneous interactions (drawn using $\sigma = 0$). The immigration rate is $\lambda = 10^{-6}$.

a fixed point for any temperature T and average interaction strength μ. The only relevant question is, at a given temperature, what is the smallest value of μ where we can find two different fixed points. Fig. 7b shows that the answer is approximately the same in most of the phase diagram, except at very low temperatures. As expected, the effect of adding damping, if any, is to move the transition to larger values of μ. Intuitively, the damping could stabilize one fixed point more than the others, preventing the algorithm from sampling them with the right probability. When one wants to correctly locate the transition between the single and the multiple attractor phases, the correct physical results are obtained by using IBMF without damping ($d = 1$).

F IBMF for graphs with correlated couplings

In Subsection 4.1, we present the results of IBMF and simulations of the gLV model on random regular graphs with asymmetric interactions. In that case, the interaction graphs were built drawing α_{ij} and α_{ji} independently for every pair of interacting species. This choice automatically sets the connected correlation $\langle \alpha_{ij}\alpha_{ji}\rangle - \langle \alpha_{ij}\rangle\langle \alpha_{ji}\rangle$ to zero.

It is important, however, to verify that the accuracy of IBMF's predictions extends to cases where the connected correlation is not zero. Fortunately, we can use a simple procedure to build correlated couplings. For each pair $i \to j$ and $j \to i$ of interacting species (let $i < j$ just to fix ideas), we do one of two things: i) with probability ϵ we choose α_{ij} from the Gaussian $\mathcal{N}(\mu, \sigma)$ and then we set $\alpha_{ji} = \alpha_{ij}$, or ii) with probability $1 - \epsilon$ we independently draw α_{ij} and α_{ji} from the same Gaussian. Evidently, the setting used in Subsection 4.1 corresponds to $\epsilon = 0$.

Fig. 8 shows the results for $\epsilon = 0.5$, chosen such that the interactions are still asymmetric,

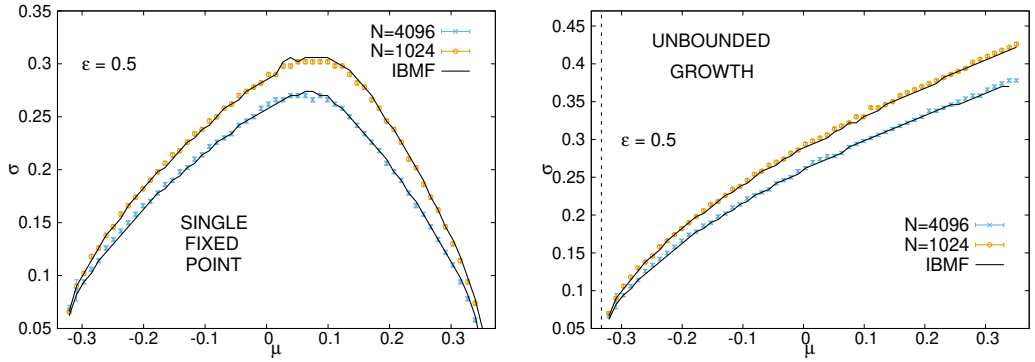

(a) (b)

Figure 8: Transitions of the gLV model for different system sizes at $T = 0$. The interactions are asymmetric and defined on random regular graphs with connectivity $c = 3$. The couplings are chosen such that with probability $\epsilon = 0.5$ we have $\alpha_{ij} = \alpha_{ji}$. Points represent the results of simulations with immigration rate $\lambda = 10^{-6}$, and lines are the predictions made with IBMF for the same sizes. Each transition was determined using 10000 graphs. Simulations are repeated for 10 different initial conditions. IBMF was run with damping (see Appendix E) for 10 different random initial conditions. **(a)** For each system size N and average strength μ, points (lines) mark the maximum value of σ such that simulations (IBMF) converged to the same fixed point in more than 50% of the interaction graphs. **(b)** Points (lines) mark the minimum value of σ such that simulations (IBMF) displayed unbounded growth (not converged) in more than 50% of the interaction graphs.

but correlated. As in Subsection 4.1, the predictions obtained with IBMF for the transitions of the gLV model are in very good agreement with simulations in this case. IBMF provides a precise description of the model's fixed points at $T = 0$, valid for different values of ϵ.

G Convergence of IBMF in directed graphs

The convergence of IBMF is sensitive to the addition of damping. In the toy model described in Subsection 4.2, using no damping ($d = 1$ in Eq. (37)) has negative implications on the convergence, and the results no longer coincide with the predictions in Ref. [39].

In Fig. 9a, we show the probability that IBMF, without damping, does not converge (P_{nc}) for large graphs with different average connectivities and interaction strengths. In this case, P_{nc} is independent of μ for all $\mu > 1$. It follows a slowly increasing function that goes from zero at $c = 0$ to one at $c = e$. This function is represented with a dashed line in the figure, and we give its precise mathematical form below. This behavior is also nearly independent of the size N, as can be seen in Fig. 9b. Only close to $P_{\text{nc}} \sim 1$, for $c \sim e$, finite-size effects make IBMF have a small deviation from the dashed line. The empirical P_{nc} is not exactly equal to one at $c = e$ for finite sizes, but the inserted graphic shows that P_{nc} increases when the number of species N increases.

To compute the function followed by IBMF for $d = 1$ and $\mu > 1$, we can do something analogous to what the authors of Ref. [39] did to predict the probability of fluctuations in simulations. Their results indicate that, for any $\mu > 1$ and in the stationary state, all species in the graph are completely polarized, *i.e.*, either they are extinct ($n_i = 0$) or they reach their carrying capacity ($n_i = 1$). First, we write the probability ϕ that a species is isolated and therefore can have $n_i = 1$. In tree-like graphs, the neighbors belong to nearly independent

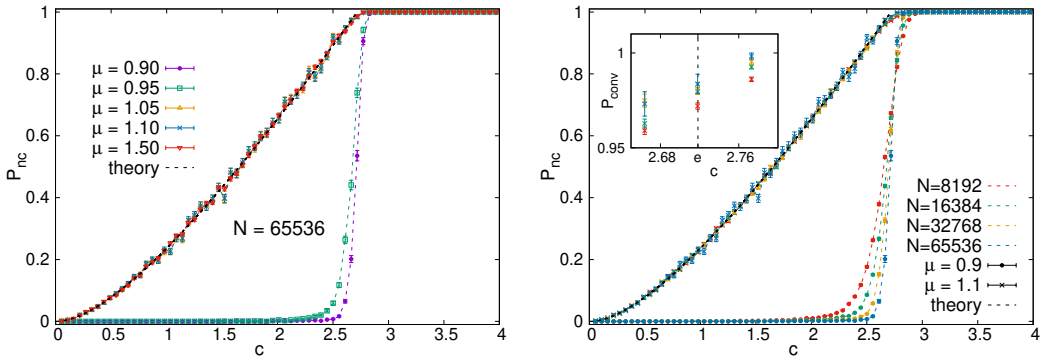

(a) (b)

Figure 9: Probability that IBMF does not converge ($P_{\rm nc}$) in instances of the toy model (directed graphs with $\sigma = 0$). Each point is obtained by running IBMF for 10000 different realizations of the interaction graph with a given average connectivity c, size N, and interaction strength μ. In both panels, IBMF is run without damping, and for $\mu > 1$ its results follow a unique function represented using a dashed line. This function is the analytical result (see the computation in the text) for the probability of having at least one isolated cycle in the graph formed by species with completely polarized abundances ($n_i = 0$ or $n_i = 1$). It is close to zero for small c and goes to one exactly at $c = e$. **(a)** IBMF convergence without damping for $N = 65536$ and different values of μ. **(b)** IBMF convergence without damping for two values of the interaction strength around the transition ($\mu = 0.9$ and $\mu = 1.1$) obtained with several system sizes. The inserted graphic shows an enlargement of the zone where $P_{\rm conv} \sim 1$.

subgraphs, and we get $\phi = \sum_{k=0}^{\infty} p(k)(1-\phi)^k$, where k is the connectivity of the species and $p(k)$ is its distribution. When $p(k)$ is Poisson, the authors of Ref. [39] show that $\phi = W(c)/c$, where c is the average connectivity and $W(x)$ is the Lambert W function. Second, we compute the probability that a directed cycle of length n is isolated. Indeed, a species in a directed cycle has one incoming edge from another species inside the cycle. For the cycle to be isolated, we need all the other incoming edges to correspond to extinct species. Since in Poisson graphs the number of these other incoming edges is also distributed as Poisson with the same mean, and the neighbors of different species in the cycle are independent, the probability that the cycle of length n is isolated is simply ϕ^n.

The number of directed cycles of length n in a Poisson graph with mean connectivity c is also Poisson distributed with mean c^n/n [56]. Thus, the number of directed isolated cycles is Poisson distributed with mean $(c\,\phi)^n/n = [W(c)]^n/n$. The probability of having at least one cycle is then:

$$f(c) = 1 - \exp\Big\{-\sum_{n=2}^{\infty} \frac{[W(c)]^n}{n}\Big\}$$
$$f(c) = 1 - \frac{c}{W(c)}\big(1 - W(c)\big), \tag{38}$$

where we used that $W(c)\,e^{W(c)} = c$. This function $f(c)$ is the one represented with dashed lines in Fig. 9.

On the other hand, for $\mu < 1$, the probability that IBMF does not converge without damping is very close to zero for all $c < e$. In this regime, the results for different system sizes N have a nice crossing point at $c > e$, as can be seen in Fig. 9a for $\mu = 0.9$. This is a familiar feature

of a first-order phase transition that happens exactly at that crossing point. To the left of the crossing, the probability that IBMF does not converge goes to zero when the number of species N goes to infinity. To the right, the probability approaches one as $N \to \infty$.

Adding damping considerably helps IBMF to converge at any μ. For the interesting case $\mu > 1$, setting $d = 0.2$ avoids the problems caused by a subset of the directed isolated cycles. Now, the probability P_{nc} does not follow a unique function $f(c)$ for all $\mu > 0$ (see Fig. 4). They coincide, instead, with the predictions made by the authors of Ref. [39]. Their computation is analogous, but with a key difference. They concluded that, for the simulations, the cycles with even length do not cause fluctuations. If we exclude the even values of n, Eq. (38) changes to:

$$\begin{aligned} f_3(c) &= 1 - \exp\Big\{-\sum_{k=1}^{\infty} \frac{[W(c)]^{2k+1}}{2k+1}\Big\} \\ f_3(c) &= 1 - \frac{c}{W(c)}\sqrt{\frac{1-W(c)}{1+W(c)}}\,. \end{aligned} \tag{39}$$

Eq. (39) gives the probability $f_3(c)$ of having at least one directed isolated cycle with odd length. Furthermore, a cycle of odd length n will be unstable, and thus will fluctuate, for all $\mu > \mu_c(n) = 1/\cos(\pi/n) > 1$ [39]. Therefore, Eq. (39) gives the probability of fluctuations for any $\mu > 1/\cos(\pi/3) = 2$. When $\mu < 2$ the cycles of length $n = 3$ are stable, but the ones with $n = 5$ are still unstable for any $\mu > 1/\cos(\pi/5) \approx 1.24$. Thus, to compute the line that corresponds to the blue points (done for $\mu = 1.5$) in Fig. 4a, we simply need to subtract the number of cycles with length $n = 3$ from the sum in Eq. (39). We get the probability:

$$\begin{aligned} f_5(c) &= 1 - \exp\Big\{-\sum_{k=2}^{\infty} \frac{[W(c)]^{2k+1}}{2k+1}\Big\} \\ f_5(c) &= 1 - \exp\Big\{\frac{1}{2}\ln\Big(\frac{1+W(c)}{1-W(c)}\Big) + W(c) + \frac{[W(c)]^2}{3}\Big\}, \end{aligned} \tag{40}$$

that we also plot using a dashed line in Fig. 4a, showing that it indeed coincides very well with the results of IBMF obtained at $\mu = 1.5$.

This is the procedure to follow for any μ. Fig. 4a indicates that, once we use damping, the convergence of IBMF stops being affected by the cycles of even length and coincides with the theoretical predictions for the probability of fluctuations as presented in Ref. [39].

H Finite size effects of IBMF on directed graphs

Fig. 10 shows the finite size effects for two values of μ around $\mu = 1.0$ when we run IBMF for the toy model defined in Subsection 4.2. We observe two distinct types of transitions in the probability that IBMF does not converge (P_{nc}). As with the simulations in Ref. [39], when $\mu > 1$ the results for finite systems do not reach $P_{\text{nc}} = 1$ exactly at $c = e$. We present numerical evidence that, for $\mu = 1.1$, the probability P_{nc} increases with the system size and the points move to the left towards the line $c = e$.

The curves for $\mu = 0.9$, instead, have a clear crossing point at $c \sim 3.28$. When the number of species N increases, the probability P_{nc} has a sharper transition between $P_{\text{nc}} \sim 0$ to the left and $P_{\text{nc}} \sim 1$ to the right of the crossing point. Therefore, the value $c \sim 3.28$ is a good estimate for the location of the transition between the single equilibrium phase and the phase with global fluctuations. Indeed, it is compatible with the results in Fig. 1 of Ref. [39].

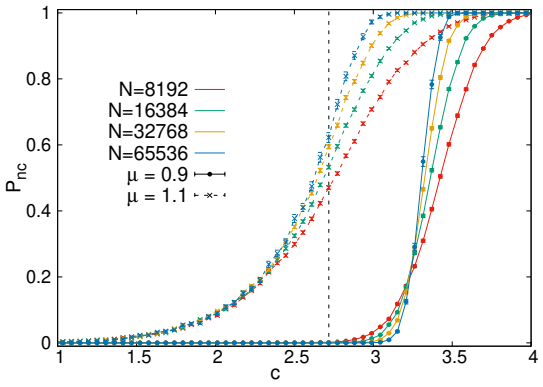

Figure 10: Probability that IBMF does not converge ($P_{\rm nc}$) in instances of the modified toy model (directed graphs with $\sigma = 0$). Each point is obtained by running IBMF for 10000 different realizations of the interaction graph with a given average connectivity c and several sizes N. IBMF is run with damping for two values of the interaction strength around the transition ($\mu = 0.9$ and $\mu = 1.1$). The vertical line marks the value $c = e$.

I Comparing runtimes of IBMF and simulations in directed graphs

One of the advantages of IBMF is that it can be implemented efficiently (see Section 3 of the SM). Here, we show that running IBMF in single graphs is considerably faster than running simulations. Although we present data for the toy model on directed graphs of Subsection 4.2, this conclusion is generally applicable to all the scenarios discussed in this article.

In Fig. 11, we compare the average wall-clock time required for IBMF to converge to a fixed point with the corresponding simulation wall-clock time. The average discards all the samples that lead to persistent fluctuations for long times. For all connectivities and for both values of μ in the figure, IBMF is consistently around 10 times faster to reach convergence. Both algorithms, available at Ref. [43], were run on a single CPU Intel Xeon Gold 6248 2.5G.

It is important to note that the iterations required by IBMF to converge do not possess a physical meaning. The process is discrete, and at each step we update one of the average abundances m_i using Eq. (17). On the other hand, the simulation involves integrating a differential equation (see Eq. (18)) whose time t does have a physical meaning and is a continuous variable. Therefore, the wall-clock time is sensitive to the precision of the integration in time. In this case, we use an adaptive step size to optimize the number of steps needed to reach convergence.

Our results indicate that, even at zero temperature, where the dynamics is simpler to simulate, it is advantageous to run IBMF instead. It gives fast and accurate predictions, as can be seen in Subsections 4.1 and 4.2.

J Zero-temperature limit of BP

We discuss here the zero temperature limit of BP for random regular graphs with symmetric and homogeneous interactions ($\sigma = 0$, see Section 4.3). We show that the exact results for the single-to-multiple equilibria [32] can be easily retrieved after properly taking the limit $T \to 0$. This has already been suggested by the numerical results in Ref. [19], obtained with the discretized version of BP, and by the results presented here in Fig. 5d.

When the temperature is small, the probability densities concentrate around the mean

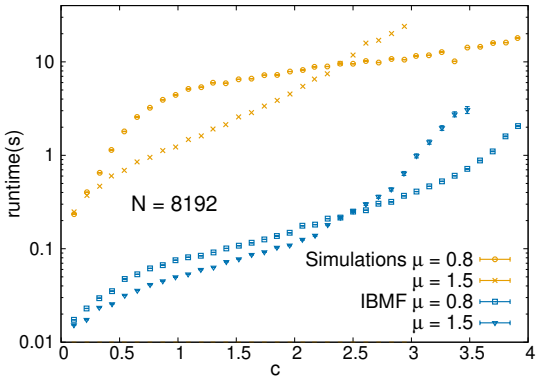

Figure 11: Average runtime (in seconds) required to obtain IBMF predictions and to run simulations for the toy model (directed graphs with $\sigma = 0$, see Subsection 4.2). Runs that show persistent fluctuations are not included. System size is $N = 8192$ in both cases. The points, with error bars, are averages over different graphs for which IBMF and the simulations converge to a fixed point. We consider 10000 graphs in total, and discarded the cases where fewer than 200 lead to convergence.

values. We can then assume that the message $\eta_{i\to j}(n_i)$ can be written as the multiplication of a Gaussian factor and a power-law factor $n_i^{\beta\lambda-1}$ as follows:

$$\eta_{i\to j}(n_i) = \frac{1}{z_{i\to j}} n_i^{\beta\lambda-1} \exp\Big\{-\frac{\beta}{2Q_{i\to j}^2}\big(n_i - m_{i\to j}\big)^2\Big\}\,. \tag{41}$$

This means that every message $\eta_{i\to j}(n_i)$ can be parameterized using the mean $m_{j\to i}$ and the variance $Q^2_{j\to i}$. The update rule (Eq. (13)) becomes:

$$\eta_{i\to j}(n_i) \sim n_i^{\beta\lambda-1}\, e^{-\beta(n_i^2-2n_i)/2} \prod_{k\in\partial i^-\setminus j} \int_0^\infty dn_k\, n_k^{\beta\lambda-1} e^{-\beta\,\alpha_{ik}\,n_i\,n_k - \beta(n_k^2 - 2n_k m_{k\to i})/2Q^2_{k\to i}}\,. \tag{42}$$

When $\beta\to\infty$, with $\beta\lambda$ finite, the integral is dominated by the value $n_k = n_k^*$ such that the argument of the exponential is maximum. Finding this maximum is equivalent to compute the minimum of $f(n_k) = n_k^2 - 2\,n_k\,m_{k\to i} + 2\,\alpha_{ik}\,n_i\,n_k\,Q^2_{k\to i}$. The result is $n_k^* = m_{k\to i} - \alpha_{ik}\,n_i\,Q^2_{k\to i}$. Therefore:

$$\begin{aligned}\eta_{i\to j}(n_i) \;\sim\; & n_i^{\beta\lambda-1}\, e^{-\beta(n_i^2-2n_i)/2}\Big[\prod_{k\in\partial i^-\setminus j}\big(m_{k\to i} - \alpha_{ik}\,n_i\,Q^2_{k\to i}\big)^{\beta\lambda-1}\Big]\times \\ & \times\exp\Big\{\sum_{k\in\partial i^-\setminus j}\frac{\beta}{2Q^2_{k\to i}}\big(m_{k\to i} - \alpha_{ik}\,n_i\,Q^2_{k\to i}\big)^2\Big\}\,.\end{aligned} \tag{43}$$

Using again that, when $\beta\to\infty$ with $\beta\lambda$ finite, the distribution $\eta_{i\to j}(n_i)$ will concentrate around $n_i = n_i^*$ such that the argument of the exponential is maximum, we get:

$$\eta_{i\to j}(n_i) = \frac{n_i^{\beta\lambda-1}}{z_{i\to j}} \exp\Big\{-\frac{\beta}{2}n_i^2\big(1 - \sum_{k\in\partial i^-\setminus j}\alpha_{ik}^2 Q^2_{k\to i}\big) + \beta\,n_i\big(1 - \sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}\big)\Big]\Big\}\,. \tag{44}$$

Comparing with Eq. (41), we can easily identify that:

$$Q^2_{i\to j} = \frac{1}{1-\sum_{k\in\partial i^-\setminus j}\alpha^2_{ik}Q^2_{k\to i}} \tag{45}$$

$$m_{i\to j} = \frac{1-\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}}{1-\sum_{k\in\partial i^-\setminus j}\alpha^2_{ik}Q^2_{k\to i}} \ . \tag{46}$$

These equations are known as relaxed Belief Propagation [57], and in this case correspond to the zero temperature expansion of BP. When we have a random regular graph with homogeneous interactions ($\alpha_{ij}\equiv\mu$ for all edges in the graph), all sites become equivalent and

$$Q^2_{\to} = \frac{1}{1-(c-1)\mu^2Q^2_{\to}} \tag{47}$$

$$m_{\to} = \frac{1-(c-1)\mu\,m_{\to}}{1-(c-1)\mu^2Q^2_{\to}} \ . \tag{48}$$

From Eq. (47), we can obtain a closed expression for the variance $Q^2_{\to}$:

$$Q^2_{\to} = \frac{1}{2(c-1)\mu^2}\left(1\pm\sqrt{1-4(c-1)\mu^2}\right). \tag{49}$$

From where it follows that, in order to have $Q^2_{\to}\in\mathbb{R}$, the strength of the interactions must fulfill the relation

$$\mu\leq\mu^*\equiv\frac{1}{2\sqrt{c-1}} \ . \tag{50}$$

This result, again, coincides with the exact relation obtained in Ref. [32] and we already presented it in Eq. (32). Finally, we can also use that

$$m_i = \frac{1-\sum_{k\in\partial i^-}\alpha_{ik}\,m_{k\to i}}{1-\sum_{k\in\partial i^-}\alpha^2_{ik}Q^2_{k\to i}} \ , \tag{51}$$

In the case with homogeneous interactions we have $\alpha^2_{ik}=\mu^2$ and $\alpha^2_{ik}=\mu^2$ for all edges in the graph. Thus:

$$m\equiv\frac{1-c\mu\,m_{\to}}{1-c\mu^2Q^2_{\to}} \ , \tag{52}$$

together with Eqs. (47) and (48) to get another exact result:

$$m=\frac{1}{1+c\mu} \ . \tag{53}$$

With Eqs. (50) and (53), we recover two known results for the single-to-multiple-equilibria transition for $\beta\to\infty$ [32]. In terms of BP, this transition is simply a boundary $\mu^*=1/(2\sqrt{c-1})$ such that, for $\mu>\mu^*$ and at low temperature ($\beta\gg1$), it is impossible to have a "Gaussian" stationary point like the one in Eq. (41) (more precisely, a Gaussian multiplied by the power $n^{\beta\lambda-1}$).

References

[1] J. Moran and J.-P. Bouchaud, *May's instability in large economies*, Physical Review E **100**(3), 032307 (2019), doi: 10.1103/PhysRevE.100.032307.

[2] T. Galla and J. D. Farmer, *Complex dynamics in learning complicated games*, Proceedings of the National Academy of Sciences **110**(4), 1232 (2013), doi: 10.1073/pnas.1109672110.

[3] J. Garnier-Brun, M. Benzaquen, S. Ciliberti and J.-P. Bouchaud, *A new spin on optimal portfolios and ecological equilibria*, Journal of Statistical Mechanics: Theory and Experiment **2021**(9), 093408 (2021), doi: 10.1088/1742-5468/ac21d9.

[4] A. Altieri, F. Roy, C. Cammarota and G. Biroli, *Properties of equilibria and glassy phases of the random lotka-volterra model with demographic noise*, Physical Review Letters **126**(25), 258301 (2021), doi: 10.1103/PhysRevLett.126.258301.

[5] G. Bunin, *Ecological communities with lotka-volterra dynamics*, Physical Review E **95**(4), 042414 (2017), doi: 10.1103/PhysRevE.95.042414.

[6] T. Galla, *Dynamically evolved community size and stability of random lotka-volterra ecosystems (a)*, Europhysics Letters **123**(4), 48004 (2018), doi: 10.1209/0295-5075/123/48004.

[7] V. Ros, F. Roy, G. Biroli, G. Bunin and A. M. Turner, *Generalized lotka-volterra equations with random, nonreciprocal interactions: The typical number of equilibria*, Physical Review Letters **130**(25), 257401 (2023), doi: 10.1103/PhysRevLett.130.257401.

[8] V. Ros, F. Roy, G. Biroli and G. Bunin, *Quenched complexity of equilibria for asymmetric generalized lotka–volterra equations*, Journal of Physics A: Mathematical and Theoretical **56**(30), 305003 (2023), doi: 10.1088/1751-8121/ace00f.

[9] L. Poley, J. W. Baron and T. Galla, *Generalized lotka-volterra model with hierarchical interactions*, Physical Review E **107** (2023), doi: 10.1103/PhysRevE.107.024313.

[10] J. W. Baron, T. J. Jewell, C. Ryder and T. Galla, *Breakdown of random-matrix universality in persistent lotka-volterra communities*, Physical Review Letters **130** (2023), doi: 10.1103/PhysRevLett.130.137401.

[11] I. A. Hatton, O. Mazzarisi, A. Altieri and M. Smerlak, *Diversity begets stability: Sublinear growth and competitive coexistence across ecosystems*, Science **383**(6688) (2024), doi: 10.1126/science.adg8488.

[12] M. Barbier, J.-F. Arnoldi, G. Bunin and M. Loreau, *Generic assembly patterns in complex ecological communities*, Proceedings of the National Academy of Sciences **115**(9) (2018), doi: 10.1073/pnas.1710352115.

[13] J. Giral Martínez, M. Barbier and S. De Monte, *Interplay of structured and random interactions in many-species ecological dynamics*, bioRxiv pp. 2024–11 (2024), doi: 10.1101/2024.11.25.624858.

[14] M. T. Pearce, A. Agarwala and D. S. Fisher, *Stabilization of extensive fine-scale diversity by ecologically driven spatiotemporal chaos*, Proceedings of the National Academy of Sciences **117**(25), 14572 (2020), doi: 10.1073/pnas.1915313117.

[15] J. A. Dunne, R. J. Williams and N. D. Martinez, *Food-web structure and network theory: the role of connectance and size*, Proceedings of the National Academy of Sciences **99**(20), 12917 (2002), doi: 10.1073/pnas.192407699.

[16] D. M. Busiello, S. Suweis, J. Hidalgo and A. Maritan, *Explorability and the origin of network sparsity in living systems*, Scientific reports **7**(1), 12323 (2017), doi: 10.1038/s41598-017-12521-1.

[17] A. M. Mambuca, C. Cammarota and I. Neri, *Dynamical systems on large networks with predator-prey interactions are stable and exhibit oscillations*, Physical Review E **105**(1), 014305 (2022), doi: 10.1103/PhysRevE.105.014305.

[18] P. Valigi, I. Neri and C. Cammarota, *Local sign stability and its implications for spectra of sparse random graphs and stability of ecosystems*, Journal of Physics: Complexity **5**(1), 015017 (2024), doi: 10.1088/2632-072X/ad2698.

[19] T. Tonolo, M. C. Angelini, S. Azaele, A. Maritan and G. Gradenigo, *Generalized lotka-volterra model with sparse interactions: non-gaussian effects and topological multiple-equilibria phase* (2025), 2503.20887.

[20] F. Roy, G. Biroli, G. Bunin and C. Cammarota, *Numerical implementation of dynamical mean field theory for disordered systems: application to the lotka–volterra model of ecosystems*, Journal of Physics A: Mathematical and Theoretical **52**(48), 484001 (2019), doi: 10.1088/1751-8121/ab1f32.

[21] S. Azaele and A. Maritan, *Generalized dynamical mean field theory for non-gaussian interactions*, Phys. Rev. Lett. **133**, 127401 (2024), doi: 10.1103/PhysRevLett.133.127401.

[22] F. Aguirre-López, *Heterogeneous mean-field analysis of the generalized lotka–volterra model on a network*, Journal of Physics A: Mathematical and Theoretical **57**(34), 345002 (2024), doi: 10.1088/1751-8121/ad6ab2.

[23] F. L. Metz, *Dynamical mean-field theory of complex systems on sparse directed networks*, Physical Review Letters **134**(3), 037401 (2025), doi: 10.1103/PhysRevLett.134.037401.

[24] Y. Wang, D. Chakrabarti, C. Wang and C. Faloutsos, *Epidemic spreading in real networks: an eigenvalue viewpoint*, In *22nd International Symposium on Reliable Distributed Systems, 2003. Proceedings.*, pp. 25–34, doi: 10.1109/RELDIS.2003.1238052 (2003).

[25] E. Cator and P. Van Mieghem, *Second-order mean-field susceptible-infected-susceptible epidemic threshold*, Physical review E **85**(5), 056111 (2012), doi: 10.1103/PhysRevE.85.056111.

[26] R. Pastor-Satorras, C. Castellano, P. Van Mieghem and A. Vespignani, *Epidemic processes in complex networks*, Reviews of Modern Physics **87**(3), 925 (2015), doi: 10.1103/RevModPhys.87.925.

[27] E. Ortega, D. Machado and A. Lage-Castellanos, *Dynamics of epidemics from cavity master equations: Susceptible-infectious-susceptible models*, Physical Review E **105**, 024308 (2022), doi: 10.1103/PhysRevE.105.024308.

[28] E. Aurell, E. Domínguez, D. Machado and R. Mulet, *A theory of non-equilibrium local search on random satisfaction problems*, Physical Review Letters **123**, 230602 (2019), doi: 10.1103/PhysRevLett.123.230602.

[29] D. Machado, J. González-García and R. Mulet, *Local equations describe unreasonably efficient stochastic algorithms in random k-sat* (2025), 2504.06757.

[30] D. Machado, R. Mulet and F. Ricci-Tersenghi, *Improved mean-field dynamical equations are able to detect the two-step relaxation in glassy dynamics at low temperatures*, Journal of Statistical Mechanics: Theory and Experiment **2023**(12), 123301 (2023), doi: 10.1088/1742-5468/ad0f90.

[31] A. Jedrzejewski and K. Sznajd-Weron, *Pair approximation for the q-voter models with quenched disorder on networks*, Phys. Rev. E **105**, 064306 (2022), doi: 10.1103/PhysRevE.105.064306.

[32] S. Marcus, A. M. Turner and G. Bunin, *Local and collective transitions in sparsely-interacting ecological communities*, PLoS Computational Biology **18**(7), e1010274 (2022), doi: 10.1088/1751-8121/ab1f32.

[33] M. Loreau and C. De Mazancourt, *Biodiversity and ecosystem stability: a synthesis of underlying mechanisms*, Ecology letters **16**, 106 (2013), doi: 10.1111/ele.12073.

[34] G. Biroli, G. Bunin and C. Cammarota, *Marginally stable equilibria in critical ecosystems*, New Journal of Physics **20** (2018), doi: 10.1088/1367-2630/aada58.

[35] F. Larroya and T. Galla, *Demographic noise in complex ecological communities*, Journal of Physics: Complexity **4**(2) (2023), doi: 10.1088/2632-072X/acd21b.

[36] G. Garcia Lorenzana, A. Altieri and G. Biroli, *Interactions and migration rescuing ecological diversity*, PRX Life **2**(1) (2024), doi: 10.1103/PRXLife.2.013014.

[37] T. Arnoulx de Pirey and G. Bunin, *Many-species ecological fluctuations as a jump process from the brink of extinction*, Phys. Rev. X **14**, 011037 (2024), doi: 10.1103/PhysRevX.14.011037.

[38] A. Altieri and G. Biroli, *Effects of intraspecific cooperative interactions in large ecosystems*, SciPost Phys. **12**, 013 (2022), doi: 10.21468/SciPostPhys.12.1.013.

[39] S. Marcus, A. M. Turner and G. Bunin, *Local and extensive fluctuations in sparsely interacting ecological communities*, Phys. Rev. E **109**, 064410 (2024), doi: 10.1103/PhysRevE.109.064410.

[40] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, Springer Series in Synergetics. Springer Berlin, Heidelberg, Berlin, Heidelberg, 4 edn., ISBN 978-3-540-70712-7 (2009).

[41] J. Pasqualini, A. Maritan, A. Rinaldo, S. Facchin, E. Savarino, A. Altieri and S. Suweis, *Microbiomes through the looking glass: Linking species interactions to dysbiosis through a disordered lotka-volterra framework*, eLife (2025), doi: 10.7554/elife.105948.2.

[42] I. Gradshteyn and I. Ryzhik, *Table of Integrals, Series, and Products*, Elsevier, 7th Edition, ISBN 9780123736376 (2007).

[43] D. Machado, P. Valigi, T. Tonolo and M. C. Angelini, *GenLotkaVolterra_SparseGraphs: Github repository for this article*, doi: 10.5281/zenodo.17612770 (2025).

[44] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes in C*, Cambridge University Press, ISBN 978-0-521-88068-8 (2007).

[45] G. N. Mil'shtejn, *Approximate integration of stochastic differential equations*, Theory of Probability & Its Applications **19**(3), 557 (1975), doi: 10.1137/1119062, https://doi.org/10.1137/1119062.

[46] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Stochastic Modelling and Applied Probability. Springer Berlin, Heidelberg, Berlin, Heidelberg, 1 edn., ISBN 978-3-540-54062-5, doi: 10.1007/978-3-662-12616-5 (1992).

[47] T. A. de Pirey and G. Bunin, *Critical behavior of a phase transition in the dynamics of interacting populations*, SciPost Phys. **18**, 051 (2025), doi: 10.21468/SciPostPhys.18.2.051.

[48] P. Valigi, J. W. Baron, I. Neri, G. Biroli and C. Cammarota, *Eigenvalue spectral tails and localisation properties of asymmetric networks*, Journal of Physics A: Mathematical and Theoretical **58**(45), 455002 (2025), doi: 10.1088/1751-8121/ae16ec.

[49] P. C. de Ruiter, A.-M. Neutel and J. C. Moore, *Energetics, patterns of interaction strengths, and stability in real ecosystems*, Science **269**(5228), 1257 (1995), doi: 10.1126/science.269.5228.1257, https://www.science.org/doi/pdf/10.1126/science.269.5228.1257.

[50] R. R. Stein, V. Bucci, N. C. Toussaint, C. G. Buffie, G. Rätsch, E. G. Pamer, C. Sander and J. B. Xavier, *Ecological modeling from time-series inference: Insight into dynamics and stability of intestinal microbiota*, PLoS Computational Biology **9**(12), e1003388 (2013), doi: 10.1371/journal.pcbi.1003388.

[51] C. Jacquet, C. Moritz, L. Morissette, P. Legagneux, F. Massol, P. Archambault and D. Gravel, *No complexity–stability relationship in empirical ecosystems*, Nature Communications **7**, 12573 (2016), doi: 10.1038/ncomms12573.

[52] M. Barbier, C. de Mazancourt, M. Loreau and G. Bunin, *Fingerprints of high-dimensional coexistence in complex ecosystems*, Phys. Rev. X **11**, 011009 (2021), doi: 10.1103/PhysRevX.11.011009.

[53] M. Barbier, G. Bunin and M. A. Leibold, *Getting more by asking for less: Linking species interactions to species co-distributions in metacommunities*, Peer Community Journal **5**, e1 (2025), doi: 10.24072/pcjournal.483.

[54] B. D. McKay, *The expected eigenvalue distribution of a large regular graph*, Linear Algebra and its Applications **40**, 203 (1981), doi: https://doi.org/10.1016/0024-3795(81)90150-6.

[55] F. Caltagirone, F. Krzakala and L. Zdeborová, *On convergence of approximate message passing*, In *ISIT 2014 - 2014 IEEE International Symposium on Information Theory*, vol. 2014, pp. 1812–1816. IEEE, IEEE, Honolulu, United States, doi: 10.1109/ISIT.2014.6875146, 5 pages, 3 figures (2014).

[56] B. Bollobas, *Random Graphs*, Cambridge Studies in Advanced Mathematics, 2nd Edition (2001).

[57] Y. Kabashima, *A cdma multiuser detection algorithm on the basis of belief propagation*, Journal of Physics A: Mathematical and General **36**(43), 11111 (2003), doi: 10.1088/0305-4470/36/43/030.

Supplemental Materials for the article

Local Fokker-Planck equations for the Generalized Lotka-Volterra model on sparse random graphs

D. Machado[1,2,3,*], P. Valigi[1], T. Tonolo[5, 6], and M. C. Angelini[1, 4]

[1]Dipartimento di Fisica, Sapienza Università di Roma, P.le Aldo Moro 5, 00185 Rome, Italy
[2]Group of Complex Systems and Statistical Physics and Department of Theoretical Physics. Physics Faculty, University of Havana. CP10400, La Habana, Cuba
[3]CNR - Nanotec, unità di Roma, P.le Aldo Moro 5, 00185 Rome, Italy
[4]Istituto Nazionale di Fisica Nucleare, Sezione di Roma I, P.le A. Moro 5, 00185 Rome, Italy
[5]Gran Sasso Science Institute, Viale F. Crispi 7, 67100 L'Aquila, Italy
[6]INFN-Laboratori Nazionali del Gran Sasso, Via G. Acitelli 22, 67100 Assergi (AQ), Italy
[*]david.machado@uniroma1.it

January 14, 2026

S1 Local closures

To write the local closures, one needs to marginalize Eq. 2 in the main text to obtain the differential equations for the local probabilities. This will be done first with $P_t(n_i) = \int_0^\infty [\prod_{k\neq i} dn_k] P_t(\vec{n})$. Marginalizing over all the abundances except n_i, one gets:

$$\frac{\partial}{\partial t}P_t(n_i) = -\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i)+\lambda\big]P_t(n_i)\Big\} + T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i)\big\}$$
$$+\frac{\partial}{\partial n_i}\Big\{n_i\sum_{j\in\partial i^-}\alpha_{ij}\int_0^\infty dn_j\, n_j P_t(n_i,n_j)\Big\}$$
$$-\sum_{k\neq i}\int_0^\infty dn_k\frac{\partial}{\partial n_k}\Big\{\int_0^\infty\Big(\prod_{r\neq k,i}dn_r\Big)\big[n_k(1-n_k-\sum_{l\in\partial k^-}\alpha_{kl}n_l)+\lambda\big]P_t(\vec{n})\Big\}+$$
$$+T\sum_{k\neq i}\int_0^\infty dn_k\frac{\partial^2}{\partial n_k^2}\big\{n_k P_t(n_i,n_k)\big\}\,. \tag{S1}$$

Assuming that $\lim_{n_k\to\infty} n_k^2\, P_t(\vec{n}) = 0$, this equation can be simplified to:

$$\frac{\partial}{\partial t}P_t(n_i) = -\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i)+\lambda\big]P_t(n_i)\Big\} + T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i)\big\}$$
$$+\frac{\partial}{\partial n_i}\Big\{n_i\sum_{j\in\partial i^-}\alpha_{ij}\int_0^\infty dn_j\, n_j P_t(n_i,n_j)\Big\}$$
$$+\sum_{k\neq i}\lambda\lim_{n_k\to 0^+}P_t(n_i,n_k) - T\sum_{k\neq i}\lim_{n_k\to 0^+}P_t(n_i,n_k)\,. \tag{S2}$$

Now, the boundary conditions that we imposed on the Fokker-Planck equation become useful. We should enforce that, for each species k, the relation $(T-\lambda)\lim_{n_k\to 0^+}P_\infty(n_k) = 0$. Because of this, the last two terms in Eq. (S2) cancel, leaving:

$$\frac{\partial}{\partial t}P_t(n_i) = -\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i)+\lambda\big]P_t(n_i)\Big\} + T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i)\big\}$$
$$+\frac{\partial}{\partial n_i}\Big\{n_i\sum_{j\in\partial i^-}\alpha_{ij}\int_0^\infty dn_j\, n_j P_t(n_i,n_j)\Big\}\,. \tag{S3}$$

Defining the conditional average

$$m_{j\to i}(n_i,t) \equiv \int_0^\infty dn_j\, n_j P_t(n_j \mid n_i), \tag{S4}$$

it is possible to write:

$$\int_0^\infty dn_j\, n_j P_t(n_i,n_j) = m_{j\to i}(n_i,t)\, P_t(n_i) \tag{S5}$$

and

$$\frac{\partial}{\partial t}P_t(n_i) = -\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\, m_{j\to i}(n_i,t))+\lambda\big]P_t(n_i)\Big\} + T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i)\big\}\,. \tag{S6}$$

Since one does not know the shape of $m_{j\to i}(n_i,t)$, Eq. (S6) cannot be solved directly. Its stationary solutions must fulfill the equation:

$$0=-\frac{d}{dn_i}\Big\{\big[n_i(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\,m_{j\to i}(n_i))+\lambda\big]P_\infty(n_i)\Big\}+T\frac{d^2}{dn_i^2}\big\{n_iP_\infty(n_i)\big\}\,, \tag{S7}$$

where $m_{j\to i}(n_i)$ is the stationary value of the conditional average $m_{j\to i}(n_i,t)$ and $P_\infty(n_i)$ is the stationary distribution of the i-th species abundance. Integrating once over n_i, gives:

$$\frac{d}{dn_i}P_\infty(n_i)=\Big[\beta(1-n_i-\sum_{j\in\partial i^-}\alpha_{ij}\,m_{j\to i}(n_i))+\frac{\beta\lambda-1}{n_i}\Big]P_\infty(n_i)\,, \tag{S8}$$

where $\beta=1/T$. This equation has the following formal solution:

$$P_\infty(n_i)=\frac{1}{Z_i}\,n_i^{\beta\lambda-1}\,\exp\Big\{-\frac{\beta}{2}(n_i^2-2n_i)\Big\}\exp\Big\{-\beta\sum_{j\in\partial i^-}\alpha_{ij}\int_0^{n_i}dx\,m_{j\to i}(x)\Big\}\,. \tag{S9}$$

Although Eq. (S9) gives a simple expression that will be useful in the future, it is not a solution to the problem because we still do not have the functions $m_{j\to i}(n_i)$. But to obtain $m_{j\to i}(n_i)$ one needs to compute the conditional probabilities, or equivalently, the pair probabilities $P_t(n_i,n_j)$. The corresponding differential equation can be obtained following a similar procedure.

$$\begin{aligned}\frac{\partial}{\partial t}P_t(n_i,n_j)=&-\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i-\alpha_{ij}n_j)+\lambda\big]P_t(n_i,n_j)\Big\}+T\frac{\partial^2}{\partial n_i^2}\big\{n_iP_t(n_i,n_j)\big\}\\&-\frac{\partial}{\partial n_j}\Big\{\big[n_j(1-n_j-\alpha_{ji}n_i)+\lambda\big]P_t(n_i,n_j)\Big\}+T\frac{\partial^2}{\partial n_j^2}\big\{n_jP_t(n_i,n_j)\big\}\\&+\frac{\partial}{\partial n_i}\Big\{n_i\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\int_0^\infty dn_k\,n_kP_t(n_k,n_i,n_j)\Big\}\\&+\frac{\partial}{\partial n_j}\Big\{n_j\sum_{l\in\partial j^-\setminus i}\alpha_{jl}\int_0^\infty dn_l\,n_lP_t(n_i,n_j,n_l)\Big\}\,.\end{aligned} \tag{S10}$$

Similarly as before, let us define the conditional average $m_{k\to i,j}(n_i,n_j,t)\equiv\int_0^\infty dn_kP_t(n_k,t\mid n_i,n_j)$. Eq. (S10) reduces to:

$$\begin{aligned}\frac{\partial}{\partial t}P_t(n_i,n_j)=&-\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i-\alpha_{ij}n_j-\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i,j}(n_i,n_j))+\lambda\big]P_t(n_i,n_j)\Big\}\\&-\frac{\partial}{\partial n_j}\Big\{\big[n_j(1-n_j-\alpha_{ji}n_i-\sum_{l\in\partial j^-\setminus i}\alpha_{jl}\,m_{l\to i,j}(n_i,n_j))+\lambda\big]P_t(n_i,n_j)\Big\}\\&+T\frac{\partial^2}{\partial n_i^2}\big\{n_iP_t(n_i,n_j)\big\}+T\frac{\partial^2}{\partial n_j^2}\big\{n_jP_t(n_i,n_j)\big\}\,.\end{aligned} \tag{S11}$$

Again, one does not know the shape of $m_{k\to i,j}(n_i, n_j, t)$, but to obtain it, one needs to solve the differential equation for $P_t(n_k, n_i, n_j)$. This process builds a hierarchy of differential equations that, in the end, goes back to the Fokker-Planck equation for the full system (Eq. (2) in the main text). This is not surprising, since no approximations have been done, the solution of these local Fokker-Planck equations must be as difficult as the solution of the original Fokker-Planck equation for the whole system.

To close the hierarchy at some point, one can make some factorization of the joint probability densities so that they are expressed in terms of probability densities of a lower level in the hierarchy. The first approximation one could do is known in the studies of epidemic spreading as Individual Based Mean Field (IBMF) [1, 2] and it is simply:

$$P_t(n_i, n_j) \approx P_t(n_i)\, P_t(n_j)\,, \tag{S12}$$

which implies that:

$$\begin{aligned} m_{j\to i}(n_i, t) &= \int_0^\infty dn_j\, n_j P_t(n_j \mid n_i) = \int_0^\infty dn_j\, n_j \frac{P_t(n_i, n_j)}{P_t(n_i)} \\ m_{j\to i}(n_i, t) &\approx \int_0^\infty dn_j\, n_j\, \frac{P_t(n_i)\, P_t(n_j)}{P_t(n_i)} = \int_0^\infty dn_j\, n_j P_t(n_j) \equiv m_j(t)\,, \end{aligned} \tag{S13}$$

where $m_j(t)$ is the expected value of the abundance n_j at time t.

The IBMF differential equation is then:

$$\frac{\partial}{\partial t} P_t(n_i) = -\frac{\partial}{\partial n_i}\Big\{ \big[n_i(1 - n_i - \sum_{j\in\partial i^-} \alpha_{ij}\, m_j(t)) + \lambda\big] P_t(n_i) \Big\} + T \frac{\partial^2}{\partial n_i^2}\big\{ n_i P_t(n_i) \big\}\,, \tag{S14}$$

which must be completed with the definition $m_j(t) \equiv \int_0^\infty dn_j\, n_j\, P_t(n_j)$.

The IBMF is the first local closure one can provide. However, it is possible to go ahead and propose closures that stop at higher levels of the hierarchy. To express $P_t(n_k, n_i, n_j)$ in terms of pair probability densities like $P_t(n_i, n_j)$, one could use information about the actual graph of interactions. In the third line of Eq. (S10), the probability density $P_t(n_k, n_i, n_j)$ is actually defined in a graph $G(V, E)$ that contains the edges $k \to i$ and $i \to j$, but not necessarily the edges $k \to j$ or $j \to k$. In tree-like graphs, with high probability when the number of species is large, given that the edges $k \to i$ and $i \to j$ are present, we will have that k and j are not directly connected by any edge. Moreover, the length of the cycles diverges with the system size N, so the only short path that connects k and j necessarily passes through i. So, expecting that it works especially well in tree-like graphs, one could propose:

$$P_t(n_k, n_i, n_j) \approx \frac{P_t(n_k, n_i)\, P_t(n_i, n_j)}{P_t(n_i)} \equiv P_t(n_k \mid n_i)\, P_t(n_i)\, P_t(n_j \mid n_i)\,, \tag{S15}$$

This is known as Pair Based Mean Field (PBMF) in the above-mentioned context of epidemic spreading [3, 4] and is a factorization of the conditional measure. Given n_i at

time t, one assumes n_j and n_k are independent. This is not strange for those who know the replica symmetric cavity method, whose algorithmic counterpart is Belief Propagation (BP). However, even in the cases where BP is the exact solution in equilibrium (provided also that equilibrium exists), this approximation is not necessarily valid for the dynamics, since temporal correlations can still forbid the factorization, and n_j and n_k are not independent given only the value of n_i at time t.

Anyways, one can assume Eq. (S15) to be valid and continue:

$$\begin{aligned} m_{k\to i,j}(n_i,n_j,t) &= \int_0^\infty dn_k\, n_k P_t(n_k \mid n_i,n_j) = \int_0^\infty dn_k\, n_k \frac{P_t(n_k,n_i,n_j)}{P_t(n_i,n_j)} \\ m_{k\to i,j}(n_i,n_j,t) &\approx \int_0^\infty dn_k\, n_k \frac{P_t(n_k,n_i)}{P_t(n_i)} \equiv m_{k\to i}(n_i,t)\,. \end{aligned} \tag{S16}$$

Thus, the PBMF differential equation is:

$$\begin{aligned} \frac{\partial}{\partial t}P_t(n_i,n_j) = &-\frac{\partial}{\partial n_i}\Big\{\big[n_i(1-n_i-\alpha_{ij}n_j - \sum_{k\in\partial i^-\setminus j}\alpha_{ik}\, m_{k\to i}(n_i,t)) + \lambda\big]P_t(n_i,n_j)\Big\} \\ &-\frac{\partial}{\partial n_j}\Big\{\big[n_j(1-n_j-\alpha_{ji}n_i - \sum_{l\in\partial j^-\setminus i}\alpha_{jl}\, m_{l\to j}(n_j,t)) + \lambda\big]P_t(n_i,n_j)\Big\} \\ &+T\frac{\partial^2}{\partial n_i^2}\big\{n_i P_t(n_i,n_j)\big\} + T\frac{\partial^2}{\partial n_j^2}\big\{n_j P_t(n_i,n_j)\big\}\,, \end{aligned} \tag{S17}$$

which must be completed with the definition $m_{j\to i}(n_i,t) \equiv \int_0^\infty dn_j\, n_j P_t(n_j \mid n_i)$.

Both closures, IBMF and PBMF, are in principle solvable, at least numerically. However, when the system is large this is a difficult task.

S2 Connections with known results

S2.1 Belief Propagation as stationary solution of Pair Based Mean Field

As mentioned before, the approximation in Eq. (S15) is valid in equilibrium in all the cases where BP is also valid, since BP respects this factorization of the conditional measure. It is no surprise then that the BP equations, introduced in Ref. [5] for symmetric interactions ($\alpha_{ij} = \alpha_{ji}$), are a stationary solution of PBMF equations for symmetric interactions. Nevertheless, proving it could be a useful exercise for the future.

BP's update rule, as presented in Ref. [5], is:

$$\eta_{i\to j}(N_i) = \frac{1}{z_{i\to j}}\frac{1}{N_i+\Delta}\exp\Big\{-\frac{\beta}{2}(N_i^2-2N_i)\Big\}\prod_{k\in\partial i^-\setminus j}\sum_{N_k=0}^{\infty}\eta_{k\to i}(N_k)\exp\Big\{-\beta\alpha_{ik}N_i\,N_k\Big\}\,, \tag{S18}$$

where the constants r and K are set to one for simplicity. The constant $z_{i\to j}$ is a normalization factor, $\Delta > 0$ is a small parameter and $\eta_{i\to j}(N_i)$ are BP's messages. One must remember that in this case $\alpha_{ik} = \alpha_{ki}$, so one could write one or the other and Eq. (S18) remains valid.

First of all, one must notice that Eq. (S18) is written in a discretized space, so N_i is actually taking discrete values. That is the reason why in Eq. (S18) one sees a sum over N_k and not an integral. It is always possible to recast this expression in terms of the continuous variables n_i. Besides substituting sums by integrals, one has to pay particular attention to the factor $(N_i+\Delta)^{-1}$. As λ in the continuous case, here Δ ensures that the divergence when N_i goes to zero is not so critical that it makes the probability distribution non-normalizable. Any positive Δ will remove this divergence at $N_i = 0$. Thus, in the continuous version, one must simply substitute the factor $(N_i+\Delta)^{-1}$ by the analogous factor $n_i^{\beta\lambda-1}$, present in Eqs. 3 and 8 in the main text.

The resulting update rule is:

$$\eta_{i\to j}(n_i) = \frac{1}{z_{i\to j}}\, n_i^{\beta\lambda-1}\, \exp\Big\{-\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \prod_{k\in\partial i^-\setminus j} \int_0^\infty dn_k\, \eta_{k\to i}(n_k) \exp\big\{-\beta\alpha_{ik} n_i\, n_k\big\}\,, \tag{S19}$$

It will be helpful to look closely at the integral

$$\mathcal{Z}_{k\to i}(n_i) = \int_0^\infty dn_k\, \eta_{k\to i}(n_k) \exp\big\{-\beta\alpha_{ik} n_i\, n_k\big\}\,. \tag{S20}$$

In understanding its meaning, it is useful to write the stationary pair probability density in terms of BP's messages:

$$P_\infty(n_i, n_k) = \eta_{i\to k}(n_i)\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i\, n_k}\,. \tag{S21}$$

Then, the conditional probability density is

$$\begin{aligned} P_\infty(n_k \mid n_i) &= \frac{\eta_{i\to k}(n_i)\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i\, n_k}}{\int_0^\infty dn_k\, \eta_{i\to k}(n_i)\, \eta_{k\to i}(n_k) e^{-\beta\alpha_{ik} n_i\, n_k}} = \frac{\eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i\, n_k}}{\int_0^\infty dn_k\, \eta_{k\to i}(n_k) e^{-\beta\alpha_{ik} n_i\, n_k}} \\ P_\infty(n_k \mid n_i) &= \frac{1}{\mathcal{Z}_{k\to i}(n_i)}\, \eta_{k\to i}(n_k)\, e^{-\beta\alpha_{ik} n_i\, n_k}\,. \end{aligned} \tag{S22}$$

So $\mathcal{Z}_{k\to i}(n_i)$ is the normalization factor of the conditional probability density. Furthermore, its derivative with respect to n_i is:

$$\begin{aligned} \frac{\partial}{\partial n_i}\mathcal{Z}_{k\to i}(n_i) &= -\beta\,\alpha_{ik} \int_0^\infty dn_k\, n_k\, \eta_{k\to i}(n_k) e^{-\beta\alpha_{ik} n_i\, n_k} \\ \frac{\partial}{\partial n_i}\mathcal{Z}_{k\to i}(n_i) &= -\beta\,\alpha_{ik} \mathcal{Z}_{k\to i}(n_i)\, \frac{\int_0^\infty dn_k\, n_k\, \eta_{k\to i}(n_k) e^{-\beta\alpha_{ik} n_i\, n_k}}{\mathcal{Z}_{k\to i}(n_i)}\,. \end{aligned} \tag{S23}$$

Remembering the definition of $m_{k\to i}(n_i)$ and Eq. (S21), one gets

$$\frac{\partial}{\partial n_i}\mathcal{Z}_{k\to i}(n_i) = -\beta\,\alpha_{ik}\,\mathcal{Z}_{k\to i}(n_i)\,m_{k\to i}(n_i)\,, \tag{S24}$$

from where it is easy to see that

$$\mathcal{Z}_{k\to i}(n_i) \propto \exp\Big\{-\beta\,\alpha_{ik}\int_0^{n_i} dx\, m_{k\to i}(x)\Big\}\,. \tag{S25}$$

Returning to Eq. (S19), the expression can be rewritten in a form that is very similar to the formal solution for $P_\infty(n_i)$ in Eq. (S9), and that is connected with the differential equation for PBMF (Eq. (S17)):

$$\eta_{i\to j}(n_i) = \frac{1}{z_{i\to j}}\,n_i^{\beta\lambda-1}\exp\Big\{-\frac{\beta}{2}(n_i^2-2n_i)\Big\}\exp\Big\{-\beta\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\int_0^{n_i} dx\, m_{k\to i}(x)\Big\}\,. \tag{S26}$$

Thus, the pair probability can be written as:

$$\begin{aligned}P_\infty(n_i,n_j) = \frac{1}{Z_{ij}}(n_i\,n_j)^{\beta\lambda-1}\,e^{-\frac{\beta}{2}(n_i^2-2n_i)}\,e^{-\frac{\beta}{2}(n_j^2-2n_j)}\,e^{-\beta\alpha_{ij}n_i\,n_j}\,\times\\ \times\exp\Big\{-\beta\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\int_0^{n_i} dx\, m_{k\to i}(x)\Big\}\exp\Big\{-\beta\sum_{l\in\partial j^-\setminus i}\alpha_{jl}\int_0^{n_j} dx\, m_{l\to j}(x)\Big\}\,.\end{aligned} \tag{S27}$$

Everything is now ready to use the update rule and obtain an expression for the derivative of the probability density $P_\infty(n_i,n_j)$ with respect to n_i:

$$\frac{\partial}{\partial n_i}P_\infty(n_i,n_j) = P_\infty(n_i,n_j)\Big[\frac{\beta\lambda-1}{n_i} - \beta(n_i-1) - \beta\,\alpha_{ij}\,n_j - \beta\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}(n_i)\Big]\,. \tag{S28}$$

Therefore

$$\begin{aligned}&T\frac{\partial}{\partial n_i}\big\{n_iP_t(n_i,n_j)\big\} - \Big[n_i(1-n_i-\alpha_{ij}n_j - \sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}(n_i)) + \lambda\Big]P_t(n_i,n_j)\\ &= P_t(n_i,n_j)\Big\{T + T\,n_i\Big[\frac{\beta\lambda-1}{n_i} - \beta(n_i-1) - \beta\,\alpha_{ij}\,n_j - \beta\sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}(n_i)\Big] -\\ &\qquad -\Big[n_i(1-n_i-\alpha_{ij}n_j - \sum_{k\in\partial i^-\setminus j}\alpha_{ik}\,m_{k\to i}(n_i)) + \lambda\Big]\Big\} = \mathbf{0}\,.\end{aligned} \tag{S29}$$

Something analogous happens with the derivative with respect to n_j. This means that, in the open region $(n_i,n_j)\in(0,+\infty)\times(0,+\infty)$, the right hand side of Eq. (S17) is equal to zero. In other words, BP is a stationary solution for PBMF when the interactions are symmetric.

S2.2 Zero temperature and large connectivity limit of IBMF

So far, we have presented two stationary distributions that are approximate solutions to the problem. First, IBMF, which is the simplest local closure one can do, but at the same time is of a general purpose. Regardless of the degree of symmetry in the interactions, the stationary solution of IBMF is:

$$P_\infty(n_i) = \frac{1}{Z_i}\, n_i^{\beta\lambda-1}\, \exp\Big\{ -\frac{\beta}{2}\,(n_i - h_i)^2 \Big\}\,, \tag{S30}$$

with $h_i = 1 - \sum_{j\in\partial i^-} \alpha_{ij}\, m_j$ and $m_j = \int_0^\infty dn_j\, n_j\, P_\infty(n_j)$.

When taking the zero temperature limit (equivalently $\beta \to \infty$), it is crucial to know how to treat the parameter λ. One must remember that λ has the role of ensuring that the distribution is normalizable, by avoiding the extinction of all the species. It is usually interpreted as an immigration rate that must be taken positive, but small, to study the phenomenology of the model in its purest version possible. Therefore, one should take $\beta \to \infty$ but keep $\beta\lambda$ finite, since $\lambda \sim 0$.

We will compare the final result of taking this limit with Eq. (13) in Ref. [6], which gives an expression for the steady state abundance of the fully connected model at $T = 0$:

$$n = \max\left(0,\, \frac{1 - \hat{\mu}\, m_\infty - \hat{\sigma}\, \zeta_\infty}{1 - \epsilon\, \hat{\sigma}^2\, \chi_{\rm int}}\right). \tag{S31}$$

Remember that μ is the average value of the couplings α_{ij}, and σ^2 is its variance. When the system is fully connected, it is necessary to rescale these parameters and send them to zero when $N \to \infty$, but keeping finite $\hat{\mu} = N\,\mu$ and $\hat{\sigma}^2 = N\,\sigma^2$. The number $\epsilon \in [-1, 1]$ is the correlation between α_{ij} and α_{ji}, or, in other words, the level of symmetry in the couplings. Besides these parameters, in Eq. (S31) one has: m_∞, which is the expected value of n; ζ_∞, which is a random Gaussian variable; and $\chi_{\rm int}$, which is the integrated response of n to the effect of a small external field.

To recover this result, one can start by taking the limit $T \to 0$ in Eq. (S30). As we discussed in Section IV of the main text, the probability density then concentrates at one point:

$$n_i = \max\left(0, 1 - \sum_{j\in\partial i^-} \alpha_{ij}\, m_j\right). \tag{S32}$$

This way of writing it ensures that n_i must always be non-negative. When the connectivity is large, the sum $\theta = \sum_{j\in\partial i^-} \alpha_{ij}\, m_j$ is distributed as a Gaussian with the following first and second moments:

$$\langle\theta\rangle = \sum_{j\in\partial i^-} \langle \alpha_{ij}\, m_j\rangle \tag{S33}$$

$$\langle\theta^2\rangle = \sum_{j\in\partial i^-} \sum_{k\in\partial i^-} \langle \alpha_{ij}\, \alpha_{ik}\, m_j\, m_k\rangle\ . \tag{S34}$$

The averages in Eqs. (S33) and (S34) are taken over the disorder in the interactions, *i.e.*, over the possible values of the whole matrix of couplings $\overleftrightarrow{\alpha}$ with elements α_{ij}. Then we must realize that m_j is actually a function of $\overleftrightarrow{\alpha}$, and that it is not possible to compute the averages in Eqs. (S33) and (S34) as if m_j were independent of the specific realization of the disorder $\overleftrightarrow{\alpha}$. However, one can explicitly average Eqs. (S33) and (S34) over part of the disorder since most of the elements of $\overleftrightarrow{\alpha}$ are indeed independent of α_{ij}:

$$\langle\theta\rangle = \sum_{j\in\partial i^-} \langle\, \alpha_{ij}\, m_j(\overleftrightarrow{\alpha})\rangle = \sum_{j\in\partial i^-} \langle\, \alpha_{ij}\, m_j(\alpha_{ij},\alpha_{ji})\,\rangle \tag{S35}$$

$$\langle\theta^2\rangle = \sum_{j\in\partial i^-}\sum_{k\in\partial i^-} \langle\, \alpha_{ij}\,\alpha_{ik}\, m_j(\alpha_{ij},\alpha_{ji},\alpha_{ik},\alpha_{ki})\, m_k(\alpha_{ij},\alpha_{ji},\alpha_{ik},\alpha_{ki})\,\rangle\ . \tag{S36}$$

The dependence of m_j on the disorder has been reduced to the only elements that are not independent of α_{ij}, which are α_{ij} itself and possibly α_{ji} (if $\epsilon > 0$). The new $m_j(\alpha_{ij},\alpha_{ji})$ is already an average over almost all the couplings, except for the ones associated with the edges $i\to j$ and $j\to i$. Therefore, the right-hand sides of Eqs. (S35) and (S36) are no longer averages over all the couplings that appear in the arguments of the functions. More specifically, the average $\langle\, \alpha_{ij}\, m_j(\alpha_{ij},\alpha_{ji})\,\rangle$ in Eq. (S35) is taken over the values of α_{ij} and α_{ji}, considering that they can be correlated. In the same way, the average in Eq. (S36) is taken over the values of α_{ij}, α_{ji}, α_{ik}, and α_{ki}. To compute those averages, it will be useful to write the couplings in terms of standardized Gaussians $x_{ij}\sim\mathcal{N}(0,1)$ as:

$$\begin{aligned} &\alpha_{ij} = \mu + \sigma\, x_{ij}, \quad \text{with} \\ &\langle x_{ij}\rangle = 0; \qquad \langle x_{ij}x_{kl}\rangle = \delta_{ij,kl} + \epsilon\delta_{ij,lk}\ . \end{aligned} \tag{S37}$$

Then $m_j(\alpha_{ij},\alpha_{ji})$ becomes a function of x_{ij} and x_{ji}. However, when the connectivity is large, it is safe to assume that the dependence of $m_j(x_{ij},x_{ji})$ on each x_{ij}, associated to only one of the many edges that contain the node j, is very weak. Taking an expansion of $m_j(x_{ij},x_{ji})$ in powers of x_{ij} and x_{ji} up to the first order leads to:

$$m_j(x_{ij},x_{ji}) \approx m_j(0,0) + x_{ij}\Big\{\frac{\partial}{\partial x_{ij}} m_j(x_{ij},0)\Big\}\Big|_{x_{ij}=0} + x_{ji}\Big\{\frac{\partial}{\partial x_{ji}} m_j(0,x_{ji})\Big\}\Big|_{x_{ji}=0}\ . \tag{S38}$$

Remembering that x_{ij} actually does not appear in the equation for dn_j/dt, but in the equation for dn_i/dt, one realizes that its effect in $m_j(x_{ij},x_{ji})$ must be even weaker than the effect of x_{ji}. Finally:

$$m_j(x_{ij},x_{ji}) \approx m_j(x_{ji}) \approx m_j(0) + x_{ji}\Big\{\frac{\partial}{\partial x_{ji}} m_j(x_{ji})\Big\}\Big|_{x_{ji}=0}\ . \tag{S39}$$

Inserting Eqs. (S37) and (S39) back into Eq. (S35) gives:

$$\langle\theta\rangle = \sum_{j\in\partial i^-} \Big\langle \left(\mu + \sigma\, x_{ij}\right) \Big(m_j(0) + x_{ji} \Big\{ \frac{\partial}{\partial x_{ji}} m_j(x_{ji}) \Big\} \Big|_{x_{ji}=0} \Big) \Big\rangle$$

$$\langle\theta\rangle = \mu \sum_{j\in\partial i^-} \langle m_j(0)\rangle + \sigma \sum_{j\in\partial i^-} \Big\langle x_{ij} x_{ji} \Big\{ \frac{\partial}{\partial x_{ji}} m_j(x_{ji}) \Big\} \Big|_{x_{ji}=0} \Big\rangle \ . \tag{S40}$$

Let the connectivity be equal to c. After identifying that $\langle m_j(0)\rangle \equiv m_\infty$, where m_∞ is the parameter in Eq. (S31) above, and by using the expression for the second moment in Eq. (S37), one gets:

$$\langle\theta\rangle = \mu\, c\, m_\infty + \sigma\, c\, \epsilon \Big\langle \Big\{ \frac{\partial}{\partial x_{ji}} m_j(x_{ji}) \Big\} \Big|_{x_{ji}=0} \Big\rangle \ . \tag{S41}$$

The derivative in Eq. (S41) gives the average response of m_j to a small perturbation in the couplings. This will be related to the response to a small external field, and to see the relation more explicitly, it is convenient to recast the differential equation for n_j (see Eq. 1 in the main text) in terms of x_{ji}:

$$\frac{dn_j}{dt} = -n_j(1 - n_j - \mu \sum_{l\in\partial j^-} n_l - \sigma \sum_{l\in\partial j^-} x_{jl}\, n_l) + \xi_j(t) + \lambda \ . \tag{S42}$$

Taking i outside the sum over the neighbors:

$$\frac{dn_j}{dt} = -n_j(1 - n_j - \mu \sum_{l\in\partial j^-} n_l - \sigma \sum_{l\in\partial j^-\setminus i} x_{jl}\, n_l - \sigma\, x_{ji}\, n_i) + \xi_j(t) + \lambda \ . \tag{S43}$$

The effect of a small perturbation x_{ji} is the same as the one provoked by a small external field with value $h_j(t) = -\sigma\, x_{ji}\, n_i(t)$. The two responses must be related by:

$$\Big\langle \Big\{ \frac{\partial}{\partial x_{ji}} m_j(x_{ji}) \Big\} \Big|_{x_{ji}=0} \Big\rangle = \Big\langle \Big\{ \frac{\partial h_j}{\partial x_{ji}} \frac{\partial}{\partial h_j} m_j(h_j) \Big\} \Big|_{h_j=0} \Big\rangle = -\sigma\, n_i\, \chi_{\rm int} \ . \tag{S44}$$

Inserting Eq. (S44) back into Eq. (S41), one finally obtains an expression for $\langle\theta\rangle$ in terms of m_∞ and $\chi_{\rm int}$, which are two of the parameters in the known result from DMFT (see Eq. (S31)):

$$\langle\theta\rangle = \mu\, c\, m_\infty - \sigma^2\, c\, \epsilon\, n_i\, \chi_{\rm int} \ . \tag{S45}$$

The same needs to be done with the second moments:

$$\langle\theta^2\rangle = \sum_{j\in\partial i^-}\sum_{k\in\partial i^-}\left\langle\left(\mu+\sigma\,x_{ij}\right)\left(\mu+\sigma\,x_{ik}\right)\left(m_\infty-\sigma\,n_i\,\chi_{\rm int}\,x_{ji}\right)\left(m_\infty-\sigma\,n_i\,\chi_{\rm int}\,x_{ki}\right)\right\rangle$$

$$\langle\theta^2\rangle = \sum_{j\in\partial i^-}\sum_{k\in\partial i^-}\Big\langle\left[\mu^2+\mu\,\sigma\left(x_{ij}+x_{ik}\right)+\sigma^2\,x_{ij}\,x_{ik}\right]\times$$
$$\times\left[m_\infty^2-m_\infty\sigma\,n_i\,\chi_{\rm int}\left(x_{ji}+x_{ki}\right)+\sigma^2\,n_i^2\,\chi_{\rm int}^2\,x_{ji}\,x_{ki}\right]\Big\rangle$$

$$\langle\theta^2\rangle = c^2\,\mu^2\,m_\infty^2+c\,\sigma^2\,m_\infty^2-\mu\,\sigma^2\,m_\infty\,n_i\,\chi_{\rm int}\sum_{j\in\partial i^-}\sum_{k\in\partial i^-}\left\langle\left(x_{ij}+x_{ik}\right)\left(x_{ji}+x_{ki}\right)\right\rangle+$$
$$+c\,\mu^2\sigma^2\,n_i^2\chi_{\rm int}^2+\sigma^4\,n_i^2\,\chi_{\rm int}^2\sum_{j\in\partial i^-}\langle x_{ij}^2\,x_{ji}^2\rangle+\sigma^4\,n_i^2\,\chi_{\rm int}^2\sum_{j\in\partial i^-}\sum_{k\in\partial i^-\setminus j}\langle x_{ij}\,x_{ji}\,x_{ik}\,x_{ki}\rangle$$

$$\langle\theta^2\rangle = c^2\,\mu^2\,m_\infty^2+c\,\sigma^2\,m_\infty^2-\mu\,\sigma^2\,m_\infty\,n_i\,\epsilon\,\chi_{\rm int}\left(4c+2c(c-1)\right)+$$
$$+c\,\mu^2\sigma^2\,n_i^2\chi_{\rm int}^2+\sigma^4\,n_i^2\,\chi_{\rm int}^2\left(c+c(c-1)\epsilon^2\right)$$

$$\langle\theta^2\rangle = c^2\,\mu^2\,m_\infty^2+c\,\sigma^2\,m_\infty^2-2\,\mu\,\sigma^2\,m_\infty\,n_i\,\epsilon\,\chi_{\rm int}\,c\,(c+1)+$$
$$+c\,\mu^2\sigma^2\,n_i^2\chi_{\rm int}^2+\sigma^4\,n_i^2\,\chi_{\rm int}^2\left(c+c(c-1)\epsilon^2\right)\quad. \tag{S46}$$

where we assumed that $\langle x_{ij}^2\,x_{ji}^2\rangle = 1$ and that $\langle x_{ij}\,x_{ji}\,x_{ik}\,x_{ki}\rangle = \epsilon^2$, if $i\neq k$.

Then, the variance of θ is:

$$s_\theta^2 = \langle\theta^2\rangle-\langle\theta\rangle^2$$
$$s_\theta^2 = c\,\sigma^2\,m_\infty^2-2\,\mu\,\sigma^2\,m_\infty\,n_i\,\epsilon\,\chi_{\rm int}\,c+c\mu^2\sigma^2n_i^2\chi_{\rm int}+\sigma^4\,n_i^2\,\chi_{\rm int}^2\,c\,(1-\epsilon^2)\quad. \tag{S47}$$

Now, to complete the procedure, one needs to rescale the average μ and the variance σ^2 of the couplings for $\langle\theta\rangle$ and $\langle\sigma^2\rangle$ to remain finite. The natural choice is

$$\hat{\mu} = c\,\mu \qquad , \qquad \hat{\sigma}^2 = c\,\sigma^2\quad. \tag{S48}$$

After neglecting the terms that go to zero when $c\to\infty$, we finally obtain:

$$\langle\theta\rangle = \hat{\mu}\,m_\infty-\hat{\sigma}^2\,\epsilon\,n_i\,\chi_{\rm int} \tag{S49}$$
$$s_\theta^2 = \hat{\sigma}^2\,m_\infty^2\quad. \tag{S50}$$

Remembering Eq. (S32), the stationary abundance for large connectivity and at $T=0$ was:

$$n_i = \max\left(0,1-\sum_{j\in\partial i^-}\alpha_{ij}\,m_j(\infty)\right)\equiv\max(0,\,1-\theta) = \max(0,\,1-\langle\theta\rangle-s_\theta\,\nu)\quad. \tag{S51}$$

where $\nu\sim\mathcal{N}(0,1)$ is a standardized Gaussian. Substituting Eqs. (S49) and (S50) into this expression, we get:

$$n_i = \max(0,\,1-\hat{\mu}\,m_\infty+\hat{\sigma}^2\,\epsilon\,n_i\,\chi_{\rm int}-\hat{\sigma}\,m_\infty\,\nu)\quad. \tag{S52}$$

Rearranging terms, we finally recover Eq. (13) of [6]:

$$n_i = \max\left(0,\, \frac{1 - \hat{\mu}\, m_\infty - \sigma\, \zeta_\infty}{1 - \epsilon\, \hat{\sigma}^2\, \chi_{\rm int}}\right) , \tag{S53}$$

where ζ_∞ is a Gaussian with zero mean $\langle \zeta_\infty \rangle = 0$ and variance $\langle \zeta_\infty^2 \rangle = m_\infty$. The latter is also consistent with the third line of Eq. (12) in Ref. [6].

S3 Efficient implementation of IBMF

To use IBMF to obtain the actual values of the averages m_i, one needs to numerically compute integrals of the form

$$I_k(\beta, \lambda, M) = \int_0^\infty dn\, n^{\beta\lambda - 1 + k}\, \exp\Big\{ -\frac{\beta}{2}\,(n - M)^2 \Big\} \ , \tag{S54}$$

with the parameter k taking the value $k = 0, 1$ in our case.

Luckily, the integral (S54) can be expressed in terms of known special functions, called parabolic cylinder functions. They have the integral representation (see 9.241 in Ref. [7]):

$$D_{-p}(z) = \frac{e^{-z^2/4}}{\Gamma(p)} \int_0^\infty dx\, e^{-xz - x^2/2} x^{p-1} \ , \tag{S55}$$

which is valid for $\mathrm{Re}[p] > 0$ and where $\Gamma(-p)$ is the Euler's Gamma function. Looking back at Eq. (S54), we can change variables making $x = n\sqrt{\beta}$ to get

$$I_k = \Big[\frac{1}{\beta}\Big]^{\frac{\beta\lambda + k}{2}} e^{-\frac{\beta M^2}{2}} \int_0^\infty dx\, x^{\beta\lambda - 1 + k} \exp\Big\{ -\frac{x^2}{2} + \sqrt{\beta\, M^2}\, x\, M \Big\} \ . \tag{S56}$$

After comparing Eqs. (S55) and (S56), we identify $p = \beta\lambda + k$ and $z = -\sqrt{\beta\, M^2}$, and we can write the integral I_k in terms of the parabolic cylinder functions as:

$$I_k(\beta, \lambda, M) = \Big[\frac{1}{\beta}\Big]^{\frac{\beta\lambda + k}{2}} \exp\Big\{ -\frac{\beta M^2}{4} \Big\}\, \Gamma(\beta\lambda + k)\, D_{-\beta\lambda - k}\Big(-\sqrt{\beta M^2} \Big) \ . \tag{S57}$$

Eq. (S57) is very convenient since we can express the first moments of the distribution $P_\infty(n_j)$ in terms of these integrals. Indeed, one has:

$$m_j = \frac{I_1(\beta, \lambda, M_j)}{I_0(\beta, \lambda, M_j)} = \sqrt{\beta}\, \lambda\, \frac{D_{-\beta\lambda - 1}\big(-\sqrt{\beta M_j^2} \big)}{D_{-\beta\lambda}\big(-\sqrt{\beta M_j^2} \big)} \ , \tag{S58}$$

where $M_j = 1 - \sum_{k \in \partial j^-} \alpha_{jk}\, m_k$.

Using again Ref. [7] (9.240), we can write the parabolic cylinder functions in terms of the more practical Kummer's confluent hypergeometric function $\Phi(a, b; z)$, which can be found already tabulated in different programming languages. The relation is:

$$D_{-\beta\lambda-k}\big(-\sqrt{\beta M^2}\big) = \frac{2^{-(\beta\lambda+k)/2} e^{-\beta M^2/4}\sqrt{\pi}}{\Gamma\Big(\frac{\beta\lambda+1+k}{2}\Big)\,\Gamma\Big(\frac{\beta\lambda+k}{2}\Big)} \times \tag{S59}$$

$$\times\Big\{\Gamma\Big(\frac{\beta\lambda+k}{2}\Big)\,\Phi\Big[\frac{\beta\lambda+k}{2},\frac{1}{2},\frac{\beta M^2}{2}\Big] + \sqrt{2\beta M^2}\,\Gamma\Big(\frac{\beta\lambda+1+k}{2}\Big)\,\Phi\Big[\frac{\beta\lambda+1+k}{2},\frac{3}{2},\frac{\beta M^2}{2}\Big]\Big\}$$

Then, the first moment can be written as:

$$m_j = \frac{\Gamma\Big(\frac{\beta\lambda+1}{2}\Big)\,\Phi\Big[\frac{\beta\lambda+1}{2},\frac{1}{2},\frac{\beta M_j^2}{2}\Big] + \sqrt{\frac{\beta}{2}}\,M_j\,\beta\,\lambda\,\Gamma\Big(\frac{\beta\lambda}{2}\Big)\,\Phi\Big[1+\frac{\beta\lambda}{2},\frac{3}{2},\frac{\beta M_j^2}{2}\Big]}{\sqrt{\frac{\beta}{2}}\,\Gamma\Big(\frac{\beta\lambda}{2}\Big)\,\Phi\Big[\frac{\beta\lambda}{2},\frac{1}{2},\frac{\beta M_j^2}{2}\Big] + \beta M_j\,\Gamma\Big(\frac{\beta\lambda+1}{2}\Big)\,\Phi\Big[\frac{\beta\lambda+1}{2},\frac{3}{2},\frac{\beta M_j^2}{2}\Big]} \tag{S60}$$

We can then set an initial condition for all the m_i, with $i = 1,\dots,N$, and iterate Eq. (S60). This is computationally fast, and the result of the iteration process depends on the relevant parameters of the model (μ,σ,β) and on the interaction graph. The code is available at Ref. [8].

S4 Continuous BP equations

The continuous BP equations for the Generalized Lotka-Volterra (gLV) model, as introduced in Eq. (13) in the main text, read:

$$\eta_{i\to j}(n_i) = \frac{1}{z_{i\to j}}\,n_i^{\beta\lambda-1}\,\exp\Big\{-\frac{\beta}{2}(n_i^2-2n_i)\Big\}\prod_{k\in\partial i^-\setminus j}\Big(\int_0^\infty dn_k\,\eta_{k\to i}(n_k)\exp\big\{-\beta\alpha_{ik}n_i\,n_k\big\}\Big)\,. \tag{S61}$$

Eq. (S61) is the update rule for the cavity marginals, or messages, $\eta_{i\to j}(n_i)$, which are the probability distributions for the abundance of species (node) i once we cut the edges with species j. As usual, we call $\eta_{i\to j}(n_i)$ *node-to-link* messages. What we immediately notice from Eq. (S61) is that $\eta_{i\to j}(n_i)$ contains the factor $n_i^{\beta\lambda-1}$, which diverges in $n_i = 0$, for $\beta\lambda < 1$. To avoid such divergence, we introduce the *link-to-node* cavity messages, which we denote $\hat\eta_{i\to j}(n_j)$ and which represent the distribution of the abundance of species j once all its links, except the one with i, are cut. These cavity marginals are defined as:

$$\hat\eta_{i\to j}(n_j) \propto \int_0^\infty dn_i\,\eta_{i\to j}(n_i)\exp\big\{-\beta\alpha_{ij}n_i\,n_j\big\}\,. \tag{S62}$$

Inserting Eq. (S61) into Eq. (S62), we get:

$$\hat\eta_{i\to j}(n_j) \propto \int_0^\infty dn_i\,n_i^{\beta\lambda-1}\exp\Big\{-\frac{\beta}{2}(n_i^2-2n_i)\Big\}\exp\big\{-\beta\alpha_{ij}n_i\,n_j\big\}\prod_{k\in\partial i^-\setminus j}\hat\eta_{k\to i}(n_i)\,, \tag{S63}$$

which is now the update rule for the link-to-node cavity marginals $\hat\eta_{i\to j}(n_j)$. Importantly, once the messages $\hat\eta_{i\to j}(n_j)$ are determined, the cavity marginals $\eta_{i\to j}(n_i)$ follow directly, as we can see by rewriting Eq. (S61) as:

$$\eta_{i\to j}(n_i) \propto n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \prod_{k\in\partial i^-\setminus j} \hat{\eta}_{k\to i}(n_i). \tag{S64}$$

We can then analyze the BP convergence directly at the level of the link-to-node cavity messages $\hat{\eta}_{i\to j}(n_j)$. As said before, the advantage is that $\hat{\eta}_{i\to j}(n_j)$ is a well-behaved function that allows us to avoid the divergence of $\eta_{i\to j}(n_i)$ at $n_i = 0$.

If we initialize all the messages $\hat{\eta}_{i\to j}(n_j)$ to the uniform distribution over n_j, the update rule (S63) in the first iteration step becomes:

$$\hat{\eta}_{i\to j}(n_j) \propto \int_0^\infty dn_i\, n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}(n_i^2 - 2n_i)\Big\} \exp\big\{ -\beta\alpha_{ij} n_i\, n_j\big\}. \tag{S65}$$

The advantage of this initialization is that we can rewrite Eq. (S65) as

$$\hat{\eta}_{i\to j}(n_j) \propto \int_0^\infty dn_i n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}n_i^2 + \beta n_i(1-\alpha_{ij}n_j)\Big\} \tag{S66}$$

$$= \frac{1}{\beta^{\frac{\beta\lambda}{2}}} \int_0^\infty dx\, x^{-p-1} \exp\Big\{ -\frac{x^2}{2} - zx\Big\}, \tag{S67}$$

where in the last passage we introduced the change of variable $x = \sqrt{\beta}n_i$, with $p = -\beta\lambda$, $z = \sqrt{\beta}(\alpha_{ij}n_j - 1)$. The integral in Eq. (S67) can be exactly solved and expressed in terms of *parabolic cylinder functions* $D_p(z)$. In particular, as it is shown in Eq. 9.241 in Ref. [7], the integral (I) is equivalent to:

$$I = \frac{\Gamma(-p)}{e^{-\frac{z^2}{4}}} D_p(z), \tag{S68}$$

where $\Gamma(-p)$ is the Euler's Gamma function. The first iteration step for the update of the message $\hat{\eta}_{i\to j}(n_j)$ then corresponds to

$$\hat{\eta}_{i\to j}(n_j) \propto \beta^{-\frac{\beta\lambda}{2}}\, \Gamma(\beta\lambda)\, e^{\frac{\beta}{4}(\alpha_{ij}n_j-1)^2} D_{-\beta\lambda}\left(\sqrt{\beta}(\alpha_{ij}n_j - 1)\right), \tag{S69}$$

which then has to be normalized.

For the following iteration steps, instead, we go back to Eq. (S63), which we rewrite as

$$\hat{\eta}_{i\to j}(n_j) \propto \exp\Big\{\frac{\beta}{2}(1-\alpha_{ij}n_j)^2\Big\} \int_0^\infty dn_i\, n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}(n_i - 1 + \alpha_{ij}n_j)^2\Big\} \prod_{k\in\partial i^-\setminus j} \hat{\eta}_{k\to i}(n_i). \tag{S70}$$

In order to analyze the integral part of Eq. (S70), which we will denote by $I_{i\to j}(n_j)$, let us split it at some $\delta > 0$:

$$I_{i\to j}(n_j) = \int_0^\delta dn_i\, n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}(n_i - 1 + \alpha_{ij}n_j)^2\Big\} \prod_{k\in\partial i^-\setminus j} \hat{\eta}_{k\to i}(n_i) +$$

$$+ \int_\delta^\infty dn_i\, n_i^{\beta\lambda-1} \exp\Big\{ -\frac{\beta}{2}(n_i - 1 + \alpha_{ij}n_j)^2\Big\} \prod_{k\in\partial i^-\setminus j} \hat{\eta}_{k\to i}(n_i). \tag{S71}$$

For $\delta \ll 1$, the first contribution, i.e. the integral between 0 and δ, can be approximated as:

$$I^{\delta}_{i\to j}(n_j) = \exp\Big\{-\frac{\beta}{2}(1-\alpha_{ij}n_j)^2\Big\}\prod_{k\in\partial i^-\setminus j}\hat{\eta}_{k\to i}(0)\frac{\delta^{\beta\lambda}}{\beta\lambda}, \tag{S72}$$

where we performed the integration of the factor $n_i^{\beta\lambda-1}$ over the interval $[0,\delta]$, while approximating the remaining part of the integrand by its value at $n_i = 0$.

Substituting Eqs. (S71) and (S72) into Eq. (S70), we obtain the following BP update rule for the cavity messages $\hat{\eta}_{i\to j}$:

$$\hat{\eta}_{i\to j}(n_j) = \frac{1}{\hat{z}_{i\to j}}\Big[\frac{\delta^{\beta\lambda}}{\beta\lambda}\prod_{k\in\partial i^-\setminus j}\hat{\eta}_{k\to i}(0) + \\ + \int_{\delta}^{\infty} dn_i\, n_i^{\beta\lambda-1}\exp\Big\{-\frac{\beta}{2}(n_i^2-2n_i)\Big\}\exp\Big\{-\beta\alpha_{ij}n_in_j\Big\}\prod_{k\in\partial i^-\setminus j}\hat{\eta}_{k\to i}(n_i)\Big]. \tag{S73}$$

After having *initialized* the messages following Eq. (S69), the update rule of the continuous BP equations is given by Eq. (S73).

This approach is general, but we specifically used it in order to study the transition from a single-equilibrium phase (where BP converges) to a multiple-equilibria one (where BP does not converge) in the presence of homogeneous ($\sigma = 0$) and competitive interactions ($\mu > 0$). The results are in Subsection IV.C of the main text. In this case, the species abundances will only rarely be larger than the carrying capacities, which in our case are all equal to 1. Some thermal fluctuations could drive species to abundances slightly larger than 1, but we can safely restrict the integral $\int_{\delta}^{\infty}$ in Eq. (S73) to the interval $[\delta, 2]$. In our implementation, we set $\delta = 10^{-4}$ and compute the integral numerically.

Let us specify that in our analysis we used a sequential update, meaning that at each iteration step k, each $\hat{\eta}^{(k)}_{i\to j}(n_j)$ is updated to $\hat{\eta}^{(k+1)}_{i\to j}(n_j)$ asynchronously. In particular, the order of updates follows a random fixed sequence of directed edges $i \to j$. This sequential approach is essential to avoid convergence inconsistencies, ensuring that BP stops converging only once the multiple-equilibria phase is reached. These issues are also discussed in Appendix C of the main text. As for IBMF and numerical integration, the code is available at Ref. [8].

References

[1] Wang, Y., Chakrabarti, D., Wang, C., and Faloutsos, C. "Epidemic spreading in real networks: an eigenvalue viewpoint". In: *22nd International Symposium on Reliable Distributed Systems, 2003. Proceedings.* 2003, pp. 25–34. 10.1109/RELDIS.2003.1238052.

[2] Pastor-Satorras, R., Castellano, C., Van Mieghem, P., and Vespignani, A. "Epidemic processes in complex networks". *Reviews of Modern Physics* 87.3 (2015), p. 925. 10.1103/RevModPhys.87.925.

[3] Cator, E. and Van Mieghem, P. "Second-order mean-field susceptible-infected-susceptible epidemic threshold". *Physical review E* 85.5 (2012), p. 056111. 10.1103/PhysRevE.85.056111.

[4] Mata, A. S. and Ferreira, S. C. "Pair quenched mean-field theory for the susceptible-infected-susceptible model on complex networks". *EPL (Europhysics Letters)* 103.4 (2013), p. 48003. 10.1209/0295-5075/103/48003.

[5] Tonolo, T., Angelini, M. C., Azaele, S., Maritan, A., and Gradenigo, G. *Generalized Lotka-Volterra model with sparse interactions: non-Gaussian effects and topological multiple-equilibria phase*. 2025. arXiv: 2503.20887 [cond-mat.stat-mech].

[6] Roy, F., Biroli, G., Bunin, G., and Cammarota, C. "Numerical implementation of dynamical mean field theory for disordered systems: application to the Lotka–Volterra model of ecosystems". *Journal of Physics A: Mathematical and Theoretical* 52.48 (2019), p. 484001. 10.1088/1751-8121/ab1f32.

[7] Gradshteyn, I. and Ryzhik, I. *Table of Integrals, Series, and Products*. Elsevier, 7th Edition, 2007. ISBN: 9780123736376.

[8] Machado, D., Valigi, P., Tonolo, T., and Angelini, M. C. GenLotkaVolterra_SparseGraphs: GitHub Repository for this article. 2025. 10.5281/zenodo.17612770.